

20
26

IRIS XTREME

AN INFORMATION RISK INSIGHTS STUDY (IRIS) ON EXTREME CYBER EVENTS



119

Cyentia
INSTITUTE

Introduction

“Extreme ways are back again.” —Moby

In late 2020, we published the “Xtreme Edition” of the Information Risk Insights Study (IRIS), which examined the ~100 costliest security incidents of the preceding five years. The study arose from a desire to focus on the extreme tail of cyber losses, which is what tends to keep risk managers and execs up at night.

This edition of the IRIS Xtreme returns once again to the tail (and tale) of the cyber loss distribution. Specifically, we examine the top 5% of financial losses across all security incidents that occurred over the last five years. These 136 extreme cyber events together accounted for \$26.4 billion in documented losses.¹ The median extreme event now costs \$100M — double what it was in our last study. Going back further, the 95th percentile of losses has risen 19-fold since 2008. By any measure, the tail of large cyber losses is growing longer and fatter.

What follows is a rigorous examination of what drives these events, what they cost, who suffers them, and what can be done about them. We look at losses in absolute terms and relative to revenue, dissect the gap between what conventional cost-per-record benchmarks promise and what the data delivers, analyze five years of SEC cybersecurity disclosures, examine the secondary impacts that extend beyond the directly affected organization, and identify the threat patterns, actors, and techniques most consistently behind extreme losses.

A study of this scope is only possible with exceptional data. As always, we leverage Zywave’s expansive Cyber Loss Data. We’ve also partnered once again with Feedly for Threat Intelligence to synthesize the broader security community’s analysis of major events. We also have a newcomer to the IRIS family, Board Cybersecurity’s structured tracking of SEC 8-K filings. The combined dataset is further enriched with Cyentia’s own research, making this the most detailed public analysis of extreme cyber loss events published to date.

The refrain of Moby’s Extreme Ways repeats “Then it fell apart, it fell apart.” And that’s what often happens when organizations suffer extreme cyber events; things fall apart quickly. But the lessons drawn from such events are valuable in helping others to avoid that fate or at least keep it together. That’s our hope with this IRIS Xtreme.

TABLE OF CONTENTS

1

EXTREME
EVENTS IN
CONTEXT

4

2

THE IMPACT
OF EXTREME
EVENTS

14

3

THE THREAT
OF EXTREME
EVENTS

25

4

NOW
WHAT?

32

A

APPENDIX A

35

B

APPENDIX B

36



The Cyentia Institute is a data-driven research firm dedicated to making cyber risk more measurable and therefore more manageable. You can learn more about us and download all reports in the IRIS series at www.cyentia.com/iris.

¹All loss amounts in this study are presented in inflation-adjusted 2025 USD.

KEY FINDINGS



Extreme cyber events keep getting more extreme. The 95th percentile of financial losses — our threshold for an “extreme” cyber event — has risen 19x since 2008, from \$15M to \$281M. And yes, that’s adjusted for inflation.



Nearly 3 in 10 extreme events cost more than the business. Losses from incidents analyzed in this study span from a rounding error (0.03% to 3,000 times annual revenue) to several hundred thousand times annual revenue. We know that seems absurd and will show what’s behind those revenue ratios.



Most “material” cybersecurity incidents aren’t material. We examined 205 incidents filed after the 2023 SEC rule mandating disclosure of material cyber incidents. Only 6% are assessed as material by the filers; the rest are deemed immaterial, indeterminate, or still under evaluation. Furthermore, there’s little overlap between “material” and “extreme” events in this study.



Extreme events aren’t just expensive; they’re distinctive. Data exposure drives roughly 90% of all security incidents, but less than 40% of extreme events. Disruption surges to 43% — more than five times its share in the overall mix. And fraud climbs 27x. Are your defenses optimized for what happens most or what hurts most?



Supply chain disruption leads secondary impacts. Nearly half of extreme events (46%) generated impacts that reverberated through the victim’s supply chain. Cyentia’s research shows that the financial cost to downstream parties can be equal or greater than those incurred by the originating organization.



Third-party risk is a major factor in extreme events. Third parties directly caused just 2% of these incidents. Indirectly, however, third-party credentials, trusted access channels, shared infrastructure, etc. contributed to a third of events and half of all losses.



\$9 of every \$10 lost in extreme events tie to Valid Accounts. Compromised credentials were the initial access vector for 58% of extreme events and account for 91% of associated losses. This isn’t a new finding, but the predominance of this technique among the largest loss events demands attention, and better yet, action.



Patchable vulnerabilities are behind 61% of stolen records. 18% of extreme events exploited known vulnerabilities — definitely not the most dominant vector. But the amount of data stolen (61% of all records) is striking, especially when you consider these weren’t zero-days and patches were readily available.

Like what you see? *Join the vision!*

We intend to continue the IRIS in the future to discover even more insights for managing information risk. If you’d like to join in that effort by contributing relevant data or sponsoring research, please reach out to us via the contact form at www.cyentia.com/iris.



1

EXTREME EVENTS IN CONTEXT

"We demand rigidly defined areas of doubt and uncertainty!"

—Vroomfondel, The Hitchhiker's Guide to the Galaxy

What's an "extreme" cyber incident?

Those familiar with the 2020 IRIS Xtreme may recall that we set a threshold of \$20M in losses or 20M data records affected for incidents to be included in the study. That worked well for the "20/20" theme featured in the IRIS that year and yielded just over 100 events for analysis.

Since that time, our primary criterion for extreme cyber events anchors on the 95th percentile for inflation-adjusted loss magnitude across all incidents over the last 15 years. That designation currently applies to any cyber event that racks up at least \$32M in known/reported losses. The top half of Figure 1 depicts this loss distribution on a log scale.



Figure 1: Distribution of cyber incident losses with the top 5% of extreme events highlighted

The bottom half of Figure 1 isolates the extreme cyber events from the broader loss distribution. It’s clear from this plot that even extreme losses have a long tail that extends into the billions of dollars. Unless otherwise noted, all analysis in this report pertains to this sample of incidents.

One may argue that the definition of an extreme event should be relative to organizational size rather than a fixed value based on the overall loss distribution. A \$32M loss means something different for a giant corporation versus a small to midsize business (SMB). It’s a valid point, and we may take that approach in future studies. For now, we’ll establish some size-specific stats that should be helpful to assessing extreme losses for your organization.

Revenue	Loss percentile	
	50%	95%
More than \$10B	\$2.2M	\$266.2M
\$1B to \$10B	\$1.8M	\$61.8M
\$100M to \$1B	\$466.7K	\$12.3M
Less than \$100M	\$357.0K	\$9.1M

Source: IRIS 2025 (Cyentia Institute)

Risk Pro Tip

The 95th percentile column in Table 1 offers good fodder for discussions with leadership. For example, how does the CFO view losses of this magnitude if they showed up in financial reports? Their reaction can reveal a lot about the firm’s ability to absorb this kind of loss and may surprise you.

Table 1: Summary loss statistics by revenue band

Some simple math from Table 1 reveals an important insight that we routinely uncover when analyzing incidents: losses don’t scale uniformly with organization size. A \$9M loss for a \$50M organization represents nearly 20% of their annual revenue—a huge hit to the bottom line. Comparatively, a \$266M loss for a \$50B organization represents 0.5% of revenue. Still painful, but probably not catastrophic.

