CORVUS

Whitepaper

Analyzing Risk Aggregation in Cyber Insurance

Author's Note

Through Corvus's partnerships with reinsurers and program managers, we have occasion to hear the thoughts, concerns and questions of some of the most experienced risk professionals in the world. It's through these conversations that we have refined our own methodologies around gathering and presenting risk data. In this paper, we will share the rationale and process behind analyzing risk aggregation in cyber insurance, including adaptations to the rise of ransomware and the hardening of the cyber market. We hope you find it informative.

Philip Edmundson
Founder and Chief Executive Officer

Madhu Tadikonda
President

Lori Bailey
Chief Insurance Officer

Table of Contents

Cyber Risk Aggregations: Mind the Gap

The management of risk aggregations is important in all lines of commercial insurance with catastrophic exposures. It's arguably job number one for the stewards of the industry who control hundreds of billions of dollars in reinsurance capacity, and whose decisions can mean the difference between the continuing functionality of a market for insurance following a severe catastrophe, and total breakdown. No small matter.

The need to understand and manage aggregation is no less acute for cyber liability than it is for any other line. Recently, rapid growth in cyber's market size, with double-digit increases in direct written premium year after year, has elevated the issue to the fore. Reinsurers now have measurable skin in the game, and with 40% of cyber risk ceded to the reinsurance market in 2020, the "dynamics of the cyber insurance market are intertwined with the behaviours and capacity of its reinsurance partners" (Berenberg). Reinsurers' ability to understand risk aggregations will define the next several years of the cyber insurance market.

The key problem — one that's a top-down driver of the persistent hard market in cyber — is that cyber risk aggregations do not play by the old rules. As we've heard repeatedly over the past few years, accurately gauging cyber risk is a serious challenge for all parties in the insurance ecosystem, one that's reflected in wide variation in loss ratios. And if it's difficult to understand risk on the account level, measuring aggregations of cyber risk across large books of business is commensurately tough. That is thanks to several exacerbating factors:

- A relatively short history of claims on which to build actuarial models
- Rapid technological evolution in both cyber threats and the defenses against them
- Innate connectivity between software products, leading to the potential for massive-scale impact from a single event that is unconstrained by geography
- A lack of discernible patterns in events: no two cyberattacks are exactly alike, making past losses a poorer predictor of future losses

These factors change the calculus for reinsurers and program managers — or rather, they should.

In practice, many program managers and reinsurers have approached the issue by attempting to apply catastrophe modeling logic that is heavily dependent on historical data. And while there have been some CAT-like events in cyber, including 2017's NotPetya ransomware and 2021's Microsoft Exchange Server vulnerability, these can be counted as near misses relative to the impact of a direct-hit "cyber hurricane" of the kind envisioned by some industry experts.

So while there have been enough large loss events in the industry to spook the market into severe hardening, there's a sense that basing risk models on past losses may still be underestimating the risk.

This gap — between what historical claims suggest aggregated risks to be, and what keen observers know is plausible in a true catastrophe — represents the primary challenge for the cyber industry as well as the greatest opportunity within it. Those who are able to move forward with confidence in managing cyber aggregations now will capture the upside of high rates and strong demand in the near term, and thereby build market share for the long term. Those who pull back out of a fear of the unknown will be left playing catch-up.

The answer to finding that confidence relies on an approach that may be more familiar than not, despite the apparent novelty of cyber risk. Reinsurers and carriers who have dealt with worsening windstorm threats in places like coastal Florida provide a blueprint for how to bring together numerous disparate factors into models that account for a great deal of uncertainty, where the future may not look exactly like the past. Standing on the shoulders of this work in property CAT modeling can help holders of risk capital in cyber short-cut the process to getting the answers they need. And it's all enabled by better collection, analysis and access of data, an area in which insurtech MGAs excel.

