Cyentia INSTITUTE

# Information Risk Insights Study

## *It's About Time*

# Introduction

> "Time isn't a straight line... It's all bumpy wumpy."[1]
>
> ~The Eleventh Doctor

Welcome to the 2025 edition of the (roughly) biennial Information Risk Insights Study (IRIS). The last one was in 2022, so it's about time we got this to you. Thanks for your patience.

Fittingly, time is of the essence in this IRIS. Not just because it's a tad overdue, but because it's literally about time—cyber risk trends over time, to be specific.

Cybersecurity is ever-changing, and there's an implicit assumption that risk is always increasing. But is it?

Are cyber events occurring at greater frequency? Is an organization more likely to have a breach now than 15 years ago? Which types of incidents have become more common over time? Have the financial impacts of cyber events increased or decreased? Are risk factors trending the same way for all sectors and sizes of organizations?

We explore these questions and more by analyzing a huge historical dataset of cyber events and losses from 2008 through 2024. As always, our goal is to dispel the fog of FUD surrounding cyber risk so you can see it more clearly and manage it more effectively. Thanks for reading!


## Acknowledgements

The Cyentia Institute wishes to acknowledge and thank the Cybersecurity Division and the Office of the Chief Economist at the Cybersecurity and Infrastructure Security Agency (CISA) for sponsoring this study. It is our sincere hope that this research will aid community efforts to manage cyber risk.


## TABLE OF CONTENTS



The Cyentia Institute is a research firm working to improve cyber risk management through our analytical services and data-driven research publications. You can download the IRIS 2025 and find related content at www.cyentia.com/iris.

# IRIS 2025

## KEY FINDINGS

The IRIS research draws heavily upon Zywave's (formerly Advisen) Cyber Loss Data, which contains over 150,000 security incidents[1] and associated financial losses[3] spanning decades. The data is compiled from publicly available sources, such as breach disclosures, public company filings, litigation details, and Freedom of Information Act requests.

It is the most comprehensive source of cybersecurity incidents and losses available. Additionally, Cyentia does extensive processing of this base dataset to extend and enrich it for cyber risk analysis use cases.

**On average, 3,000 significant security incidents are publicly reported or discovered each quarter. That's a 650% increase over the last 15 years.**

**Cyber events affecting smaller businesses are far more common overall, but relative to population size, the rate among the largest corporations is 620 times higher.**

**The annual probability of any given organization experiencing a cyber event has almost quadrupled since 2008.**

**The probability of a <$1B firm suffering an incident has more than doubled, while the annual likelihood for a $100B+ organization has fallen 50%.**

**Losses from a typical security incident have absolutely exploded, rising 15-fold from a median of $190K to almost $3 million!**

**The cost of more extreme "tail loss" events is also up 5-fold, ballooning to $32 million.**

**Cyber events aren't just costing more—they're hurting the bottom line more than ever before. We've seen an 8-fold increase in costs as a proportion of annual revenue.**

**Median losses for professional services firms are up 25x over the last 15 years! Alternatively, there's been a huge decrease in loss magnitude among retailers.**

**Compromising user credentials remains the most common intrusion technique over the last decade, fluctuating between 43% and 60% of all incidents.**

**Exploitation of web applications is up 6x for smaller firms, while targeting third-party relationships has doubled among large organizations.**

## 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.

119 Cy cyentia

# Q1 ARE SECURITY INCIDENTS BECOMING MORE COMMON?

To many of you, the answer to this question seems so obvious that it's hardly worth asking. But we're not ones to let any assumption go unchallenged. As it turns out, this one is solidly backed by historical data—at least in terms of reported incidents[4]. Figure 1 shows a 650% increase in the average number of incidents added to the public record each quarter in 2024 (~3,000) versus the rate set 15 years ago (~450).

But there's a lot more going on than simply "incidents are way up!"[5] The proliferation of large-scale data breaches combined with the rolling out of breach disclosure laws certainly drove the steady climb early in this timeframe. The plateau beginning in 2013 corresponds with the emergence of advanced persistent threats (APTs)[6] that employed a "low and slow" rather than "smash and grab" strategy. The reacceleration circa 2019 was spurred by the rapid rise of ransomware (see Figure 2) and exacerbated by the COVID-19 pandemic. We could go on, but you get the point. These trends have reasons.