NOTE: Losses analyzed in this study tend to reflect direct losses that are easier to quantify (e.g., response costs or lost revenue) and identify from public records (e.g., class action suits or U.S. Securities and Exchange Commission (SEC) filings). Indirect and intangible impacts often aren’t captured. Thus, the values reported herein likely represent a more conservative lower bound of financial losses associated with cyber events.

Are extreme events getting more...extreme?

We defined extreme events as those landing in the top 5% of losses, but that threshold isn't a fixed dollar amount. It's a moving target. Where it moves tells an important story.

To get a sense of how that target has shifted, we can compare losses among extreme events from our prior study against the same figure today. In the 2020 IRIS Xtreme, the median loss from extreme events was \$47M. In this edition, that's climbed to \$100M. So, by that metric, the “typical”² extreme event has doubled in cost between these two studies.

But medians among two five-year samples of loss events aren't the best measuring stick, so let's use a more rigorous approach. Figure 2 applies time series modeling to track how the 95th percentile and median of losses have evolved over the last 15 years.




Figure 2: Trend analysis of median (50th) and 95th percentile losses for all incidents

The result is eye opening: top-end losses have risen 19x since 2008, from \$14.7M to \$281.3M. That makes the overall extreme benchmark of \$32M seem almost quaint by comparison. For context, typical (median) losses saw a 12-fold increase over the same period, which is sobering in its own right for the broader spectrum of cyber incidents.

So what's driving this? No single factor explains a 19x increase, but several forces have been converging for years. Attack surfaces have expanded faster than controls could keep pace—more cloud infrastructure, more third-party dependencies, more APIs, more entry points, more to lose when something goes wrong.

²Extreme events, by definition, aren't “typical.” The midpoint of the top 5% of losses is a rather odd metric, but it does provide a simple point of comparison of loss magnitude for the two time periods.

Attacks have become more disruptive over time, and attackers have diversified their strategies to maximize profits from and impact on their victims.

There's also a measurement effect worth acknowledging. Improved disclosure requirements and more scrutiny on organizational impacts have surfaced losses that once went unreported entirely. Some of the increase reflects rising costs; some of it reflects better information. Both matter for risk management, and both argue for the same conclusion: your loss estimates need to be recalibrated on a regular basis to remain anchored to our evolving understanding of reality.

The broader lesson here echoes what we said in the original Xtreme and what IRIS 2025 reinforced throughout: these thresholds aren't static. A risk model built on what was considered an extreme loss five years ago is already navigating from an outdated map.

How common are extreme events?

Since we've narrowed extreme incidents to the top 5% of losses, they are by definition, rare. But that doesn't really convey how often they occur in the overall population. We identified 136 of them over our five-year period, averaging a handful of extreme cyber events added to the public record per quarter. So, not very common in the grand scheme of cyber things.

Figure 3 depicts the trend (such that one exists) for extreme incidents over the last five years. There appears to be a buildup over the first two years, then an 18-month trough, and ending with several above-average quarters. Ignore the dip at the end; that's due to the reporting lag that's inherent to cyber loss events.




Figure 3: Number of publicly-known extreme cyber events occurring each quarter

At this point, it's worth stepping back and asking, "Are there really only a handful of extreme cyber loss events each quarter?" We suspect the publicly available data undershoots reality, at least to some extent. One key factor behind this suspicion is that only a small percentage of security incidents in our dataset include any information on financial losses. And those that do often provide a less-than-complete accounting of all costs.

The counterargument to these points is that major cyber events tend to get attention. That attention, in turn, drives public and regulatory scrutiny that tends to flush out details that would otherwise go unreported. In other words, public knowledge of extreme loss events is likely more comprehensive than that of run-of-the-mill security incidents.

So, are there extreme security incidents we don't know about? Definitely so. Do we know enough of them to glean useful insights about managing the long tail of cyber risk? Absolutely.

What's riskier—a critical hit or death by 1,000 cuts?

Over the years, we've had conversations (often prompted by the classic likelihood x impact risk matrix) that go something like "Should we be more concerned about a major catastrophe or a bunch of minor loss events that mount up over time?" In other words, a critical hit or death by a thousand cuts? We've run some basic "what if?" scenarios to explore that question but never had the opportunity to answer it in true quant fashion. Until now. What better opportunity than a study focused on the risk of extreme cyber events?




Figure 4: Simulated likelihood of experiencing extreme loss amounts via one large incident or multiple smaller incidents

Figure 4 presents results from simulations incorporating Cyentia’s overall probability and loss models for cyber events across varying time windows.³ The outcomes of each run are tallied to determine how many years it takes for the cumulative losses from smaller incidents to exceed expected losses from a single extreme loss event (\$32M). The chart indicates that, on average, it takes 7 years of “minor cuts” to equal the damage caused by one “critical hit.” It’s also clear in Figure 4 that the risk of extreme events always exceeds that of minor incidents. In some respects, this resembles the Taleb distribution, named after the author of *The Black Swan*. Such a distribution exists in an investment portfolio that takes small regular gains while hiding a small but catastrophic risk of total loss.

This dominance of extreme events, however, does lessen as organizations get larger. That shift is clear in Figure 5, which recalibrates the lines for organizations over \$100B in annual revenue. Here, the damage taken from numerous smaller cuts does indeed eclipse that of the rarity of critical hits within a few years. Granted, not many of those companies exist, but definitely good to know if you’re responsible for managing risk for the upper echelons of the Fortune 500.




Figure 5: Simulated likelihood of experiencing extreme loss amounts via one large incident or multiple smaller incidents for very large firms

Speaking of managing risk, that’s going to look different for these two classes of incidents, regardless of organization size. Most organizations will find it easier to weather a series of smaller losses spread over time than a single extreme loss. For these periodic events, focus on layered preventive controls and efficient (preferably automated) detection and response. You’re optimizing for recurring incidents.

Mitigation still matters for rare high-impact events, but resilience and antifragility become a much more effective way to respond. Look to cyber insurance and cash reserves to cover the financial outlay and ensure operational recovery capabilities prevent one bad day from turning into an existential threat.

³This uses the same approach to probability modeling described in Appendix 1 in the IRIS 2025 (www.cyentia.com/iris)

Bottom line: If your security budget is 90% prevention and 10% resilience, does that match what this data is telling you?

Who's prone to extreme cyber events?

We established in prior sections that extreme cyber incidents are relatively rare. But are some organizations more prone to these tail events than others? Figure 6 has thoughts on the matter.

Try to ignore the different colors comprising the bars in Figure 6; we'll get to that in a moment. For now, just note the overall length of the Financial sector's bar. About 40% of extreme cyber events in the last five years occurred in that single industry. That's not a statistical quirk. It's well known that cybercriminals (and state actors) tend to follow the money. Banks, investment firms, and — here comes that gray swath — cryptocurrency exchanges make very tempting targets.




Figure 6: Proportion of extreme incidents by industry

Which brings us to the elephant in the chart. About one-third of all extreme incidents in this study involved cryptocurrency thefts and fraud. Those events account for the majority of extreme losses recorded in the Financial sector and a meaningful slice of those in Information Services as well.