In this paper, we'll explore what kinds of cyber data is available to use in analysis, how it can be analyzed with parallels to property insurance modeling, and how MGAs enable the flow of information with partners, including a review of Corvus's approach.

Cyber CAT & The Property Precedent

Cyber risk may seem like a horse of a different color, but in fact the management of cyber risk aggregation follows the same storyline as catastrophe risk aggregations for natural disasters like earthquakes, floods and hurricanes. We look at those events in probabilistic frames such as 1 in 100 years or 1 in 500 years, and we can do the same for cyber events. Here we'll walk through an outline of the modeling method, noting the parallels to a proven CAT modeling application: quantifying the risk of windstorm insurance to properties in hurricane-prone areas.

As with property insurance, the primary goal is to avoid excessive aggregations of risk. There's no question as to whether a cyber book of business will experience some large losses in the future — the goal is to avoid outsized losses relative to peers in the industry. This means examining the book of business closely through a number of different lenses and forms of analysis, which fall into two categories: deterministic and probabilistic.

Setting the Table

We begin with a deterministic approach: if X happens, then it will cost Y. While it's conceptually simple, this piece is complex in practice; and when done for cyber risks, it's only possible with a rich set of data (more on that in the following section).

First, we must document exactly what is insured, what is not, and what limits are offered for each piece of coverage, to narrow down the universe of risk. This is easy enough, assuming the book of business has been properly segmented. Then comes analyzing what is happening within that insured universe — and here's where we can borrow directly from the experience of modeling windstorm risk.

For windstorm risk, it all starts with geography. Winds from hurricanes and tornadoes do by far the most damage within a relatively narrow storm track. If a book of business is over-indexed in an area with measurable risk to windstorm damage, this has obvious implications to overall aggregation of risk. That's why geography has always been, and remains, the table stakes for windstorm risk assessment.

Think of industry class designation, a staple of cyber underwriting, as the equivalent to this rather basic geographic problem. We have enough claims history in cyber insurance to know the hazard levels of each industry class, and seek to avoid over-concentration in any one industry — especially high-hazard industries.

For instance, we know that managed service providers (MSPs) are hazardous because an issue can spread quickly from their own systems to their customers', leading to the potential for larger third-party risk. As we saw with the Kaseya incident in July 2021, which involved software used by MSPs, popular software products within an industry can contribute to aggregated losses. Trends in which threat actors copy others' successful exploits of certain types of organizations is another way that a single industry can suddenly see a spate of attack activity.

Analyzing a book of business with an eye toward careful treatment of certain industry classes like MSPs is similar to how property insurers would avoid over-insuring a particular zip code in South Florida: a basic measure to limit the impact a direct hit would have on your overall book of business. Table stakes.

The Finer Points

In property insurance, the next level down from geography is to look at the details of how buildings are built — roof design, building materials, age. When we look at a cyber book of business we take a similar approach, examining the kinds of technologies present in insured IT systems to gain a deeper understanding of where losses come from, and the unusual places they might accrue within a book of business.

- Just as some buildings, such as those made of masonry or concrete, are less vulnerable to windstorms, IT systems that are built on infrastructure that is inherently more resilient — think modern cloud systems that introduce redundancy and backups by default — are a safer bet to withstand a cyber incident with less severe losses.
- Just as newer buildings are more resilient, because they've been built to newer building code regulations, so, too, are IT systems: it's risky to harbor lots of old, legacy software that may be neglected by its original creators (whereas patches are issued more frequently for newer software) or simply becomes unmanageable because of its age and complexity.
- Risk is even wrapped up into design considerations: certain types of roofline designs handle high winds better than others. Likewise it's possible to architect an IT system in such a way that lateral movement by attackers within the system is significantly hindered, reducing attack severity. "Security by design" principles, such as least privilege access, articulate this kind of thinking.