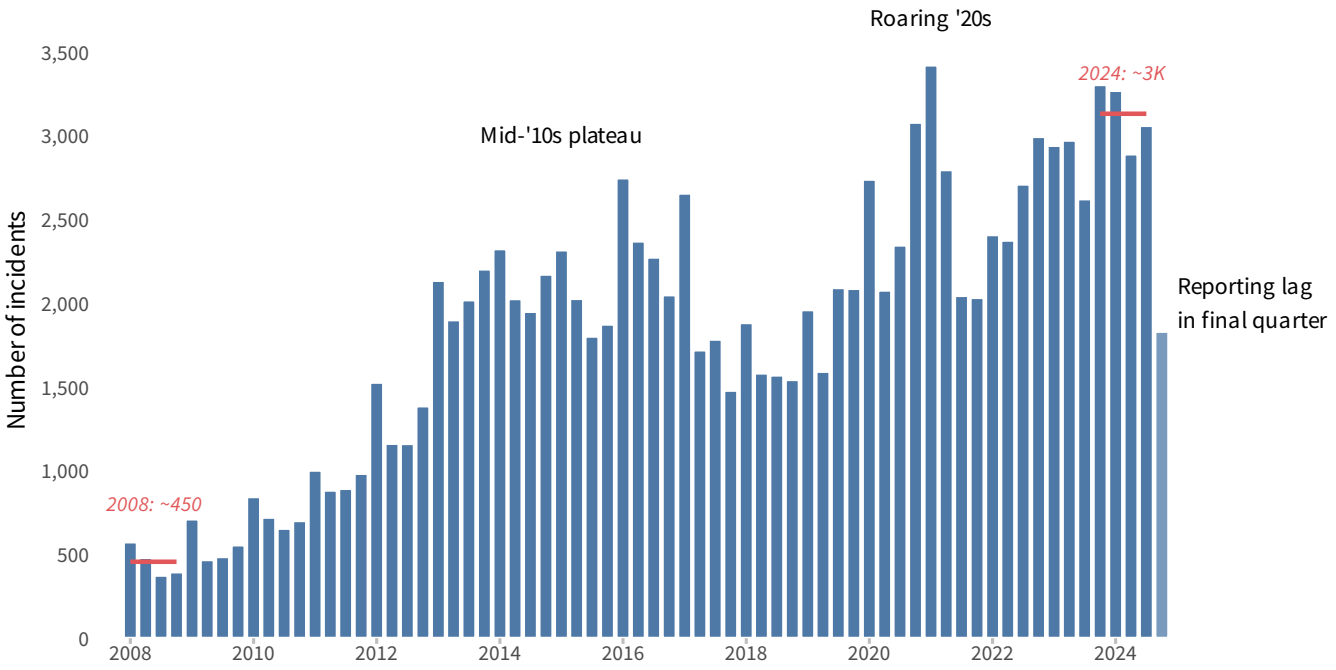


*Source: IRIS 2025 (Cyentia Institute)*

Figure 1: Number of security incidents publicly reported or discovered each quarter

1 "Actually, from a non-linear, non-subjective viewpoint, it's more like a big ball of wibbly-wobbly, timey-wimey stuff." - The Tenth Doctor

2 We use the terms security incident, cyber event, and loss event interchangeably. This refers to actual incidents that compromised the confidentiality, integrity, or availability of a firm's information assets.

3 We use the terms losses or costs to refer to the financial consequences of incidents.

4 This entire report is based on analyzing incidents that make their way into the public record through outward signs or impacts, mandatory reporting, voluntary disclosure, company filings, public lawsuits, etc.

5 Yes—we're aware that the "incidents are way down" in the last quarter of 2024. But we're fairly confident that number will go up once the reporting lag catches up and flushes all those as-yet-unknown events into the open. Spoiler alert: we test (and confirm) this in Q7.

6 To be clear, we're not saying APT attacks started in 2013. But that's when Mandiant's APT1 report published and community awareness of these events ballooned. This slowed the rate of publicly reported incidents because attackers (even cybercriminals) were more discrete and much of the threat intel and incident response community was focused on APTs rather than standard cybercriminals.

We just mentioned the rise of ransomware, which prompts a related question: Are all types of incidents trending the same way? The crisscrossed lines in Figure 2 are sufficient for a definitive "nope," but let's highlight some of these trends that meaningfully impact organizations' security strategies.

Are all types of incidents following the same trend?

The data in our analysis is clear:

**"NO"**

At the top of Figure 2, below, system intrusion (unauthorized access to systems, applications, or networks) has long reigned supreme among incident patterns[7]. The particular techniques attackers use to infiltrate networks and systems have undoubtedly changed, but we'll dig into that later (see Q6). For now, simply observe that the most common category of incident experienced by organizations hasn't really changed in the last 15 years.

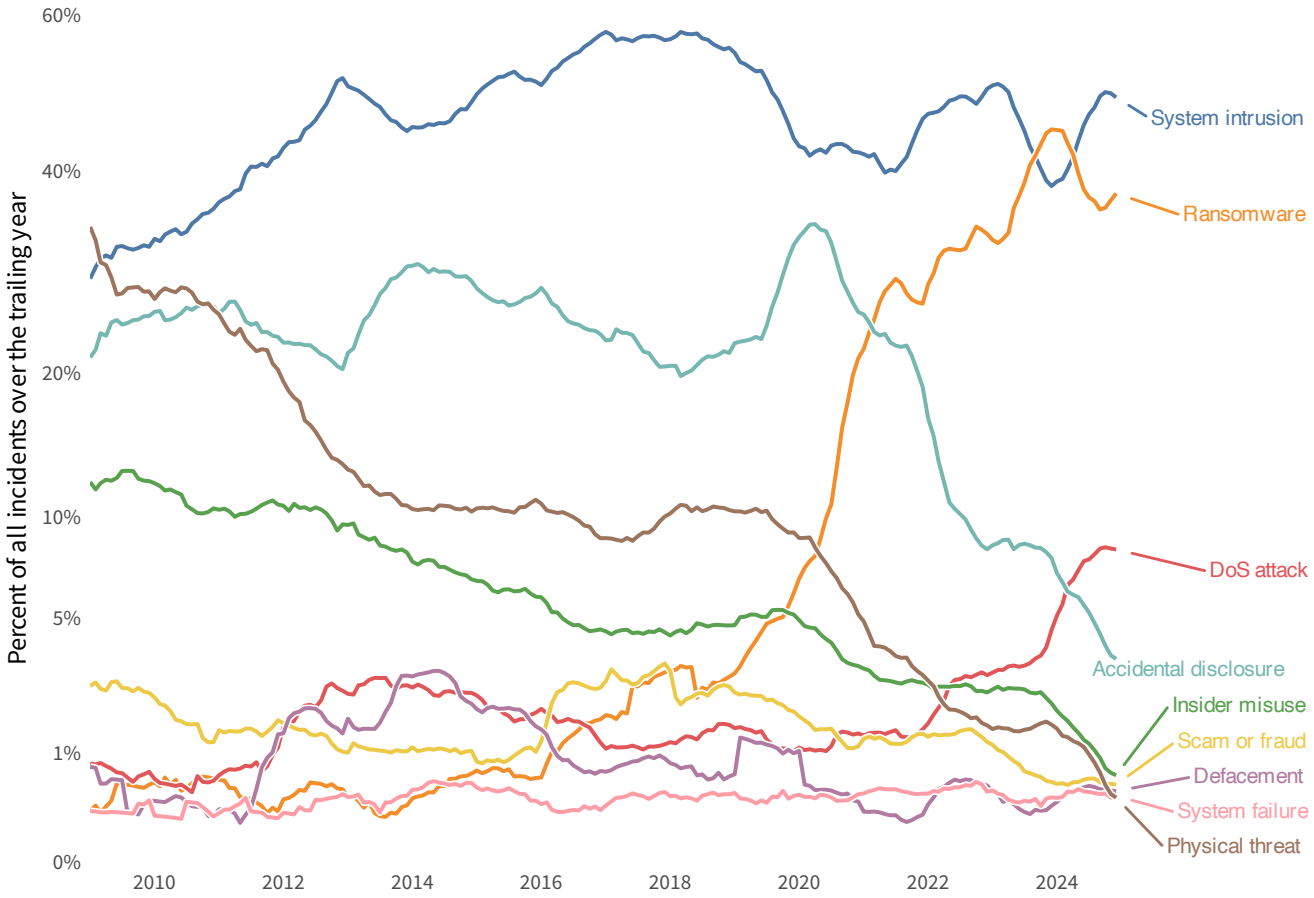


*Source: IRIS 2025 (Cyentia Institute)*

Figure 2: Relative frequency of incident patterns over time (rolling 12mo. 2009-2024)

