KOVRR
Cyber Decisions. Financially Quantified.

Cyber Portfolio Analysis: Translating Control Gaps Into Financial Exposure

An Investment Perspective for Private Equity

www.kovrr.com

John Zuska, CISSP —— In Partnership with Kovrr

About the Authors

John Zuska, CISSP, is a cybersecurity executive with over 25 years of experience advising enterprise clients across regulated industries, including finance, healthcare, and energy. He served as Chief Information Security Officer at Z Capital Group, a US-based private equity firm, where he held security responsibility for both the firm and its full portfolio of companies, spanning hospitality, industrial manufacturing, gaming, and consumer brands, including Mrs. Fields Cookies and TCBY Yogurt. In that role, Mr. Zuska was among the early practitioners to apply financial risk modeling to cybersecurity due diligence in a PE context, translating control gaps into quantified financial exposure that investment teams could incorporate into deal analysis.

Prior to his work in private equity, Mr. Zuska held cybersecurity and systems administration roles in financial services, including positions at Citigroup and Bank of America. He subsequently built an extensive consulting career serving as a Virtual CISO and strategic advisor to organizations across regulated industries. He holds the CISSP designation alongside certifications including PMP, JBoss, and Red Hat, and is a frequent speaker at industry events, and serves as a board advisor for several organizations.

Kovrr is a leading provider of AI and cyber risk governance and control solutions, specializing in financial risk quantification. Led by CEO and Co-founder Yakir Golan, Kovrr's platform empowers CISOs, GRC leaders, portfolio performance teams, and finance leaders to translate complex cybersecurity and AI-related data into financially quantified insights, enabling smarter decisions around budgeting, insurance, regulatory compliance, and strategic investment.

Golan began his career in the Israeli intelligence forces and later built multidisciplinary expertise across software engineering, product development, and enterprise risk strategy. Over the past decade, he has worked closely with CISOs, risk executives, and boards to strengthen how organizations govern cyber and AI-related risk at scale.

Built on insurance-grade models continuously validated against data from millions of global companies, Kovrr's patented financial risk quantification technology engine helps organizations adopt a Shift Up approach, elevating cyber and AI risk considerations into the boardroom, moving beyond qualitative assessments, and managing digital exposure with the same rigor applied to any other enterprise risk.

Cyber Due Diligence in Private Equity

When a private equity (PE) firm acquires a company, it is not only buying revenue streams and assets, but also inheriting the organization's cyber risk. Every unpatched vulnerability and misconfigured access control becomes the PE's problem the moment the deal closes. Rather than being mere annoyances that stakeholders have to deal with later, however, these issues can translate into material financial events for investors. A breach or supply chain compromise caused by one of these shortcomings can disrupt operations and create reputational damage that affects both valuation and exit consequences.

Consequently, cyber due diligence has emerged as a vital part of the investment process. During acquisitions, chief information security officers (CISOs), technology leaders, or risk managers at the PE firm will now evaluate the target company's cybersecurity risk, exploring key factors including controls maturity, vulnerabilities, regulatory compliance, data protection practices, incident response readiness, and many more. These assessments often reference established cybersecurity frameworks such as the NIST Cybersecurity Framework (CSF) or ISO 27001, which provide a structured way for these stakeholders to identify potential weaknesses before the deal closes.

For investment teams, the outcome of this work is to help mitigate post-acquisition risk and illuminate opportunities to improve operational resilience. When conducted effectively, it protects investor returns and ensures that cyber exposure is reflected appropriately in deal negotiations. When performed superficially, though, the process often fails to deliver the type of insight necessary for optimizing terms and conditions. In order to be useful, cybersecurity findings must be connected to the valuation models and investment analyses that ultimately drive deal decisions.

The Limitations of Traditional Cyber Risk Evaluation

For years, cybersecurity assessments conducted across the PE space and other industries alike have relied on traditional evaluation methods. The most common approach has been some variation of a red-yellow-green (RYG) rating, often paired with a likelihood-versus-impact matrix. Controls are classified as high, medium, or low risk, while threats are plotted on a grid and compiled into a report. Although these models can help security teams prioritize issues internally, they fail to provide portfolio managers with the information needed to make more informed investment decisions.