With just these few examples, you can get a sense of the array of dimensions through which we can view a book of cyber or tech E&O business — but also how much information is required to do it effectively. Being able to quantify and make judgements about something like the types of software present within the insured population involves gathering reams of data on a regular basis, and that requires technological solutions, which we'll discuss next.

Primer: The Cyber Data Smorgasbord

To understand how new forms of data can resolve questions around cyber aggregation, we have to know what's on the buffet table.

Cyber data includes claims history and firmographic data, like we see in other lines, but also significant pre- and post-claim data that can be gathered with automated, non-permissioned scans. On top of these forms of data MGAs and carriers can layer on data gathered throughout the policy period, either through digital questionnaires or third-party security tools. The result is a corpus that far exceeds what any static application could provide.

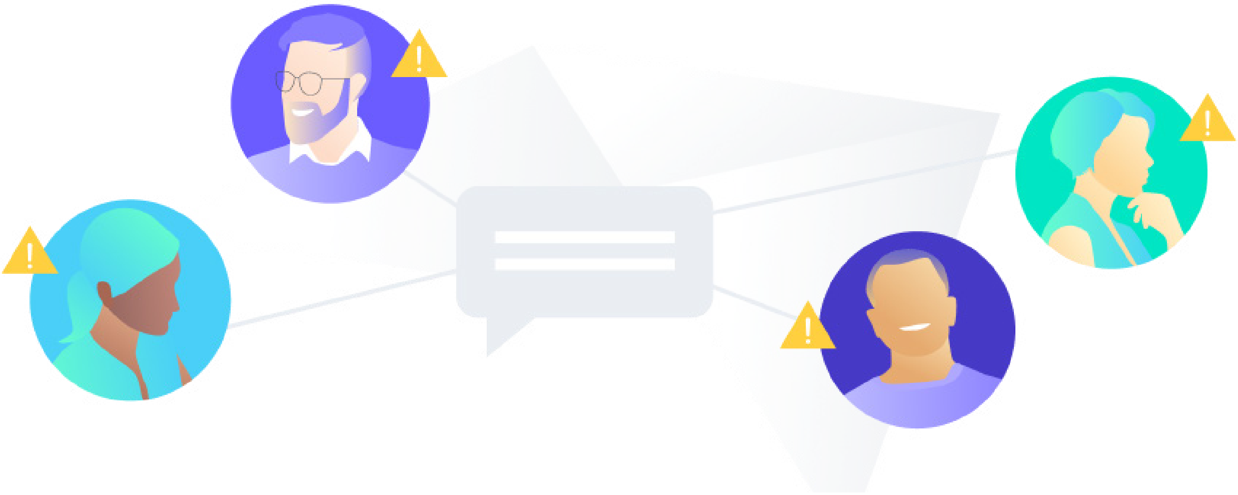

Cyber Aggregation Metrics

Here are a few examples of the kinds of data that can be used to determine aggregations, all of which can be gathered through a sophisticated external scan:

Email service provider / email security tools

As the most common digital communication tool, business email is a locus for attempted attacks and the efforts to prevent them. Phishing and other social engineering tactics center on email communications, and the sheer number of credentials that exist provides good odds for attackers seeking to brute-force credentials or stuff them with known passwords.

Knowing what email providers — and what versions of those email providers — are used within a book of business provides insights into overall risk relative to benchmarks. Particularly looking at changes over time, for instance, Corvus found that the use of the least-risky email providers and tools increased measurably as ransomware rose. On the aggregation side, it's possible to measure the impact of a zero-day vulnerability that affected a major email provider, or certain versions of that email provider's software, for example.

Industry-specific Software

The hazard levels of various industries are, of course, one of the key underwriting criteria used in cyber as elsewhere in commercial insurance. These are determined based on prior claims, and with the increase in overall cyber claims across the industry over the past few years the data is becoming more refined.

But when considering aggregated risk, we don't need to stop there. Before any major claims trends appear in the data, we can marry experts' insight about what kinds of industry-specific software may be popular in a given industry and analyze aggregation points that exist for that industry as a result — and in turn, look at how an event in that industry would impact the overall book of business.