Now let's also observe what has changed. The aforementioned rise of ransomware in recent years is, in fact, unprecedented. So is accidental disclosure's precipitous drop over roughly the same period. We could go into far more detail on ransomware trends, but we've already done that in another IRIS.

7 See Appendix 2 for definitions of these incident patterns.

We'd love to think that "oopsies" as a major cause of data disclosure are a thing of the past, but we suspect human nature will reassert itself at some point.

Physical threats and insider misuse show a marked downward trend over the years. Compared to more scalable remote alternatives, the "hands-on" approach to data theft has fallen out of favor. Data handling regulations and endpoint protections—such as encryption at rest—have further contributed to its decline. Insider misuse never rises above fourth place among all incident patterns, which goes against the long-standing "employees are the enemy" mentality. Sure, employees are often targeted in cyberattacks, but they're usually not acting with malicious intent.

## Key Risk Insight

If it seems like a lot more incidents are happening these days, it's not just recency bias.

The overall rate has seen more than a sixfold increase over the last 15 years.

> *The data shows how fluid and contextual the cyber threat landscape really is and how important your firmographic footprint is to that, as we will show throughout this report.*
>
> *Quantifying that risk, especially at the board level, means understanding these patterns as time sensitive, not timeless.*
>
> *Today's dominant risk may be tomorrow's footnote, and cyber risk models need to keep pace.*
>
> *Further, if your security strategy isn't recalibrating with these changes in risk, you're planning for a past that no longer exists.*

***~ Jack Freund***
*Executive Fellow | The Cyentia Insitute*

# Q2 DO INCIDENT TRENDS DIFFER ACROSS ORGANIZATIONS?

This seems like another obvious answer on the surface because it's well known that certain organizations make more attractive targets, some have poor defenses, and others are just plain unlucky. But what we're really after here is whether there are inherent differences between different types of firms. Figure 3 attempts to open the door to that question by comparing trends across organizations grouped by their annual revenue.

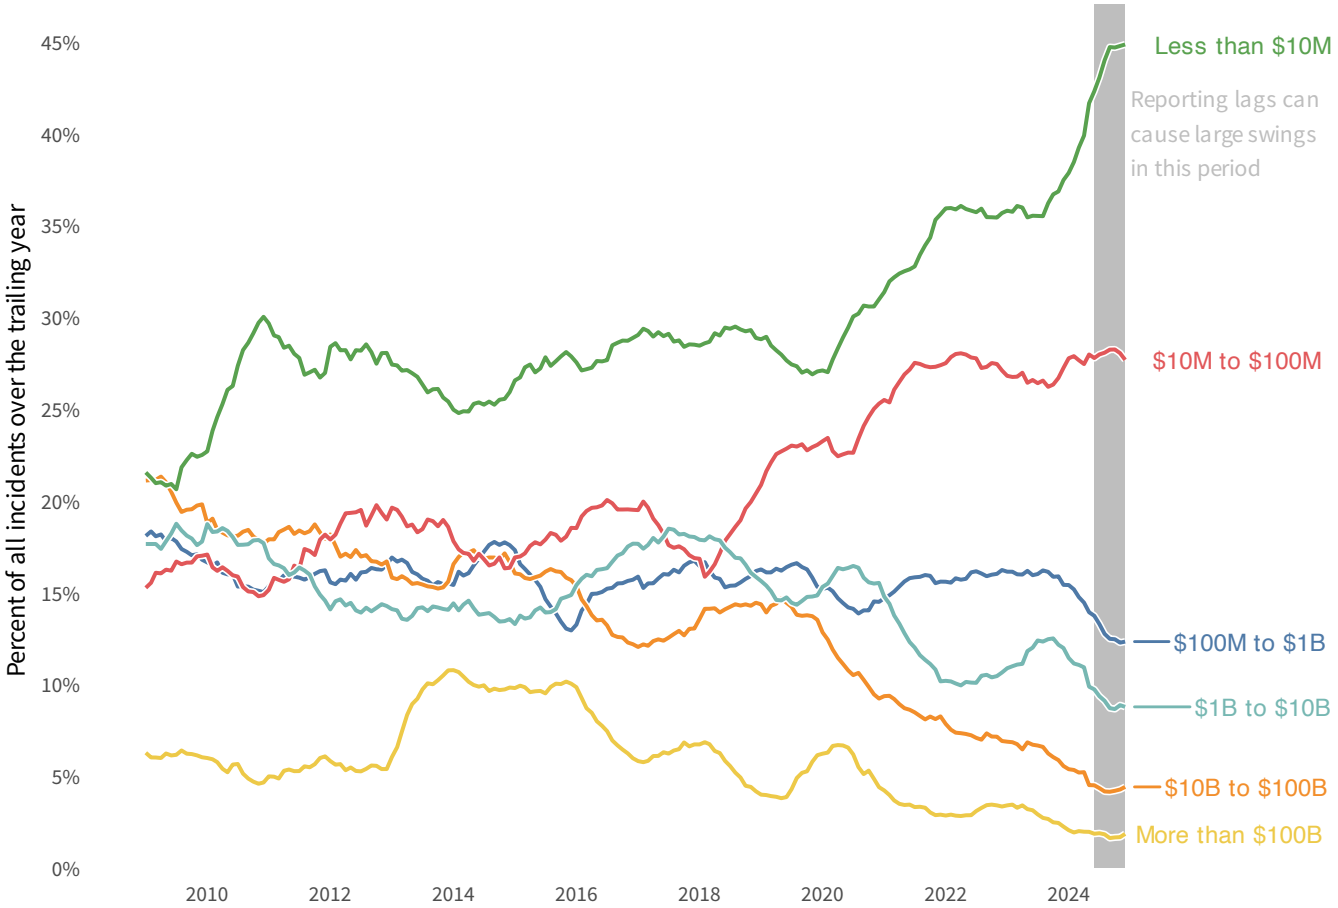


*Source: IRIS 2025 (Cyentia Institute)*

Figure 3: Proportion of all incidents in each revenue tier (rolling 12mo. 2009-2024)

Rebutting the "Who would attack little 'ol us?" argument, smaller businesses (<$100M annual revenue) see the biggest overall share of incidents. What's more, that share is growing over time. The proportion of events affecting larger organizations (>$1B annual revenue), on the other hand, appears to be declining over the last 15 years.