As Mr. Zuska observed firsthand, both RYG ratings and likelihood-versus-impact matrices are fundamentally qualitative. They categorize risks but do not quantify the potential business impact. A "yellow" rating, for example, might represent a $200,000 remediation cost, or it might conceal a $5 million average breach exposure. The label does not communicate which, leaving stakeholders with more questions than answers. Moreover, this subjective scoring introduces inconsistent evaluations across portfolio companies, making it difficult to integrate cybersecurity findings into financial valuation models.

> "A technology person understands what a high-risk rating means. But if you give that to a portfolio manager at a PE firm, they look at it and say — okay, great. What the heck does that mean to me?"
>
> **John Zuska, CISSP**, Former CISO, Z Capital Group

Modeling Cyber Risk as Financial Exposure

Accounting for this issue has the fairly straightforward solution of shifting from these dated, qualitative classifications to financial risk modeling. Instead of determining whether a risk should be labeled as high, medium, or low, CISOs, with cyber risk quantification (CRQ) and AIRQ models, can forecast the probability and severity of potential incidents, allowing them to translate control gaps into measurable financial exposure that partners can then incorporate directly into investment analyses.

One of the primary methods of measurement used in the modeling process is the Monte Carlo simulation, a statistical approach that models tens of thousands of potential cyber incident scenarios to produce a probability distribution of outcomes, or a loss exceedance curve (LEC), from which the average annual loss and extreme tail-risk scenarios can be extracted. With these data points, it becomes much easier to visualize the tangible business implications of acquiring the new portfolio company, ensuring deals can be priced accordingly.

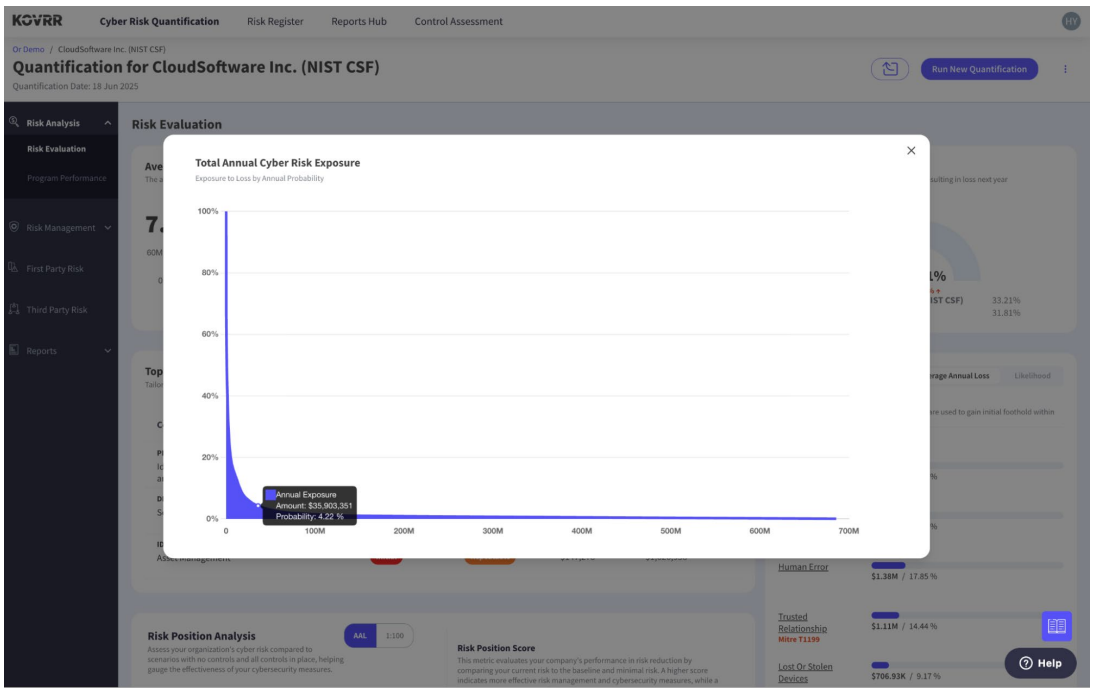

***Figure 1**. Example loss exceedance curve generated through Kovrr's cyber risk quantification platform, illustrating modeled cyber loss distributions.*