The Enabler: External IT Scans

Initially viewed skeptically, then accepted, and now practically required, external IT scans are one of the most powerful tools for gathering data to drive cyber risk analysis. They work by scanning the externally-facing facets of any IT system, gathering information about the kinds of software being used to communicate with the outside world, such as web applications and email, and the infrastructure used to host it all. A scan equates to the kind of visibility threat actors have as they survey the web for easy marks.

While the readouts from these scans may be less comprehensive than security professionals (and underwriters) might prefer to see on an individual account they are working, this approach's advantages are evident at larger scales. The fact that scans are easy to execute, can be done repeatedly at regular intervals, and are non-intrusive (and thus don't require permission), means that their use for large-scale data science is incomparable. They can gather data about policyholders, but also non-policyholders (to form a control sample), and establish time-series data to show trends.

That means a program manager or reinsurer can learn a lot about their book of business, and how the book compares to the broader population of organizations.

Cloud Service/Web Hosting Provider

While there is a long tail of fragmentation in web hosting providers, there's also a large pareto-style congregation at the top with very few providers holding a massive share of the industry. For obvious reasons, hosting providers are fundamental to the basic operations of millions of organizations. Assigning probabilities to scenarios when quantifying aggregated risk requires understanding exactly how much of a book of business will be affected by scenarios that impact customers of giants such as Amazon Web Services.

Virtual Private Networks (VPNs)

As remote and hybrid work expanded during the pandemic, so did the usage of VPNs. As this category continues to grow in popularity, understanding how it factors into risk aggregations will be increasingly important to overall risk assessments. In late 2021 VPN software from Palo Alto Networks was the subject of a zero-date vulnerability, underscoring that this kind of software can become a target, even if installed properly and kept up to date.

Patch management practices

Rather than looking specifically at a piece of software's share within a book of business, we can also measure how segments within the book deal with "hygiene" practices like patch management. By measuring proportions of software that is out of date — and how far out of date — across a "basket" of common pieces of software, we can learn what percentage of the book is likely to be vulnerable to a future attack that exploits older versions of a particular software product, or conversely how quickly the book of business is likely to rectify a zero-day after it's announced.

Corpus Maximus

Where information cannot be gained by the scan, it can be requested in a way that automatically builds it into the existing database, enriching the existing foundation data. Corvus has proven this model by gathering information directly from policyholders through a digital dashboard, offering it as a path to getting bespoke recommendations for improving security and discounted services. Since launch, uptake among new policyholders has grown significantly.

All of this — regular scan data and survey data — layers on top of firmographic data, encompassing a large corpus of data for measuring cyber aggregations. But the data is only half the battle. Next is putting it into action by assigning probabilities to loss events.

The Probability Layer

We've covered the simple deterministic portion of measuring aggregations, as well as some of the many sources of data that can fuel those quantifications. Next is the probabilistic layer — defining the hundred-year storm scenario versus the five hundred year storm scenario.

How likely is a catastrophic supply chain cyberattack to occur in the next year? The next five?

This is where, again, we can borrow from sophisticated modeling techniques. In the realm of windstorm risk, firms can lay a carrier's book of business in the path of a hypothetical storm, calculating precise loss estimates for storms of varying severity and that might track just a few miles in one direction or another. While there's no equivalent visual representation in cyber, the approaches are similar.

Like in property modeling, the method used by specialists in the space like CyberCube (with whom Corvus has a partnership), involves building many plausible scenarios to cover the full range of possibilities, then assigning probabilities within that range.