> The data in our analysis is clear: smaller businesses—those under $100M in annual revenue—account for the largest share of incidents, countering the idea that they're too minor to target.

There's more to this story, however, as astute readers have probably already discerned. The obvious objection to the prior chart's depiction of trends is that it does not account for the number of firms that exist in each revenue tier[8]. Sure, more incidents affect small businesses, but they vastly outnumber large corporations.[9] What happens when we factor in the relative number of firms in each group? Figure 4 gives the answer—a complete reversal of fortune!

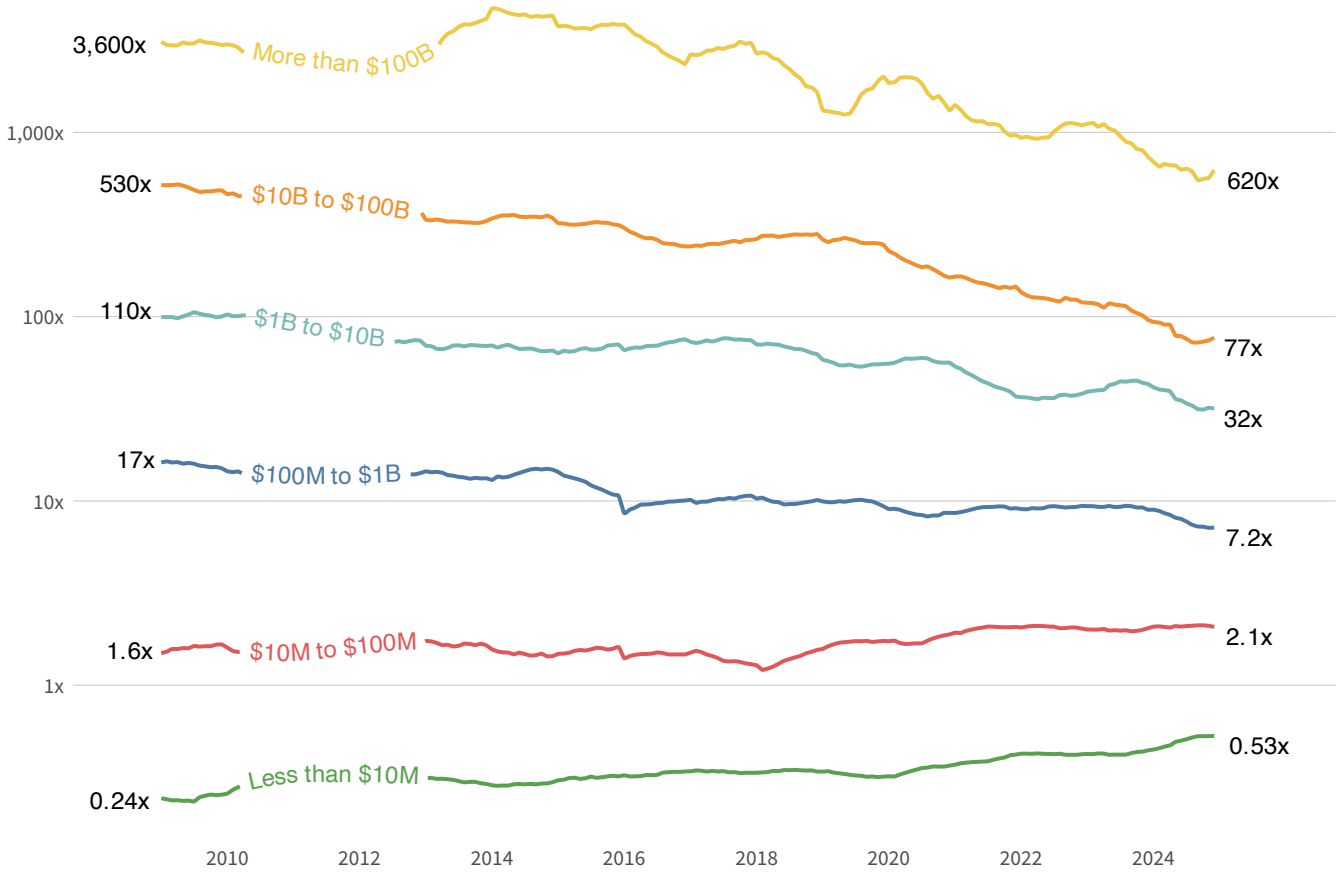


*Source: IRIS 2025 (Cyentia Institute)*

Figure 4: Relative number of incidents to number of firms in each revenue tier (rolling 12mo. 2009-2024)

8 We use data from Dun & Bradstreet for the number of organizations in each revenue tier.

9 The Small Business Administration estimates that 99.9% of all businesses are small. https://advocacy.sba.gov/2024/07/23/frequently-asked-questions-about-small-business-2024/

The multiples shown in Figure 4 compare the number of incidents across a revenue tier with the number of organizations within it. The higher the multiple, the higher the average rate of incidents per organization in each tier. While larger organizations show a declining trend in the relative number of incidents, they remain disproportionately affected by them. The $100B+ tier has experienced 620 times more incidents than the number of megacorporations in this segment. Though the smallest firms experience the largest number of incidents in absolute terms, only a fraction of them (0.53x) are actually affected.