Now look at just the red portion of the bars. The spread among the top several industries is less than 10 percentage points. In other words, absent crypto hacks, extreme events are far from being a Finance-specific problem. Everyone's represented here.

Manufacturing tied for the #2 spot deserves mention. In the IRIS 2025, the Information sector had 4 to 5 times the number of security incidents compared to Manufacturing. But the latter is far more represented among

extreme events. As will be seen shortly, much of that ties to operational and supply chain disruptions that disproportionately impact manufacturers.

If your organization has nothing to do with cryptocurrency, you may be wondering whether a third of this report is about to become irrelevant to you. Hold that thought — we'll address it directly in the next section.

Do crypto hacks obscure insights from this study?

Fair question, and we're not going to pretend we didn't ask it ourselves. The honest answer is: it's complicated. But probably less so than you'd expect.

This isn't our first crypto rodeo in the IRIS series. The 2020 Xtreme identified cryptocurrency theft as a notable incident pattern — 16 events out of 103. That number has grown substantially in the current dataset: 43 of 136 events. Crypto incidents have gone from a novel subset to a meaningful portion of the study population, and that warranted some serious discussion on our end before proceeding.

The first and most important question was whether crypto incidents distort our overall findings. If crypto hacks were dramatically more, or less, costly than other extreme events, they could skew summary statistics and make the findings less actionable for the non-crypto reader. Figure 7 addresses this directly.




Figure 7: Comparison on reported losses among crypto and non-crypto incidents

Crypto hacks not your concern? Don't worry—this report provides many non-crypto insights! Plus, we'll let you know if results warrant making a distinction.

While the sample of crypto incidents does carry a moderately higher median loss (\$126M vs. \$92M for non-crypto), the overall distributions and 95th percentiles are similar. On the loss dimension at least, crypto and non-crypto extreme events look more alike than different.

That said, financial losses aren't the whole story, and there are areas of this report where crypto and non-crypto events part ways more dramatically. When those differences are meaningful, we'll call them out. If not, we won't force the contrast to avoid losing the forest for the TREES.⁴

How is “extreme” different from “material”?

In 2023, the Securities and Exchange Commission (SEC) [adopted new rules](#) requiring public companies to disclose “material” cyber incidents under Item 1.05 of Form 8-K. The criteria for materiality are not specific to cyber events; they derive from long-standing cases in securities law. In short, the incident is material if a reasonable investor would likely consider it important in making investment decisions. How does the nature and magnitude of these disclosed “material” events compare to our corpus of “extreme” events? That’s a great question we take up in a later section.

⁴TreesCoin (TREES) is not a major cryptocurrency, but we simply couldn't resist the pun.

What can we learn from extreme crypto incidents?

from Cyentia Institute Executive Fellow, Tony Martin-Vegue

Don't skip the prior section. You see "crypto" and think it doesn't apply because most of us don't work for cryptocurrency exchanges. But look closer at the causes, attack vectors, and consequences. There's more here than you'd expect.

Read these events through the lens of business email compromise, wire fraud, and AI-powered phishing. As attackers get more sophisticated and find new ways to steal money directly, crypto heists offer a preview of where things are heading.

Gone in 60 Seconds: In crypto heists, the money is gone in minutes. Most organizations assume they'll have time to detect and respond. Crypto hacks show what happens when that window doesn't exist. Do your controls assume response time you might not have?

Direct access to value: Most incidents like data breaches and ransomware involve a monetization step after the initial attack before the attacker sees a payday. Crypto hacks skip the monetization step because the attack IS the theft. But any organization with wire transfer capabilities, treasury functions, or BEC exposure faces a version of this. The pattern is the same, just slower. And with AI-generated deepfakes making impersonation easier, that gap is closing. Does your risk model include cash as an asset at risk?

Concentration risk: Many crypto hacks wiped out entire exchanges in a single event: one attack, total loss.

The lesson we can derive here is about single points of catastrophic failure. One admin account with domain-wide privileges, one vendor your entire operation depends on (mega cloud vendor, anyone?), or one person authorized to approve wire transfers without dual authorization, just to name a few.

Crypto exchanges learned the hard way that "one key to rule them all" is an existential vulnerability. Look at your own environment: where does one failure equal total loss?

More reading on the largest crypto thefts: [Crypto's Biggest Hacks and Heists \(Reuters\)](#)

2

THE IMPACT OF EXTREME EVENTS

"Oh yes, the past can hurt. But from the way I see it, you can either run from it or learn from it."

— Rafiki, *The Lion King*

What are extreme losses relative to revenue?

Having defined extreme cyber events and provided some context about them, let's turn to the question of impact — specifically, how hard these incidents actually hit the organizations on the receiving end. We'll work through several lenses, beginning with the most fundamental: relative financial impact.

A \$100M loss — roughly the median for extreme events in this sample — means something very different depending on who's writing the check. For a sprawling megacorporation, it might be a painful but manageable quarter. For a mid-market firm, it could quickly turn the balance sheet from black to red. That's why, throughout the IRIS series, we express losses as a percentage of the victim organization's annual revenue. Absolute numbers grab headlines; relative numbers tell the story.

Figure 8 makes this conversion for all extreme events in the current study. The range is remarkable, spanning from a small drop in the revenue bucket (0.03%) to well over 3,000 times annual revenue!



Figure 8: Extreme event losses as a percentage of annual revenue

A third of extreme events cost less than 1% of annual revenue, the kind of hit that stings but doesn't scar. Another quarter fall in the 1-10% range, which is likely material, but survivable. Things get more serious for the 13% of events in the 10-100% range, where losses start to reshape organizational trajectories. But the

truly striking finding sits at the far right of the chart: nearly 3 in 10 extreme events (29%) generated losses that exceeded the victim's entire annual revenue. That's a bill for a single event that's larger than the whole business.

If loss multiples in the thousands or millions sound like funny money, you're onto something. A closer look at Figure 8 reveals that the most extreme ratios cluster around cryptocurrency exchange hacks (highlighted in gray), where stolen assets can dwarf the exchange's operating revenue by orders of magnitude. But it would be a mistake to write off the rest of the chart as a crypto curiosity. A meaningful number of non-crypto events also push well past the 100% threshold. The crypto hacks definitely stand out, but they're not outliers.

What's the primary outcome of extreme events?

Another lens through which we can assess the impact of security incidents is by their high-level outcomes. We broadly categorize events based on whether they resulted in the exposure of sensitive information, operational disruptions, or fraudulent activity. Think of it as a business-oriented version of the classic “CIA Triad” of confidentiality, integrity, and availability.

Incidents can and do involve multiple outcomes, but one tends to define the primary character of the event. Figures 9 and 10 compare how these three categories shake out for all security incidents versus extreme events specifically — first by event count, then by total losses.



Figure 9: Proportion of all incidents (left) and extreme events (right) by primary outcome

The contrast is striking. Among all incidents in the broader dataset, data exposure reigns supreme: roughly 90% of events are primarily exposure-driven. Disruption and fraud are statistical afterthoughts by comparison. If you managed your security program based solely on what happens most often, you'd optimize almost entirely for breach prevention.

Extreme events tell a very different story. When we filter down to the tail of the loss distribution, exposure drops to under 40% of the population. Disruption surges over 43% — more than five times its share in the overall incident mix. Fraud climbs to 24%, a 27-fold increase from its baseline representation.