> For cyber risks, history is not a predictor of the future in terms of modeling as threat actors and the methods they deploy are constantly changing."
>
> CyberCube Report: Drawing from the Experience of Nat Cat Modeling

Tomorrow's Catastrophe

As CyberCube readily admits, "the lack of a formally identified historical catastrophic insured cyber event along with the associated claims data showcases the difficulties in modeling [frequency and severity]. Event duration and correlation are not as easily defined as the duration of a hurricane passage and subsequent flooding, for instance." But there are a number of near-catastrophic events that have occurred in the recent past — hacks that looked like they could become "the one", like SolarWinds — and in this case the more recent that the near-miss event took place, the higher the relevance of the data which can be used to inform models.

Because of the dynamism in the environment, with changing preferences around IT products and services and trends in which attack styles are most lucrative, the most likely way a catastrophic event will unfold in the next year isn't necessarily well modeled by hacks that occurred in 2006, or even 2020. This contrasts with modeling for nat cats, where for instance floods that occurred in those years could be an indicator of future trends.

Instead, seeing how attacks unfolded in real life, with data from as recently as this month, provides a better starting point. Scenarios that envision events that look considerably like recent ones, but with a few details changed — what CyberCube calls "exaggeration of the present" — may be more informative than building on what a "typical" event has looked like over the longest time period possible, even if that means having fewer total data points. It does not take a significant stretch of the imagination to reconceive of a given event, with minor changes creating far more significant consequences. Where limited historical events exist for certain types of cyber catastrophe, analogous events are used to examine the potential impact of such an event. For example, an electricity power outage could illustrate the potential impact of a cloud outage on a given portfolio of companies.

In addition to the approach of using recent historical data to identify trends and running exercises to find out whether a repeat, or something similar, is plausible, CyberCube has a few other lenses through which it views scenarios. These include reviewing how new technologies are evolving, a skewed distribution analysis that focuses on the targets most likely to result in catastrophic loss, and changes in large-scale crime motivations or current tools and tactics used.

While there are still unknowns in modeling, as there always are, the sheer volume of data and the ability to integrate new information rapidly are factors helping to make sensible estimations about cyber risk aggregations. This means adjustments can be made quickly by program managers and reinsurers to limit exposure as things change and are continually reassessed. That is where the relationship with leading-edge MGAs is beneficial.

MGAs: Ushering in a new paradigm in data transparency

MGAs have always offered opportunities for program managers and reinsurers to enter specialized markets, and with its newness and complexity it's no surprise that Cyber has proven to be a fertile area for the model. But while the majority of program managers still view market access as the most desirable trait for an MGA (Clyde study), it's increasingly clear that there's more for MGAs to offer when it comes to technology-enabled approaches to sharing data.

While it may not have been initially expected that Cyber MGAs would provide their program partners with vastly different data than their forebears in other lines of insurance, the variety and immediacy of cyber data represent a step-change in data's utility for reinsurers. With claims data that develops quickly as quickly as cyber's does, rapid action can be taken to mitigate risk — if provided with sufficient context to understand what those claims may mean moving forward.

As the breeding ground for new underwriting methodology in cyber, MGAs are positioned to not only deliver on their original promise of providing access to difficult/niche markets, but also extend their value beyond sales/underwriting to being a new kind of data provider for their program managers, in delivering real time data that enriches the understanding of aggregation.

The Corvus Approach

While a number of insurtech MGAs have pursued technology-led approaches to underwriting and gathering data, Corvus is unique in its "productized" delivery mechanism of data to our risk capital partners and the breadth of data shared.

Corvus combines our comprehensive and proprietary cyber claims database with complete measurement of key points of cyber risk aggregation, along with conventional firmographic metrics (industry, state, limits of insurance, risk hazard classifications). This corpus of data is packaged and delivered through an online dashboard purpose-built for managing cyber risk aggregation. The Risk Aggregation Platform^{SM} is updated in real time, inclusive of even emerging threats that have not yet resulted in a claim with a Corvus policyholder.