Let's turn next to frequency-based disparities among different industries. Since we've established the importance of adjusting for the number of firms in each segment, we can skip to the punch line. Figure 5 groups sectors[10] based on their relative event frequency.[11]

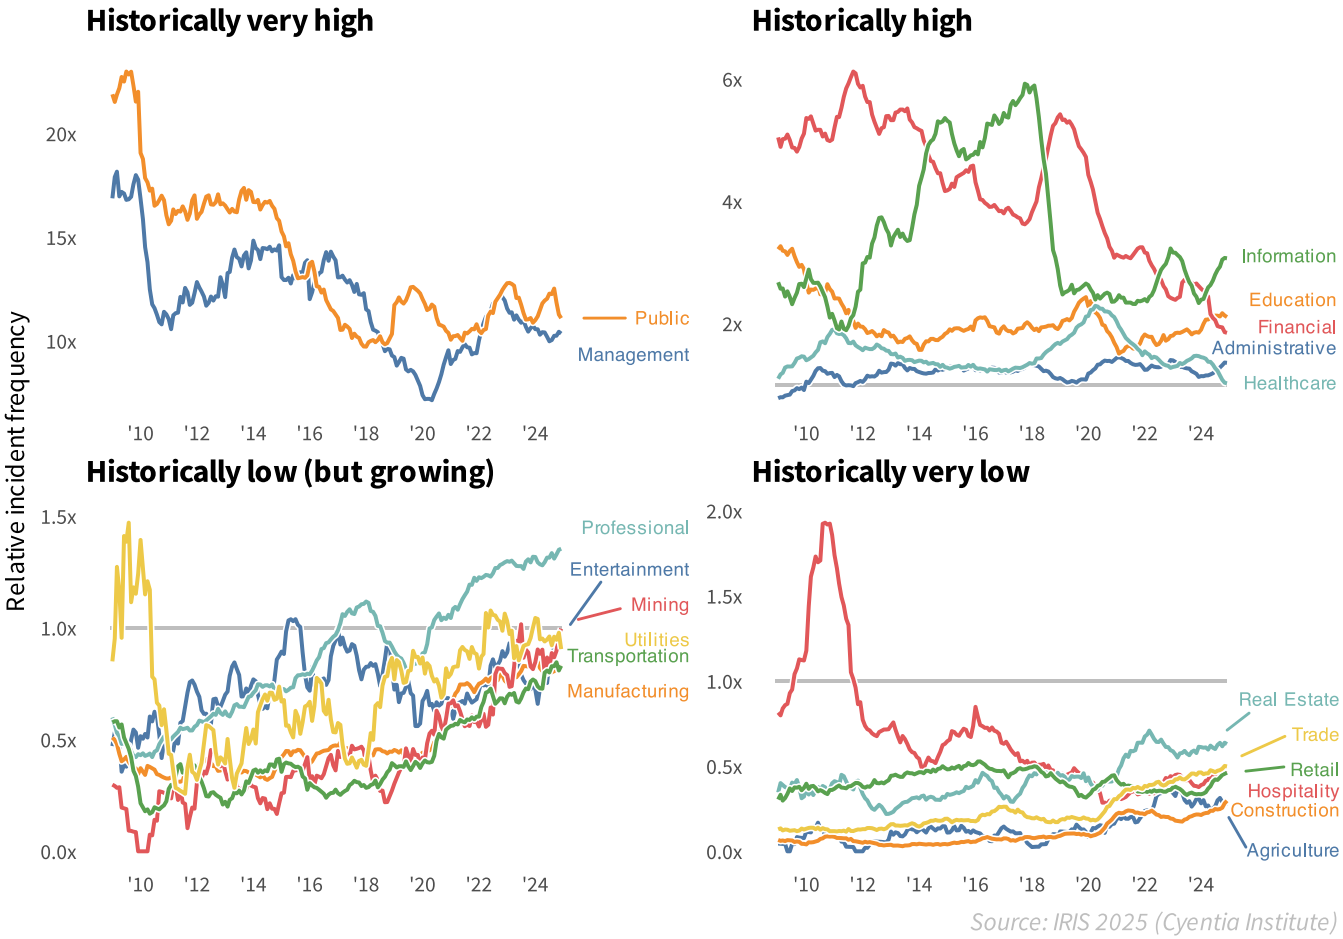


Figure 5: Relative number of incidents to number of firms in each sector (rolling 12mo. 2009-2024)

The Public and Management sectors are the only two that have historically exhibited a very high relative incident frequency. For the former, we attribute that to mandatory disclosure requirements that typically exceed those in the private sector.

10 Sectors throughout this report use the North American Industry Classification System (NAICS). Our labels are short versions of NAICS sectors—generally the first word of the official sector name.

11 A multiple >1 indicates higher relative incident frequency based on the number of firms in a sector; <1 indicates the opposite (low relative frequency).

The Management sector is a bit of an oddball in NAICS, consisting mainly of holding companies. We suspect part of what's going on here is that incidents affecting their subsidiaries are being attributed to them as the parent entity.

Moving to the upper-right panel, the Finance sector has historically seen a high share of incidents relative to the number of firms that exist. But that rate has fallen over time, perhaps due in part to the industry's outsized security budgets. The Information sector continues to experience an elevated incident rate, yet is currently well below its high-water mark. Together, these industries control money and data flowing through the economy, so it's no surprise they receive more than their fair share of cyberattacks.

Energy and supply chain sectors are creeping up in incident frequency—Utilities, Mining, Manufacturing, and Transportation are no longer sitting safely below the line.

Industries with historically low relative incident frequencies are split into two groups—those likely to remain low for the foreseeable future and those that will soon cross over the line of demarcation. We find it unsettling to see energy and supply chain sectors such as Utilities, Mining, Manufacturing, and Transportation (which includes oil and gas pipelines in NAICS) increasing in relative frequency. The Professional sector has already crossed that line, which is quite concerning given that they offer advice and services to the rest of us.

## Key Risk Insight

Incidents involving small and midsize businesses (SMBs) are far more common overall, but the relative incident frequency among large enterprises is much higher.