As with most statistical modeling techniques, CRQ draws on a combination of real-world data and environment-specific information. Cyber risk models harness industry benchmarks such as IBM's Cost of a Data Breach Report and Verizon's Data Breach Investigations Report to generate sector-specific loss figures and breach patterns. For example, in 2025, healthcare breaches averaged more than $7 million per incident, while financial services breaches typically amounted to $5.5 million.

While these industry reports provide important baseline assumptions, they represent only a starting point for financial cyber risk modeling. Public breach reports capture broad trends, but they rarely reflect the full financial consequences of cyber incidents or the specific conditions within an individual organization. A breach affecting a global hospital network, for example, carries very different operational and financial implications than one affecting a small healthcare provider, even if both technically fall within the same sector statistics.

To produce more accurate estimates, the modeling process must incorporate additional layers of data that reflect how cyber events actually unfold in practice, including variables such as the organization's control maturity, the scale and sensitivity of its data environment, its reliance on third-party vendors, and the regulatory frameworks governing its operations. These factors heavily influence both the likelihood of an attack occurring and the magnitude of the financial and operational impact if one does materialize.

Increasingly, this customization layer is where specialized CRQ platforms have become particularly valuable. Solutions like the one from Kovrr incorporate vast cyber event datasets alongside large volumes of insurance claims data, providing a far more actuarial perspective on cyber loss patterns. In fact, Kovrr is uniquely positioned in this regard, with access to one of the industry's largest collections of cyber incident and insurance claims data used to calibrate its financial risk models. Because insurance claims capture the full scope of incident costs, they offer a more comprehensive and accurate view of the economic consequences of cyber events than breach reports alone.

By integrating these datasets into probabilistic models, CRQ platforms can produce loss distributions that more closely reflect the financial reality of cyber risk. For PE investors, the result is a better understanding of how specific control gaps translate into potential financial exposure, allowing cybersecurity findings to be evaluated alongside other financial risks during the deal process. This quantitative perspective ensures cyber risk can be incorporated into the same financial frameworks used to assess other risks across the investment lifecycle.

> "Cyber risk quantification is a force multiplier. You take a PE firm CISO, and suddenly, with CRQ, they're producing the kind of intelligence that used to require hiring a full-time consultant at hundreds of thousands, if not millions, of dollars."
>
> **John Zuska, CISSP**, Former CISO, Z Capital Group

Portfolio-Level Cyber Risk Aggregation

Quantifying cyber risk at the individual company level provides a highly valuable frame of reference. Mr. Zuska's experience managing risk across an entire PE portfolio makes this point clear. Still, a greater strategic benefit for private equity firms emerges when those insights are evaluated across the entire portfolio. Given their inherent structural nature, PE firms rarely manage cyber exposure within a single organization. Instead, stakeholders are tasked with overseeing multiple portfolio companies simultaneously, which often vary across different industries, technology stacks, and regulatory environments.

With traditional cyber assessments, companies are evaluated in isolation and then given a red-yellow-green rating or a maturity score. However, those qualitative scores cannot be meaningfully aggregated. If one company is labeled “high risk” and another “medium risk,” there is no objective way to combine those results or measure the total exposure they represent. Portfolio managers cannot multiply or sum color-coded ratings in the same way they would financial metrics. Consequently, these assessments provide little visibility into how cyber risk behaves across the portfolio as a whole.

Cyber risk quantification changes this dynamic by enabling firms to model financial exposure across multiple portfolio companies simultaneously. Once cyber risk is expressed in monetary terms, it becomes possible to identify concentrations of risk that may not be visible when companies are evaluated individually. This view is particularly important when portfolio companies rely on common technologies, cloud providers, or third-party vendors.