The same reversal holds when we look at losses rather than event counts in Figure 10: exposure cedes ground, while disruption and fraud claim a substantially larger share of the financial damage. (Yes, a good chunk of those fraud events involve cryptocurrency theft. You knew that was coming.)




Figure 10: Proportion of all losses (left) and extreme losses (right) by primary outcome

The practical implication here is worth noting. Extreme cyber events aren't simply bigger versions of ordinary incidents — they're structurally different, driven by a more diverse mix of threat scenarios. That diversity has real consequences for how organizations prepare. Avoiding or mitigating isn't a simple "just do this one thing."

Do bigger data breaches lead to bigger losses?

Readers of the 2020 and 2022 IRIS editions may recall our analysis demonstrating that the number of data records compromised in a data breach tells you essentially nothing about the financial losses that follow. In case you're wondering if that's changed (it hasn't), here's a [2026 update to that analysis](#). Despite the enduring popularity of "cost per record" as a benchmark, using it to estimate losses results in massive errors. We've said it before, and the data keeps saying it with us.

Maybe the cost-per-record model works better for predicting losses from extreme events, right? It's a reasonable hypothesis. It's also wrong. If anything, predictions are worse at the extremes. But don't take our word for it; let's look at the data.




Figure 11: Correlation between data records stolen and reported losses for extreme events

Figure 11 plots losses against records exposed for the extreme events in this study that involved a known data count. Cryptocurrency thefts are excluded — those stolen bits do carry a per-unit value, at least until the market disagrees. What remains is a scatter plot that is, charitably, a random cloud of dots. The regression line threading through it is almost perfectly flat, producing the same loss prediction for a breach of a thousand records as for a breach of a billion.

Will that stop headlines from dutifully reporting the updated cost-per-record figure? Almost certainly not. But perhaps it'll give you good reason to raise an eyebrow the next time someone reaches for that number to build a risk estimate.

What types of losses do extreme events incur?

Ideally, every major security incident would come with a comprehensive accounting of all losses incurred by the organization affected. That rarely happens. For 16% of extreme events in this study, only a single overall loss figure was reported. Half of incidents yielded just one identifiable loss type, and fewer than 20% reported three or more. That's frustrating from an analytical standpoint, and we'd argue it hinders the broader industry's ability to learn from and prepare for catastrophic events. There's ample room — and genuine need — for improvement.



Figure 12: Categorized loss distributions for extreme events

With those limitations in mind, Figure 12 and Table 2 present what we were able to piece together about the types of losses associated with extreme events.

Financial theft and fraud leads the field in both frequency and magnitude, with a median loss of \$105M — well above every other category. As Table 2 makes clear, that's largely a crypto story: financial theft and fraud accounts for 91% of their reported losses. For these events, it's not really a loss type so much as the entire point.

	Non-crypto		Crypto	
	Frequency	Losses	Frequency	Losses
Business interruptions & lost sales	29.2%	20.4%	14.3%	4.6%
Response & notification costs	45.8%	17.8%	4.8%	—
Fines & judgements	33.3%	10.0%	7.1%	1.1%
Financial theft & fraud	18.1%	6.9%	97.6%	91.2%
Cost to replace IT & people	19.4%	6.7%	—	—
Increased security investments	26.4%	3.0%	—	—
Ransom/extortion payment	15.3%	2.6%	2.4%	0.0%
Brand or reputational damage	4.2%	0.9%	—	—

Table 2: Frequency and losses for different types of incident impacts; events may have more than one impact so amounts will not sum to 100%

Strip out the crypto hacks, and a different set of categories takes center stage. Among non-crypto extreme events, the big three are business interruptions and lost sales, response and notification costs, and regulatory fines and judgments. Response costs are the most commonly reported loss type, reflecting the need to contain and recover from a major event. But business interruptions drive the largest share of actual losses, consistent with our earlier finding that disruption is the defining outcome of non-crypto extreme events. Fines and judgments round out the top three, a reminder that regulators often arrive after the attackers have left.

Though not a direct impact of the incident, increased spending on personnel and technology to improve security is quite common. This often comes in the form of a “here’s how we’re going to do better” commitment that follows major incidents and represents a real, if forward-looking, cost.

Brand or reputational damage appears in just three incidents in the entire dataset. For anyone who has sat through an executive risk briefing, that number warrants a pause. Reputational harm is perhaps the most commonly expressed concern when discussing potential impacts from major cyber events. It's the wildcard that often sends scenarios to the top of risk registers regardless of their measured financial profile.

We anticipate some "yeah, but" reactions to that finding, and two counterarguments in particular are worth addressing.

The first: Brand damage is hard to quantify — perhaps it's not reported as a loss. That's a fair argument. We'll search for signs of indirect and intangible losses in the sections that follow.

The second: What about stock price drops and credit rating impacts — aren't those forms of reputational damage? They are, and they're important enough to deserve a deeper discussion. We'll get there!

What secondary impacts are most common?

Not every consequence of a major cyber event shows up on a balance sheet. Beyond the direct costs of financial theft, business interruption, and incident response arise intangible impacts, secondary losses, and externalities that reverberate well beyond the original victim organization. We looked for signs of such impacts across all extreme events in this study, and what we found is summarized in Figure 13.

Supply chain impacts were the most pervasive, occurring in nearly half of all extreme events (46%) and accounting for over half (55%) of all associated losses. The forms these impacts took varied: production outages that disrupted the flow of goods, service failures that left dependent partners stranded, and ripple effects that spread to interconnected third and fourth parties. This last dynamic is well-documented in Cyentia's [Ripples Across the Risk Surface](#) research, which found that the financial costs incurred by downstream parties can equal or even exceed those borne by the primary victim. A breach doesn't stay where it first hits.




Figure 13: Frequency of secondary impacts

Government hearings were the second most common secondary impact, triggered by 38% of extreme events. That's a meaningful signal in its own right. It indicates that nearly four in ten extreme incidents attracted sufficient public or political attention to warrant formal legislative or regulatory scrutiny. The direct costs of responding to such hearings (executive time, legal preparation, testimony, remediation commitments made under oath) rarely make it into loss tallies, but they're real and often substantial. Organizations building risk scenarios around extreme events should treat government engagement as a likely cost at the tail of the distribution, not an unlikely edge case.

Further down the chart, the numbers get smaller but more consequential. About 9% of extreme events involved the affected organization filing for bankruptcy, and near 7% resulted in the company ceasing to exist entirely. Those numbers may seem small if you've heard the "60% of companies go out of business after a breach" [myth](#) that persists in the security community. The reality is that going belly up after an incident is very rare in the grand scheme of things, but that risk is significantly elevated when we cross into extreme event territory.

Credit rating impacts appear in 3% of extreme events. Not common, but notable enough to bridge directly into the question of how financial markets and lenders register the impact of major security incidents. We'll elaborate on that topic in the next section.

Executive churn rounds out the chart at 2%. It's sometimes tough to tell if the firing/leaving was truly a result of the breach or simply coincidental timing, so we took a conservative route of looking for a direct causal relationship. It's a reminder that extreme events occasionally claim careers as well as capital.

Breached into Oblivion?