In addition to making this richly segmented data transparent and easy for our partners to analyze, we've taken an industry-first step toward making this data actionable. Within the platform, users can alter the caps on specific combinations of risk factors within the book of business to respond to changing levels of risk, for instance from a newly-discovered vulnerability in a certain type of software, or notable trends in attack patterns. This is available to the risk capital partners of Corvus in real-time allowing for active management of aggregations during the contract year of a program.

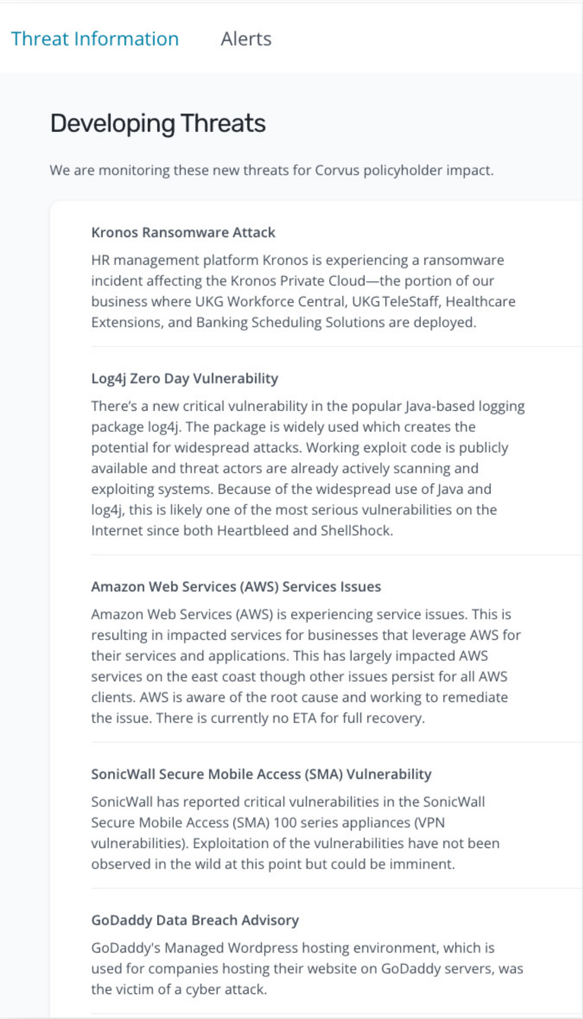

Developing Threats

We are monitoring these new threats for Corvus policyholder impact.

Kronos Ransomware Attack

HR management platform Kronos is experiencing a ransomware incident affecting the Kronos Private Cloud—the portion of our business where UKG Workforce Central, UKG TeleStaff, Healthcare Extensions, and Banking Scheduling Solutions are deployed.

Log4j Zero Day Vulnerability

There's a new critical vulnerability in the popular Java-based logging package log4j. The package is widely used which creates the potential for widespread attacks. Working exploit code is publicly available and threat actors are already actively scanning and exploiting systems. Because of the widespread use of Java and log4j, this is likely one of the most serious vulnerabilities on the Internet since both Heartbleed and ShellShock.

Amazon Web Services (AWS) Services Issues

Amazon Web Services (AWS) is experiencing service issues. This is resulting in impacted services for businesses that leverage AWS for their services and applications. This has largely impacted AWS services on the east coast though other issues persist for all AWS clients. AWS is aware of the root cause and working to remediate the issue. There is currently no ETA for full recovery.

SonicWall Secure Mobile Access (SMA) Vulnerability

SonicWall has reported critical vulnerabilities in the SonicWall Secure Mobile Access (SMA) 100 series appliances (VPN vulnerabilities). Exploitation of the vulnerabilities have not been observed in the wild at this point but could be imminent.

GoDaddy Data Breach Advisory

GoDaddy's Managed Wordpress hosting environment, which is used for companies hosting their website on GoDaddy servers, was the victim of a cyber attack.