In these situations, a single vulnerability or cyber event can propagate across several portfolio companies at once, creating systemic exposure for the PE firm. A compromised software supplier, for example, may simultaneously affect multiple portfolio companies that rely on the same platform. Quantified risk models make it possible to simulate these cascading scenarios and estimate the potential financial impact across the portfolio.

This perspective also supports the adoption of standardized cybersecurity expectations across portfolio companies. Many PE firms now establish baseline security requirements during the integration process. When combined with quantified risk modeling, these initiatives allow firms to measure how improvements in cybersecurity posture reduce financial exposure across the entire portfolio over time. With the portfolio-wide lens, PE firms gain the ability to identify shared vulnerabilities and manage cyber exposure with the same discipline applied to financial risk across their investments.

How Quantified Exposure Shapes Insurance Strategy

Once cyber risk is quantified and aggregated across the portfolio, it can begin to influence the cyber insurance strategy. In practice, the question is rarely whether cyber insurance should be purchased, because for most mid-market companies, the alternative is effectively self-insurance. The financial burden of a cyber incident escalates quickly. Forensic investigations alone can amount to hundreds of thousands of dollars, and after regulatory response, customer notification, legal exposure, and operational downtime are added

to the equation, the total cost can run into the millions before a company has recovered a single dollar.

This reality explains why cyber insurance has become a standard risk management tool across private equity portfolios. But buying insurance and buying the right insurance are two separate things. Generic policies sized to generic assumptions leave coverage gaps in exactly the places a quantified assessment would have flagged. Dollar-based risk metrics, conversely, give insurers credible, specific data, and when an organization's potential losses are expressed as probabilistic financial outcomes, investment teams can evaluate insurance policies with a greater degree of accuracy.

Expected annual loss estimates and tail-risk scenarios provide a clearer framework for determining whether existing policies offer sufficient coverage or whether additional protection is required. This data also improves the dialogue between portfolio companies and insurers. Organizations that can demonstrate a quantified understanding of their cyber risk profile and highlight the controls they've implemented to mitigate exposure earn greater confidence from carriers. In many cases, that transparency leads to more favorable premium pricing and broader coverage terms.

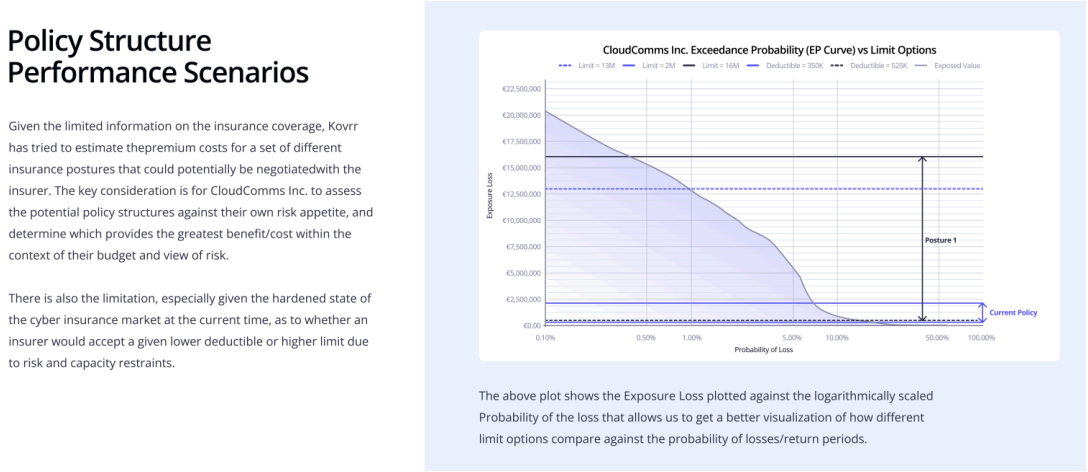

***Figure 2.** Cyber insurance stress testing using quantified risk scenarios. Advanced CRQ platforms generate reports like this to evaluate how simulated cyber losses interact with policy deductibles and coverage limits.*