Adrian Sanabria maintains an [ongoing public investigation](#) into companies destroyed — and not destroyed — by cyber incidents. It's a useful empirical check on many of the catastrophic-sounding headlines we hear about, but rarely get the full story of the true aftermath. The reality, as with so much in this space, is more nuanced when you look at the actual cases.

ON STOCK DROPS AND CREDIT DOWNGRADES

The prior section noted that brand damage rarely surfaces as a quantified loss in extreme cyber events. Stock price drops offer counterargument to that. They're visible, dramatic, and they feel like the market rendering a verdict on reputational harm. So, let's take them seriously.

The academic literature on this is actually reasonably settled. Studies consistently show that stock returns decrease in the short term following cyber events, but the effect tends to reverse within roughly two weeks for the average incident.⁵ The dip can be larger and longer for more severe incidents (up to 250 days post-announcement), but even then, recovery typically follows.

What the market appears to be pricing in, however, is worth examining carefully. The evidence suggests that investors aren't primarily reacting to the economic damage already done; they're reacting to uncertainty about future vulnerability. Markets treat a disclosed breach as a forward-looking indicator rather than a backward-looking loss tally. In that sense, a stock price movement following a breach resembles a risk assessment more than a damage invoice.

This means stock prices are a noisy proxy for actual organizational impact. Market, media, and organizational factors all shape the reaction as much as the underlying incident does. Organizations that respond well tend to fare better with investors regardless of incident severity. There are more direct signals that measure whether a cyber event has genuinely impaired an organization's ability to generate value: realized changes in revenue (particularly in affected business lines), customer retention, market share shifts, remediation and legal costs, and credit rating changes.

On that last point, formal credit downgrades following cyber incidents remain rare, but the cases that have occurred are instructive. S&P downgraded the City of Hamilton, Ontario, following a 2024 ransomware attack. This wasn't due to financial impairment; the city emphasized that its finances remained strong following the incident. The downgrade was attributed to audit reporting failures caused by the attack. The Co-operative Group in the UK received a similar signal following their cyber incident in April, 2025. S&P downgraded the group to 'BB-' from 'BB', reflecting the financial strain, but maintained a stable outlook based on expected recovery by the end of the year.

There have also been a few incidents where agencies publicly flagged a cyber event as “credit negative” or potentially rating-relevant. Such was the case after MGM's 2023 ransomware attack. Moody's warned the event was negative for MGM's credit standing due to significant revenue losses, operational disruptions, and reputational risks, but a full downgrade was not issued.

These cases illustrate that when credit agencies act, it's often the secondary effects of an incident — governance failures, reporting lapses, operational credibility concerns — that drive the rating action. The rarity of direct, cyber-driven downgrades, like the rarity of quantified reputational losses, may say something meaningful about how lenders assess the long-run consequences of these events. Or it may reflect that the bar for a formal rating action is simply higher than the bar for executive anxiety. Either way, it's a more rigorous signal than stock price movement. And the signal is still mostly quiet.

⁵There are many referencable studies on this topic, but here's a recent one to start with: HACKED: Understanding the stock market response to cyberattacks. <https://www.sciencedirect.com/science/article/abs/pii/S1042443124001483>

What can we learn from SEC 8-K filings?

We touched on the 2023 [SEC ruling](#) on reporting “material” cyber incidents earlier and distinguished their disclosure criteria from our definition of extreme events. That section ended with a promise that an actual analysis of those disclosed material events was forthcoming. Now we’re back to make good on that promise.

To help us with that, we’ve partnered with [Board Cybersecurity](#), which, among other sources, [tracks security incidents](#) reported via SEC 8-K filings. They contributed 205 incidents that were reported to the SEC, including 34 via Item 1.05 (material cyber incidents), 33 via Item 7.01 (material nonpublic information), and 145 via Item 8.01 (other events).⁶

Reported losses are lacking: Only 22 of the 205 incidents included an estimate of financial impact in the filing(s).⁷ There are two ways to look at this.

On one hand, 10% of incidents with loss values is significantly higher than other datasets we’ve analyzed. Combine that with the scrutiny of SEC reporting, and these filings hold a lot of promise for extracting high-confidence data on financial losses. Even better as the number of reported events grows over time.

On the other hand, we find it a bit of a letdown. Given that these are ostensibly material cyber incidents (or materially important info), we’d expect companies to have counted the costs and reported them clearly and comprehensively. That’s not happening — at least not via Items 1.05, 7.01, and 8.01.⁸ We could dwell on that further to speculate reasons, but let’s keep going to see what else we can learn.

Losses are high but often not extreme: With the caveat that n=22 makes this more of a curiosity than a basis for credible model parameters, the loss figures that do appear in these filings tell an interesting story about where 8-K disclosures sit in the broader landscape. The median reported loss of \$13.5M sits well above the ~\$600K median across all security incidents in the IRIS 2025 dataset — confirming that these are genuinely significant events by ordinary standards.

But they’re nowhere near the \$100M median of the extreme events we’ve been analyzing throughout this report. Material incidents, it turns out, occupy a real but distinct tier: serious enough to clear the SEC’s investor-relevance bar, but not at the level of tail losses that define this study.

Most material incidents aren’t material: While the reporting of financial losses might be rarer than we’d hope, many 8-K filings do include the companies’ assessment of materiality of the incident. This is valuable information, even if qualitative in nature.

Of the 205 incidents in this dataset, only 7 are currently assessed as material by the reporting organizations, with 5 more assessed as likely to be material once the dust settles (6% of the total). The other 94% are either assessed as not material or still under evaluation.

⁶Yes, the numbers sum to more than 205; some incidents are reported under multiple 8-K Items.

⁷Board Cybersecurity combines information on a common incident across multiple filings, so if losses were reported or updated in a later 8-K filing, that’s accounted for here.

⁸Financial losses associated with cyber events are reported in the regular 10-Q and 10-K filings. We did not integrate those because our primary focus here is on the new rule for material cyber incidents.




Figure 14: Loss distributions for all incidents, extreme incidents and 8-K filings

Put plainly: the vast majority of "material cybersecurity incidents" filed with the SEC aren't, by the filers' own reckoning, material. This almost certainly reflects the confusion and overcaution that characterized early compliance with the new rule — companies filing defensively under Item 1.05 before the SEC clarified that doing so for undetermined or immaterial incidents was itself problematic. This likely also stems from companies not updating their estimates after the initial disclosure. The 8-K corpus, at least for now, is as much a record of regulatory uncertainty as it is a record of genuine material risk.

Material incidents won't ruin your rep: Yes, yes — this heading smacks of hyperbole. But it's based on the fact that nearly all impacts assessed as material were operational in nature. Reputational and brand damage — the boardroom bogeyman that routinely overrides quantitative risk assessments, elevates incidents to the top of risk registers, and justifies security spending that might not otherwise survive scrutiny — shows up as a "possibly material" impact in exactly 2 of the 205 events.

This dearth of material reputation hits may be because the impacts are hard to quantify. Maybe reputational damage is omitted because others aren't reporting it either. It's also something companies wouldn't want to disclose to investors if not required to do so. Whatever the reasons behind it, the gap between how prominently reputation risk features in executive threat discussions and how rarely it materializes in formal disclosure is, to put it mildly, worth noting.