# Q3 IS THE PROBABILITY OF INCIDENTS INCREASING?

This question may initially sound similar to the previous two, but those involved tallying the total number of incidents reported across many organizations. Here, we explore how the likelihood of a single organization having an incident is changing over time. If that nuance isn't quite clear, think of it like this: what are the chances your firm will suffer a significant incident this year?

We'll begin by changing our perspective from the past to the future—except how the future was modeled in the past… over time. Jeepers, this timey-wimey stuff is confusing, isn't it? Maybe a chart will help; Figure 6 tracks the modeled probability[12] of a typical organization[13] experiencing an incident in the next 12 months.

It is understandable if cybersecurity folks can't hold back an "I told you so!" here because overall incident probability has almost quadrupled over the last 15 years. We could stop there, issue a press release, and bid you adieu until the next installment, but we're just getting started.

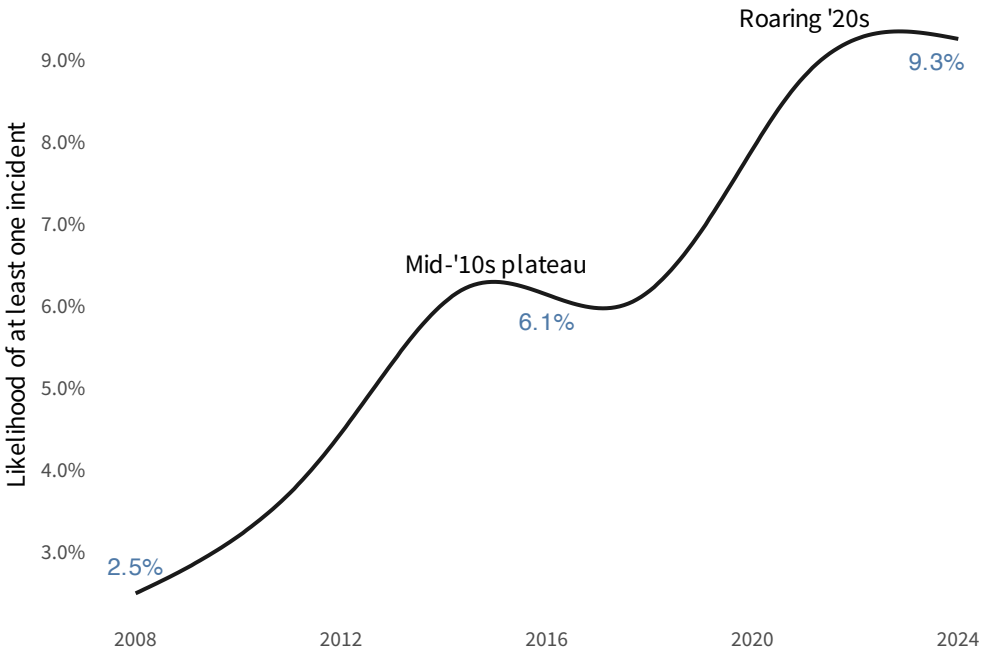


*Source: IRIS 2025 (Cyentia Institute)*

Figure 6: Historical probability of a firm having an incident in the next year

The probability that a typical firm will experience a significant security incident has almost quadrupled over the last 15 years.

12 See appendix for details on our approach to modeling annualized incident probability.

13 We use "typical" to remind readers that this model doesn't account for the many factors that would make incidents more or less likely for a particular organization. We'll look at some of those later.

**"But wait,"** we hear you saying, **"doesn't the probability for different types of cyber events change over time?"**

You're not wrong. We simply can't cram everything into this one study.

**New studies** are always in the pipeline — we regularly publish **extra analysis** like that on our website at **www.cyentia.com/iris**.

What if we told you that the probability of a <$100M firm suffering a security incident has more than doubled, while the chance of a $100B+ megacorporation suffering an incident has dropped by a third over the same time frame? Well, that's exactly what Figure 7 tells us.