For example, we analyze Cloud Service Providers, not only at the aggregate level but also at the operational level. That means that we can measure and limit exposure for key regional hubs for AWS, Google Cloud and Microsoft Azure. Comparing this to Florida windstorm aggregation, this analysis allows us to be certain that we are not insuring all of our Florida windstorm in one zip code.

Other key variables like Virtual Private Network (VPN) tools are also a source of potential risk aggregation. Corvus both measures the presence of the various brands of VPN tools and ranks them qualitatively. Combined, these inputs lead us to various strategies for managing aggregate VPN exposure in our overall book of business. Furthering the Florida windstorm analogy, this feature is like managing construction types of buildings, all with a strong opinion about the resistance of various types of construction to windstorms.

The potential to manage Cyber Risk using Property Catastrophe modeling techniques is significant. We also measure eCommerce vendors, web-hosting services, FTP servers, and other variables. Taken together, this can provide more efficient use of reinsurance capital with more transparency and visibility. It may lead, as it has for Property Insurance, to the development of Insurance Linked Securities.

Adding to our in-house efforts to aggregate and illuminate as much data as possible for our partners, Corvus has added modeling capabilities by following the P&C industry's best practices for managing Windstorm and Earthquake catastrophic risk management. Just as Property insurers work with aggregation modeling companies like AIR Worldwide and RMS, Corvus manages Cyber risk aggregation in concert with the catastrophic loss scenarios of CyberCube, the recognized leader in the space. This combination will be critical for successful Insurance Linked Securities.

"CyberCube recognize the power of differentiated real-time data collection in improving the transparency of cyber risk to capital partners. Partnering with Corvus enables us to support the improved understanding of complex cyber catastrophe modeling, and increase confidence in the market" says Oliver Brew, the company's Head of Client Success.