Operational and downstream impacts dominate: The fact that most losses aren't expected to be material doesn't mean these incidents had no impact on the companies reporting them. Even incidents that don't rise to material significance can leave real marks, and the filings reflect that.

Overall, 91 of the 205 events (44%) include confirmed references to at least one type of operational impact. Disruptions to business operations were the most commonly cited, followed by system outages, and impacts to various business activities. Taken together, these paint a picture of incidents that hit operations hard enough to document, even when executives don't expect the financial fallout to be investor-relevant.




Figure 15: Impacts from material incidents reported in SEC 8-K filings

A supply chain thread also runs through the impact categories. Disruptions to order placement and fulfillment feature prominently, suggesting that many of these incidents propagate beyond the directly affected organization into its customer and partner relationships — a theme that resonates with the broader third- and fourth-party risk dynamics documented in this study and elsewhere in the IRIS series. These filings may understate material financial risk, but they're quietly telling a story about operational fragility that deserves its own attention.

3

THE THREAT OF EXTREME EVENTS

"Persistence without insight will achieve the same outcome."

— The Armorer, *The Mandalorian*

This is not a threat intelligence report. We're not reversing malware, tracking malicious infrastructure, or piecing together evidence for adversary attribution. But threats are a key component of cyber risk, and understanding them is crucial to mitigating that risk effectively. To that end, this section provides observations on the who, what, and how behind these extreme events.

Who's behind extreme events?

Researching cyber adversaries is a field all unto itself. We're not doing first-hand attribution here; it's not necessary to our goals. Major incidents like those analyzed in this study tend to attract that level of investigation from qualified agencies and analysts. And we're more than happy to stand on their shoulders for a better view.

We're also glad to leverage partnerships like Feedly for Threat Intelligence to enrich the base cyber loss data sourced from Zywave. This made it considerably easier to synthesize open source intel shared on these extreme events and integrate that community knowledge into our analysis.

Figure 16 provides a high-level breakdown of the primary threat actors behind extreme events using three categories that readers of the Verizon DBIR will find familiar: external, internal, and third party (or "partner" in DBIR lingo).




Figure 16: Primary threat actors in extreme events

External actors steal most of the limelight here. They're behind 95% of extreme cyber events over the last five years and 86% of related losses. Show-offs. And that's not just a peculiarity that emerges at the extremes—they're predominant among regular ol' security incidents too.

Digging into external actors, organized criminal groups are behind the majority of events, outnumbering state-affiliated actors roughly four to one overall. Among cryptocurrency hacks, however, that ratio reverses. The states take the lead, and North Korea's well-documented crypto capers represent a large part of that. Financial losses tend to be significantly higher for incidents involving state-affiliated actors (\$108M median vs. \$60M for organized criminals).

Internal actors account for just 3% of extreme events and about that same proportion of losses. The old myth that [80% of attacks originate from insiders](#) has largely receded in recent years as threat research continually reinforces an outsider majority. But holdouts remain, and these results are unlikely to change their minds any more than the existing evidence has.

To be clear: employees undoubtedly contribute to breaches through poor security decisions, falling for phishing lures, neglecting system updates, and otherwise creating the openings that external actors exploit. That's real risk.⁹ But these scenarios involve insiders as contributing factors rather than primary actors, and Figure 16 deals with the latter. That's a meaningful distinction when allocating defensive resources to reduce insider risk.

Third parties, as primary actors, appear modest at 2% of frequency. They do manage to punch above their weight, though, causing 11% of all recorded losses. But Figure 16 is just the tip of the iceberg when it comes to the role of third parties in extreme events.




Figure 17: Third parties as indirect contributors to extreme events

Below the surface, the third-party risk picture changes dramatically. Third parties indirectly — though substantially — contribute to a third of extreme events and half of all losses. This includes misuse of stolen third-party credentials, access or propagation via trusted applications, and disruption of shared infrastructure. Combined with findings presented earlier on supply chain impacts and ripple effects, third-party risk emerges as a major factor in extreme cyber events.

⁹The Verizon DBIR reports that the “human element” contributes to about three-quarters of all breaches.

What scenarios drive extreme events?

Effective risk mitigation is scenario-dependent. The controls that neutralize a rogue insider have little overlap with those that stop a ransomware gang, and physical defenses won't thwart system intrusions. That's why we classify all incidents into broad patterns based on common actors, actions, and technical impacts.¹⁰ Understanding which patterns dominate extreme events matters practically, not just analytically.

Figures 18 and 19 compare how these patterns distribute across all incidents versus extreme events, first by frequency, then by losses. The format should be familiar: overall on the left, extreme on the right.

On frequency, the first read of Figure 18 might seem unremarkable — system intrusions and ransomware lead both populations, and the ordering doesn't change much as you move down the chart. But the annotations earn a second look.

Fifty-seven percent of extreme system intrusions are cryptocurrency exchange hacks. Strip those out, and the system intrusion bar on the right side of the chart shrinks considerably, leaving ransomware as the more dominant scenario among extreme events.

The growth of system failures among extreme events is less obvious but worth noting. Disruptions are expensive, and extended outages are more likely to accumulate losses large enough to qualify as extreme. Accidental disclosures, meanwhile, contract at the extremes in a way that makes intuitive sense. Such “own goals” are certainly embarrassing, but don't usually generate losses that qualify as extreme.




Figure 18: Prevalence of incident patterns among extreme events

¹⁰Definitions for the incident patterns shown here are provided in Appendix B.

We underlined “usually” there for a reason, which becomes clear as we switch to the breakdown of losses in Figure 19. Accidental disclosures cause more than their fair share of losses relative to their frequency among extreme events. Most of those losses come from one particularly costly, unusual incident.




Figure 19: Losses associated with incident patterns among extreme events

System intrusions and ransomware exhibit similar levels of dominance for losses as we saw with frequency. Crypto theft takes an even larger share of extreme losses among intrusions, leaving ransomware as the undisputed heavyweight risk champion for extreme cyber events.

The scam or fraud pattern presents an interesting inversion. It's a substantial loss driver in the overall population, where the aforementioned “thousand cuts” metaphor applies. But it constricts at the extremes, where individual BEC events rarely generate losses large enough to qualify.

The risk management takeaway is clear. System intrusions and ransomware represent the highest exposure by a wide margin, and that remains true whether you're building controls for the full incident distribution or specifically hardening against tail events. The next section examines the ATT&CK techniques behind these two patterns in more detail, providing sharper guidance on where mitigation investments are most likely to move the needle.

How do attackers gain initial access?

Given the dominance of system intrusions and ransomware documented in the previous section, the natural next question concerns how attackers slipped through defenses. Reconstructing a full kill chain from publicly available information is difficult. Organizations disclose what they must, not necessarily what would be most instructive to the community. But the method of initial access is one of the more reliably reported details, whether by the victim organization or researchers who investigated the incident.

This serves our purposes well because preventing the initial intrusion is the first thing to consider when evaluating how to avoid becoming victims of extreme cyber events like these. If they can't get in, everything else they might do becomes moot.¹¹

We'll leverage the [initial access techniques](#) defined by MITRE ATT&CK, a widely-used framework of adversary tactics, techniques, and procedures (TTPs). While it could be argued that there's only one true initial access technique per incident, we allowed more than one where appropriate. Table 3 presents the results across frequency, losses, and data records stolen.