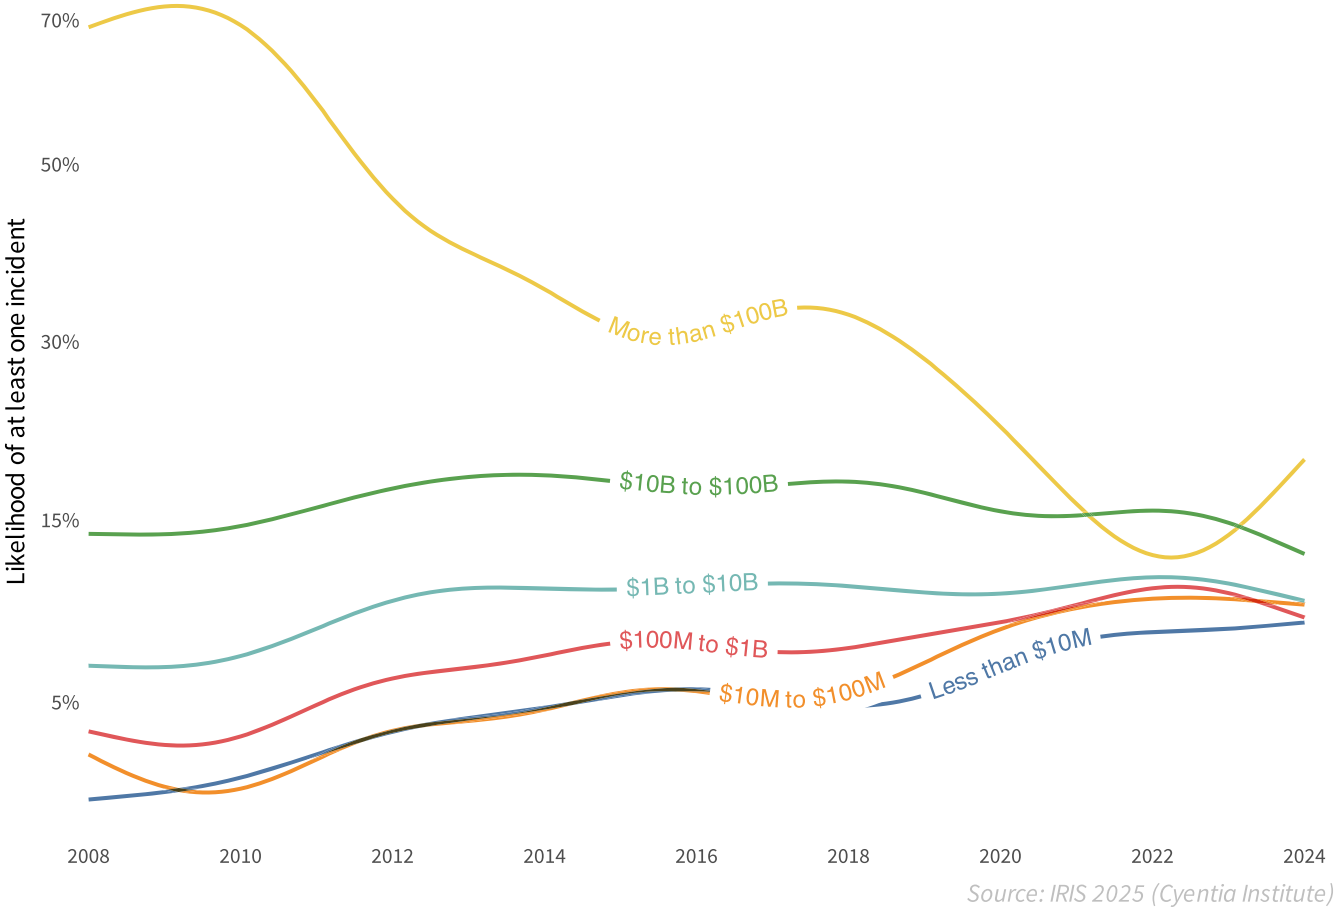


Figure 7: Annualized incident probability for firms in each revenue tier (rolling 12mo.)

Unfortunately, our dataset is silent on the underlying factors behind these trends, so all we can offer is some speculation. Perhaps cybercriminals have shifted to more volume-oriented ("low-hanging fruit") strategies over time. Maybe the pace of digitalization has outpaced SMBs' ability to defend their growing attack surfaces, while the bigger enterprise security budgets offset that. Maybe increased regulatory pressures on large corporations are gradually hardening enterprise security architectures. Whatever the cause(s), these are important trends that are worth further research by our industry.

So, an organization's size matters when evaluating the likelihood of incidents. Now, let's see what happens when we treat industry as a feature of interest. Figure 8 paints that picture.

Allow us to briefly describe what you're looking at. Sectors are sorted in descending order by the latest probability estimate. So, a typical manufacturing firm has an 11% chance of having a security incident in the next 12 months—up from ~2% 15 years ago.

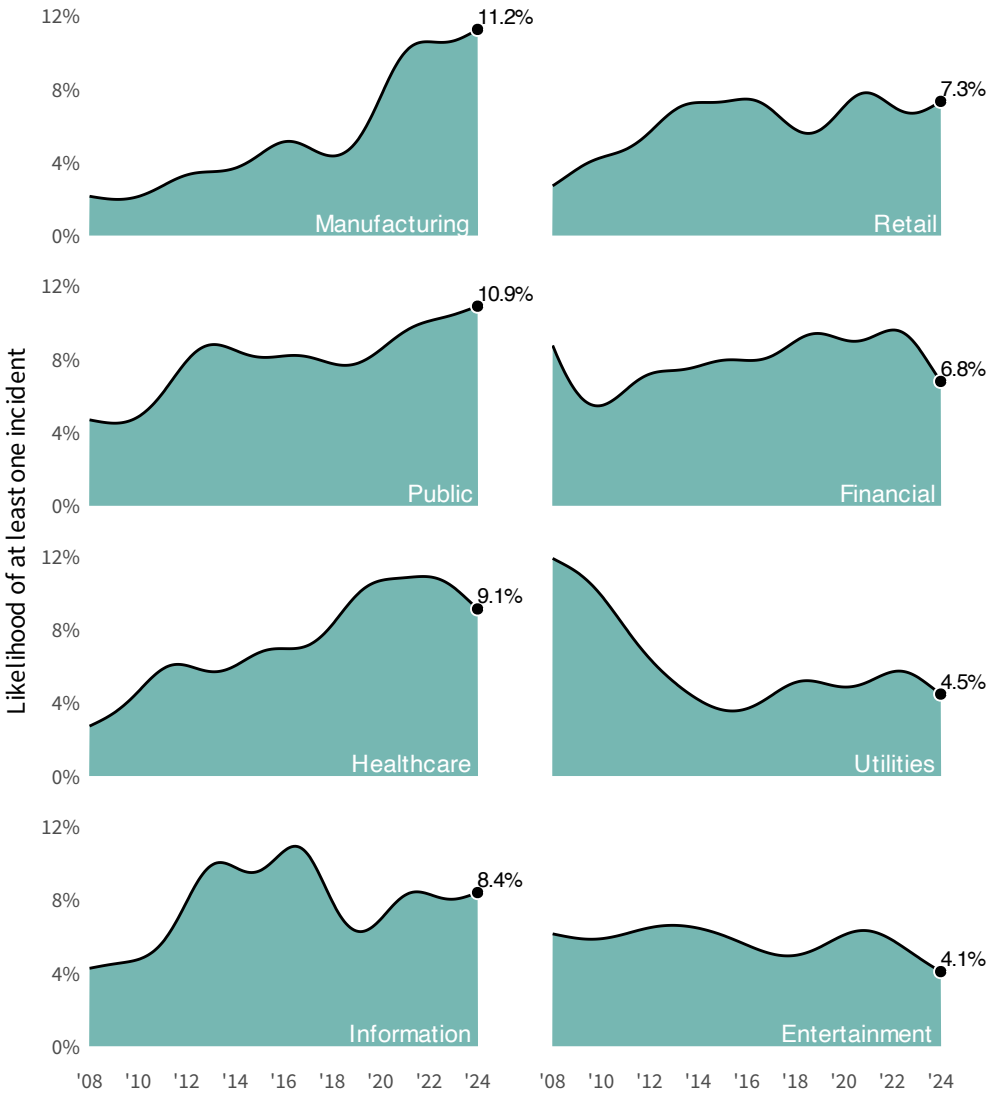


*Source: IRIS 2025 (Cyentia Institute)*

We could spend oodles of time combing through historical evidence behind the peaks and troughs for certain industries, but we'll leave that to eager readers. Suffice it to say that incident probability trends can be significantly different depending on firmographics (which are reflective of evolving business models, changes in the threat landscape, shifting adversary goals, etc.). That's intuitive, but perhaps seeing this confirmed will help validate the need to incorporate such factors into your cyber risk assessments.