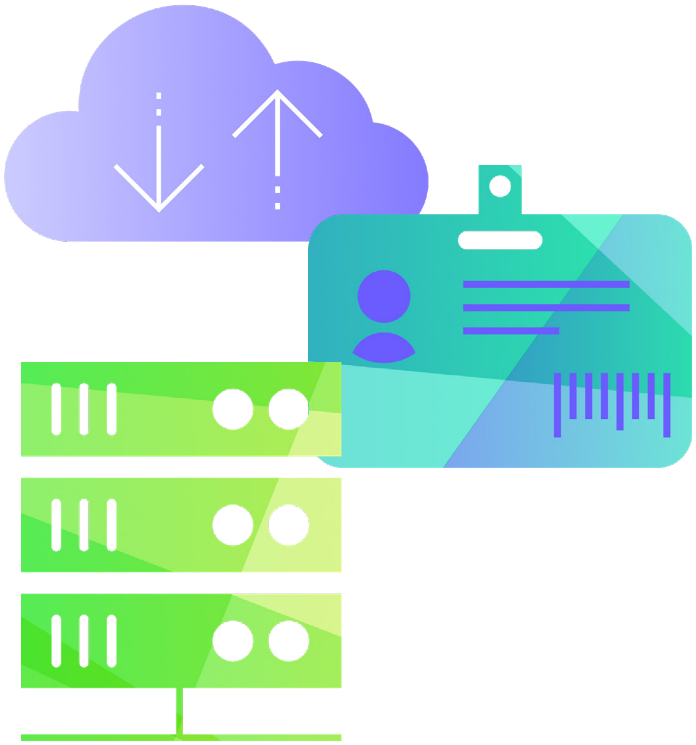

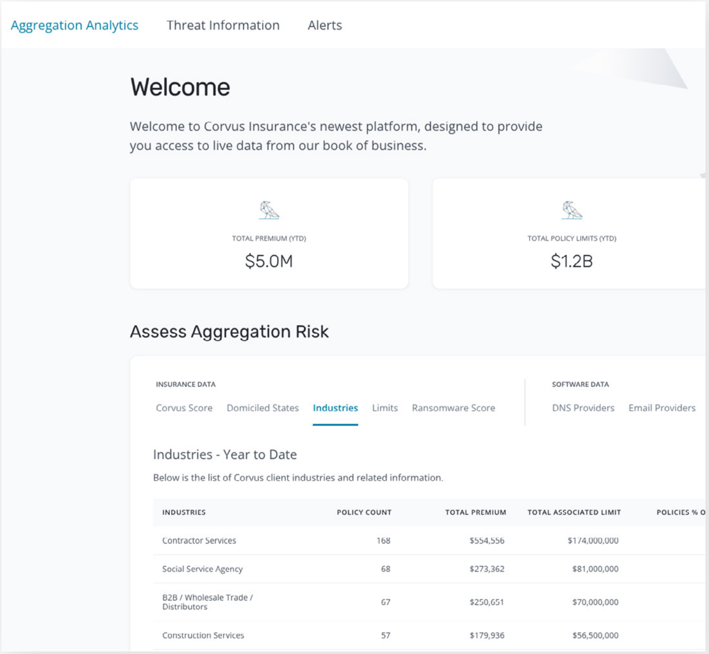

Welcome

Welcome to Corvus Insurance's newest platform, designed to provide you access to live data from our book of business.

TOTAL PREMIUM (YTD)
$5.0M

TOTAL POLICY LIMITS (YTD)
$1.2B

Assess Aggregation Risk

Industries - Year to Date

Below is the list of Corvus client industries and related information.

INDUSTRIES	POLICY COUNT	TOTAL PREMIUM	TOTAL ASSOCIATED LIMIT	POLICIES % O
Contractor Services	168	$554,556	$174,000,000	
Social Service Agency	68	$273,362	$81,000,000	
B2B / Wholesale Trade / Distributors	67	$250,651	$70,000,000	
Construction Services	57	$179,936	$56,500,000	

Cyber's Questions; Property's Solutions

There's much that's unprecedented about cyber risk, and much has been written (rightly so) about the challenges it poses for insurers, program managers, and reinsurers. But we believe that this territory is better charted than many believe. Windstorm modeling proves that with enough motivation and enough data, it's possible to build models even for notoriously unwieldy and changeable risk factors. Within cyber, we look at the volume of unusually detailed data about books of cyber business; the pace at which new, real-world data is ingested into CAT models; and the ability of program managers to review and take action on new information in real time — and see opportunity. These factors all point to cyber risk being not just modelable, but manageable.

CORVUS

About Corvus

Corvus Insurance is building a safer world through insurance products and digital tools that reduce risk, increase transparency, and improve resilience for policyholders and program partners. Corvus's market-leading specialty insurance products are enabled by advanced data science and include Smart Cyber Insurance®, Smart Tech E&O℠, Smart Cargo®, and a suite of products for Financial Institutions. Our digital platforms and tools enable efficient quoting and binding and proactive risk mitigation.

Corvus and its subsidiaries offer insurance products in the US, Middle East, Europe, Canada, and Australia. Current insurance program partners include AXIS Capital, Crum & Forster, Hudson Insurance Group, certain underwriters at Lloyd's of London, R&Q's Accredited, SiriusPoint, and Skyward Specialty Insurance. Corvus Insurance was founded in 2017 and is headquartered in Boston, Massachusetts with offices across the US and in London, UK. For more information, visit corvusinsurance.com.

CyberCube

About CyberCube

CyberCube is on a mission to deliver the world's leading analytics to quantify cyber risk.

We help the cyber insurance market grow profitably using our world-leading cyber risk analytics and products. The combined power of our unique data, multi-disciplinary analytics and cloud-based technology helps with insurance placement, underwriting selection and portfolio management and optimization. Our deep bench strength of experts from data science, security, threat intelligence, actuarial science, software engineering, and insurance helps the global insurance industry by selecting the best sources of data and curating it into datasets to identify trustworthy early indicators of risks and to build forward-looking views of them. www.cybcube.com