Valid accounts leads by a margin that should stop you in your tracks. Though we've gotten used to seeing this technique at the top whenever we do this kind of analysis, it's still shocking that more than \$9 out of every \$10 lost in the most impactful incidents of the last 5 years tie back to the compromise and reuse of legitimate credentials. It's hard to imagine a clearer or louder risk signal. Identity and access management isn't just one of many domains of cybersecurity; it's the primary battleground.

	Frequency		Losses		Records	
	Rank	Percent	Rank	Percent	Rank	Percent
Valid accounts	1st	58.1%	1st	91.0%	2nd	39.3%
Exploit public facing app	2nd	40.4%	3rd	44.2%	1st	87.1%
Phishing	3rd	39.0%	4th	37.9%	5th	9.0%
Trusted relationship	4th	19.9%	2nd	64.9%	3rd	21.6%
External remote services	5th	16.9%	5th	21.5%	4th	9.8%
Supply chain compromise	6th	6.6%	6th	9.7%	7th	1.2%
Drive by compromise	7th	5.1%	7th	5.2%	6th	7.2%
Content injection	8th	2.2%	8th	2.0%	8th	0.0%
Hardware additions	9th	0.7%	9th	0.2%	9th	0.0%

Table 3: Initial access techniques observed in extreme events; multiple techniques may apply to an incident so columns may not sum to 100%

Phishing often gets credited as the primary siphon for valid accounts, but the difference in frequency and losses between these two techniques proves it's not the only one by far. Among these extreme events, phishing was as likely to be a delivery mechanism for ransomware and other malware as it was a tool for harvesting credentials.

Exploiting public-facing applications ranks second in frequency at 40% and third in losses at 44%. But its most striking statistic is in the records column: 87% of all stolen data across these extreme events involved this technique. When attackers find a vulnerability in an internet-exposed application, they don't just infiltrate the front end. They often get back-end access as a bonus due to trust architecture between web servers and databases.

¹¹ Just to be clear—we're not claiming that preventing lateral movement and post-compromise detection isn't important. Such capabilities are critical to an effective defense.

External remote services offer another fast-track to penetrating the inner sanctum of corporate networks. Especially when those remote access tools and accounts are provisioned for administrators. At 17% of incidents and 22% of all losses, it's clear these internet points of presence are a large and tempting attack surface for adversaries.

Trusted relationships reveal the widest frequency-losses ratio in the table. At 20% of events, they rank fourth among initial access techniques. But the 65% in the losses column reflects how trusted third-party access tends to be broad, deep, and difficult to detect once abused. This is where our decision to allow multiple initial access techniques becomes analytically meaningful: valid accounts and trusted relationships were frequently used together, and that pairing was both common and very costly. A vendor (or an attacker) with valid credentials using legitimate access channels is about as close to an invisible threat as one can get.

Supply chain compromise sits at 7% frequency and 10% of losses, numbers that might seem modest given how prominently supply chain risk has featured throughout this report. The apparent discrepancy is definitional: MITRE's [definition of supply chain compromise](#) refers specifically to adversaries corrupting hardware or software components, a narrower scope than the third-party risk dynamics we've discussed elsewhere. The bigger supply chain picture weaves in all these threads together to see the complete arc.

Beyond ATT&CK initial access techniques, three cross-cutting attack characteristics round out the picture of how extreme events unfold. They appear in Figure 20 by relative frequency (x), losses (y), and data records affected (size).




Figure 20: Common threat observations among extreme events

Malware was involved in 52% of incidents and contributed to just over half of all financial losses. That's a big footprint, but it's also notable that malware apparently wasn't necessary for success half of the time. Attackers who can accomplish their objectives using tools and access that already exist in the environment have little incentive to introduce foreign code that might trigger detection.

Speaking of, **living off the land** (LOTL) techniques were used in 38% of incidents and accounted for that same proportion of losses. This includes using the victim's own legitimate system tools to move laterally, escalate privileges, exfiltrate data, etc. Attackers borrowing what's already in the environment are in a good position to quietly exfiltrate data at scale, which accounts for the high numbers of records stolen. The stats here likely understate reality, since LOTL activity is not often reported in post-incident public disclosures.

Exploits of **known vulnerabilities** (CVEs) were comparatively less common at 18% of incidents and 16% of losses, but again, this is likely understated due to limited reporting. The more alarming story is that 61% of stolen data records involved known vulnerabilities. The word "known" carries weight here. They were vulnerabilities with published CVEs—most with available patches—that remained unaddressed long enough to be weaponized against organizations suffering some of the largest cyber losses in recent history.

For defenders, the actionable thread connecting these findings runs directly through MITRE ATT&CK's mitigation and detection guidance. The initial access techniques documented in Table 3 each map to specific, testable controls — from multi-factor authentication and privileged access management for valid accounts, to patch cadence and attack surface reduction for public-facing applications and known CVEs. MITRE's enterprise mitigation library and detection strategies are the natural next stop for translating this analysis into a prioritized defensive agenda.

What can you do to help prevent attackers from gaining a beachhead in your environment? Well, that could be an entire report on its own, but the [Mitigations](#) and [Detections](#) projects from MITRE are excellent resources to get you started.

4

NOW WHAT?

*A practitioner's guide to doing something with all this,
by Tony Martin-Vegue*

The IRIS Xtreme gives us something rare in cybersecurity: loss data from the tail of the distribution. We at Cyentia don't want this report to sit on a shelf. We want you to take the findings and put them to work.

This report pairs nicely with the IRIS 2025 (and upcoming IRIS 2027), where we cover typical cyber events. It turns out we need to understand both typical and extreme events to make well-rounded security decisions, because we should respond to and invest in them differently.

Think of it like adverse events in your house. Broken door knobs, leaky faucets, and drafty windows are typical. We make repairs, invest in better fixtures, mitigate the everyday wear and tear. But we can never lose sight of the extreme events: a total loss from a fire or natural disaster. We plan for those differently. We buy homeowner's insurance and install smoke detectors that we hope rarely go off. We don't just buy a nicer door knob and call it a day. Too many organizations respond to extreme cyber risk the same way they respond to everyday incidents; better patching, another tool, tighter configurations. That's buying a nicer door knob when what you need is a fire suppression system. The IRIS Xtreme report helps us plan for those improbable but very costly events; the ones that won't happen often but can change the trajectory of an organization when they do.

With that, let's look at core decisions you can make with this data. These are all things you can do tomorrow.

Decision 1: "What loss range should we plan for?": Your CFO, CEO and Board may be pondering this very question right now, and most security teams don't have a good answer. Every organization has some notion of what a bad cyber event would cost them, but without a mature risk quantification program, the answer is often an adjective ("high!") or, even worse, anchored to the last headline or borrowed from a vendor report. The IRIS Xtreme data lets you replace that with something defensible, even without a mature risk quantification program.

The median extreme loss in this study is \$100M, more than double the \$47M median from the 2020 IRIS Xtreme. But the headline number isn't the point. What matters is the range and how it maps to organizations like yours. Losses as a percentage of revenue span from less than 0.03% to over 3,000 times annual revenue. As we consistently find in the IRIS series, larger firms cluster at the lower end of that scale, while SMBs absorb far greater relative damage.