The same logic extends to the portfolio level. When a PE firm can present aggregated, quantified cyber risk data across its portfolio companies, it fundamentally changes its position in negotiations with insurers. Carriers can price risk more accurately, creating the conditions for bulk insurance programs that reflect real-time risk rather than conservative assumptions, thereby resulting in materially lower premiums and more consistent coverage terms across the portfolio. For mid-market firms that have historically managed insurance company by company, this represents one of the more tangible financial benefits of moving to a quantified approach.

Mr. Zuska's View From the Field: A Transaction Walkthrough

The difference between qualitative and quantified cyber assessments becomes most apparent during transactions. Mr. Zuska encountered this dynamic directly during his tenure advising private equity cybersecurity strategy. In one transaction, the firm was evaluating the acquisition of a logistics software company. During the due diligence process, a standard cybersecurity assessment was conducted. The resulting report classified most control areas as "yellow," suggesting moderate risk but nothing severe enough to threaten the transaction. From a traditional security perspective, the company appeared to be in relatively typical condition for a mid-market technology provider.

However, when the findings were run through a financial cyber risk model, a more specific exposure began to emerge. Mr. Zuska's analysis identified weaknesses in third-party vendor access controls that could potentially allow attackers to move laterally through the company's environment. While this control gap appeared manageable within the qualitative assessment, the quantified analysis estimated that a supply chain–related cyber event could result in $5-7 million in potential financial losses, depending on the scope and duration of the incident.

Armed with those numbers, Mr. Zuska worked with the investment team to incorporate cyber exposure directly into the financial analysis of the deal. During negotiations, the firm leveraged the quantified insights and adjusted its offer downward by approximately 5-7%, a reduction that reflected the actual cost of remediation and contingent liability rather than a subjective risk classification. The deal still closed, but with terms that more accurately reflected the risk profile of the company.

Post-acquisition, Mr. Zuska led the prioritization of remediation efforts around vendor access management and oversaw the implementation of stronger controls for third-party authentication and monitoring. The company's cyber insurance program was also updated to better account for supply chain–related incidents. Within the first year of ownership, these measures reduced the company's modeled cyber exposure by roughly 40%. The improvements not only strengthened operational security but also simplified future diligence processes when the company was evaluated during later investment discussions.

In fact, the improved cybersecurity posture, documented and benchmarked against recognized frameworks such as NIST and CMMC, contributed to stronger buyer confidence and higher exit multiples. Buyers conducting their own diligence clearly saw what remediation steps had been taken and what the remaining risk profile looked like. That transparency carried real financial value. Companies that can demonstrate resilience and regulatory readiness often see enterprise value increase by 10-20% during exit negotiations. In this particular case, translating cyber risk into financial terms helped turn a potential liability into a measurable driver of value.

Cyber Risk as an Investment Variable

Cyber risk in private equity is a financial problem, requiring a financial framework to manage it effectively. Red-yellow-green assessments were a reasonable starting point when cyber was a technical concern at the edges of deal-making. However, considering the far more material role cybersecurity now plays in valuation, transaction negotiations, insurance strategy, and exit outcome, the qualitative approach no longer suffices. Translating cyber risks, from vulnerabilities and control gaps to operational disruptions and regulatory exposure, into measurable financial exposure allows investment teams to incorporate cyber risk into the same analytical frameworks used to evaluate every other dimension of a deal.

In Mr. Zuska's assessment, private equity firms that adopt this approach early are better positioned to negotiate sharper deals, build more resilient portfolios, and exit at stronger multiples. The visibility gained from cyber risk quantification platforms like the one from Kovrr allows investment teams to compare exposure across portfolio companies, identify systemic vulnerabilities, and prioritize remediation where it will have the greatest financial impact. Cyber risk, expressed in financial terms, stops being a liability to manage around and becomes a variable that can be actively optimized, like any other driver of investment value.