That's why Table 1 in this report is one of the most actionable things we've published. It breaks down loss thresholds by revenue band, so you can see what the 95th percentile loss looks like for an organization of your size. A \$9M loss might be extreme for a \$50M company, which is nearly 20% of annual revenue. For a \$50B company, the extreme threshold is much higher in absolute terms but much smaller relative to the business.

How to help your CFO sleep better tonight: take your current "worst case" cyber loss estimate and check against Table 1. Find your revenue band. If your worst-case ransomware scenario is \$5M and you're a multi-billion dollar company, you're off by an order of magnitude. On the other hand, if every scenario on your risk register produces a billion-dollar loss, you're likely catastrophizing. Either way, you're making downstream decisions like insurance limits, capital reserves, and security investments off a bad number. Set an upper and lower bound for the extreme loss scenarios that you can justify with this data.

Decision 2: "Are we focused on the right threats?": Risk registers are full of scenarios that feel big, maybe even existential, but may not reflect where extreme losses concentrate. This report gives you empirical data to pressure-test your priorities.

System intrusions and ransomware represent the highest exposure by far, and that finding holds whether you're looking at all incidents from the IRIS 2025 report or just extreme events from this one. That consistency is worth noting, because it means the threats driving everyday security operations are the same ones driving catastrophic losses. You don't need a separate "extreme event" strategy for threat prioritization. You need to make sure the threats you already know about are getting the investment they deserve.

But here's what does change at the extreme end: the type of impact shifts dramatically. Data exposure accounts for roughly 90% of all security incidents. Among extreme events, it drops below 40%. Disruptions and fraud pick up the slack, playing a much greater role. If your risk register is dominated by data breach scenarios and you haven't modeled operational disruption or large-scale fraud, you have a blind spot where extreme losses land.

Compare your current risk register against Figures 18 and 19 in this report. Look at where intrusions, ransomware, and system failure scenarios sit relative to the investment they're getting. If you're spending assessment cycles on threats that barely register in the extreme loss data while underinvesting in the scenarios that drive the biggest losses, your portfolio is misaligned with reality.

Decision 3: "Do our risk transfer mechanisms match real life exposure?": Most cyber insurance purchasing decisions are based on broker benchmarks. This report gives you the data to have a more grounded conversation.

Start with Table 1. Find your revenue band and look at the 95th percentile loss. That's the number your insurance program should be built around. If your policy limit is well below that threshold, you're covering the routine and leaving the extreme unprotected. If it's well above, you might be overpaying for coverage.

Then look at the forms of loss in Table 2. Response and notification costs are the most frequently reported form of loss, but business interruptions and lost sales account for the largest share of total losses. That difference matters because your policy may cover one but not the other, or may have sublimits that effectively cap coverage for the category that drives the biggest losses. Ask your risk team or your broker to map your current coverage against the forms of loss in this report and identify the gaps.

Finally, consider capital reserves. The "critical hit vs. death by 1,000 cuts" analysis in this report shows that for most organizations, a single extreme event represents more expected annual loss than years of smaller incidents combined. If your financial planning thinks of cyber losses as a steady operational expense rather than a low-frequency, high-severity exposure, your organization may not have the cash reserves to absorb an extreme event. That's a conversation for your CFO, and this data gives you the numbers to start it.

Wrapping Up: There are more decisions where these came from, and we hope the data in this report sparks a few we haven't thought of yet. That's the whole point of the IRIS series: give practitioners real numbers so they can make better calls. If you've read this far and you're already thinking about how your risk register, your insurance program, or your next board presentation needs to change, we've done our job. Now go do yours.

APPENDIX A: METHODOLOGY

This research draws heavily upon three datasets:

1. Zywave Cyber Loss Data
2. Feedly for Threat Intelligence
3. Board Cybersecurity Incident Tracker

Zywave's Cyber Loss Data forms the foundation of our analysis because of its breadth (over 150,000 security incidents), history (spanning decades), inclusion of financial losses, and firmographic data. The data is compiled from publicly available sources, such as breach disclosures, company filings, litigation details, and Freedom of Information Act requests. It is used extensively by cyber (re)insurers to assess risk and price premiums. More info: <https://www.advisenltd.com/data/cyber-loss-data/>

Feedly for Threat Intelligence is a platform that continually scours thousands of open trusted web sources to extract actionable intelligence for cybersecurity teams. The platform uses over 1,000 pre-trained AI models to filter out noise and identify cyber threats, including attacks, threat actors, TTPs, and IOCs. We used Feedly for Threat Intelligence in this research to enrich extreme events with contextual threat information. More info: <https://feedly.com/threat-intelligence>

Board Cybersecurity is a governance-oriented platform that aggregates and structures cyber risk information from SEC disclosures, state AG notifications, news reports, and governance filings. As of this writing, this spans over 12,000 companies and roughly 9,000 incidents. Board Cybersecurity's Incident Tracker provided the corpus of 8-K filings analyzed in this study. More info: <https://www.board-cybersecurity.com/>

Additionally, Cyentia does extensive processing of all these sources to extend and enrich the combined dataset for our research and services. This is done using a combination of classification models, natural language processing (NLP), AI-supported evidence gathering, taxonomy mapping, malware behavioral analysis, and manual research by our analysts.

While the incident and loss data derived through these sources is vast, we're not claiming that it is all-inclusive. We can only analyze incidents that make their way into the public record through outward signs or impacts, mandatory reporting, voluntary disclosure, etc. That's not all the events that occurred over the timeframe, of course, but we have high confidence that significant cyber events, like those we focus on in this study, are well represented.

APPENDIX B: INCIDENT PATTERNS

All security incidents in our historical dataset are assigned one of the following patterns using a combination of natural language processing techniques and human expert assessment.

- 1. Accidental Disclosure:** Data stores that are inadvertently left accessible to unauthorized parties, typically through misconfigurations on the part of the data custodian.
- 2. DoS Attack:** Any attack intended to render online systems, applications, or networks unavailable, typically by consuming processing or bandwidth resources.
- 3. Defacement:** Any unauthorized content modification to an organization's website or other online assets.
- 4. Fraud or Scam:** Any incident that primarily employs various forms of deception to defraud the victim of money, property, identity, information, and so on.
- 5. Insider Misuse:** Inappropriate use of privileged access, either by an organization's own employees and contractors or a trusted third party.
- 6. Physical Threats:** Threats that occur via a physical vector, such as device tampering, snooping, theft, loss, sabotage, and assault.
- 7. Ransomware:** A broad family of malware that seeks to encrypt data with the promise to unlock upon payment or seeks to completely eradicate data/systems without the pretense of collecting payment.
- 8. System Failure:** All unintentional service disruptions resulting from system, application, or network malfunctions or environmental hazards.
- 9. System Intrusion:** All attempts to compromise systems, applications, or networks by subverting logical access controls, elevating privileges, deploying malware, and so on.



THE CYENTIA INSTITUTE IS A WIDELY-RESPECTED, RESEARCH AND DATA SCIENCE FIRM
WORKING TO ADVANCE CYBERSECURITY KNOWLEDGE AND PRACTICE.

We accomplish that goal through collaborative research publications like the
IRIS series and analytic services that help our clients manage cyber risk.

Visit cyentia.com for more information.