Figure 8: Annualized incident probability for firms in each sector (rolling 12mo.)

Sorry to disappoint if you were hoping to see updated versions of the old school IRIS charts/tables for incident likelihood by sector and revenue tier. Since this version of the IRIS focuses on trends over time, Figures 7 and 8 replaced those. But we recreated some of the key figures and stuck them in Appendix 3 as a thank you to our loyal readers.

## Key Risk Insight

Overall, the chances of any given organization experiencing an incident have gone up.

But that trend has flattened or even reversed in some sectors and size tiers.

*Our analysis in this section focuses on the likelihood of experiencing at least one security incident within a year. It is possible, of course, for organizations to suffer multiple incidents, and veteran IRIS readers may recall that we've supplied probability tables for two, three, or five incidents in a 12-month period.*

*We decided not to include that in this edition because a) it's already a long report, and b) feedback suggested most cyber risk models focused on single-event likelihood. However, we have not abandoned that concept and will continue that research outside of this study.*

*Visit www.cyentia.com/iris to get additional analysis of multi-incident probabilities.*

# Q4 HAVE SECURITY INCIDENTS GOTTEN MORE COSTLY?

We've covered the evolving frequency and likelihood of cyber events—now it's time to talk dollars and cents. Let's begin by establishing the distribution of financial losses from cyber events using Figure 9, which reproduces a classic IRIS chart with the latest and greatest data.[14]

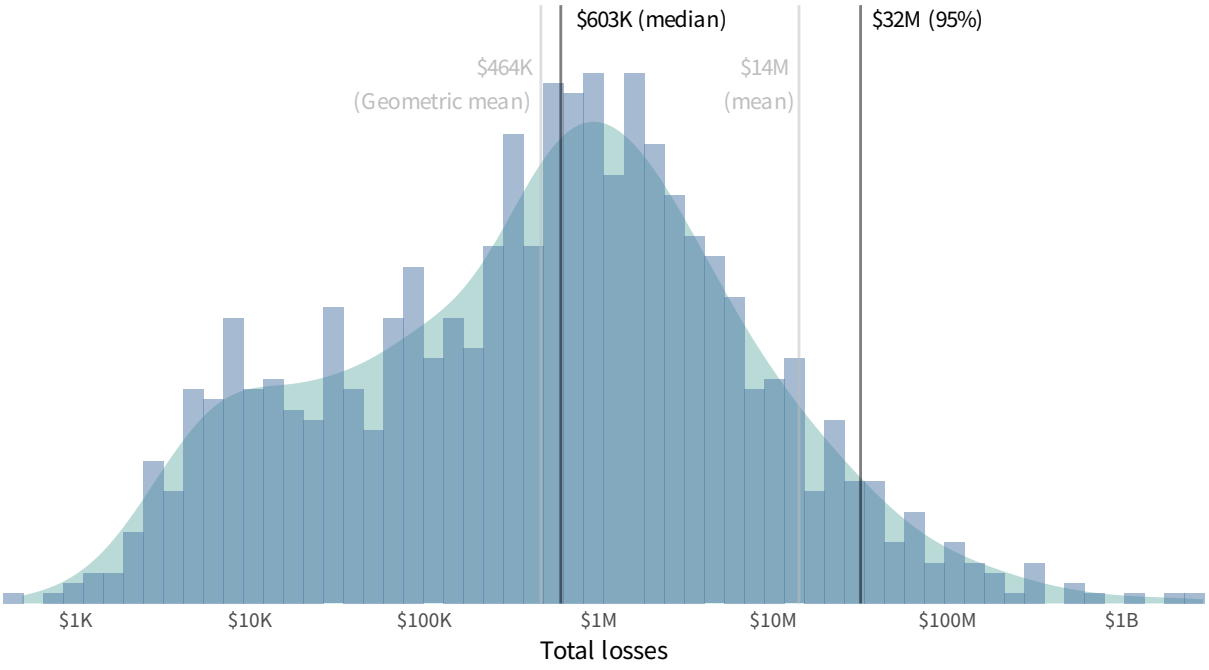


*Source: IRIS 2025 (Cyentia Institute)*

Figure 9: Distribution of reported losses for security incidents from 2015 to 2024

The typical (median[15]) incident costs about $600K, while more extreme (95th percentile) losses swell to $32 million. Note that Figure 9 is plotted on a log scale, so the tail of large losses is longer than it appears. If it's not too much to "shoulder," also note the bump in the lower half of the distribution. We'll explore that later.

NOTE: Losses analyzed in this study tend to reflect direct losses that are easier to quantify (e.g., response costs or lost revenue) and/or 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, this represents a conservative view of financial losses associated with cyber events.

14 All loss amounts considered in this report have been converted to 2024 dollars to adjust for inflation.

15 Prior IRIS used the geometric mean for a typical loss. Since the growing "shoulder" in lower part of the distribution pulls the geomean down, the median is better central measure for the updated distribution.

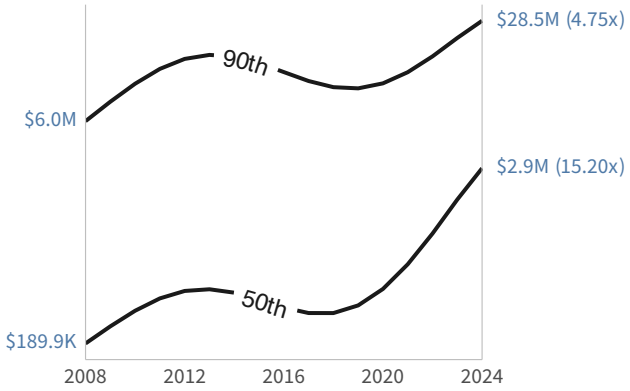


*Source: IRIS 2025 (Cyentia Institute)*

Figure 10: Trend analysis of median and 90th percentile losses

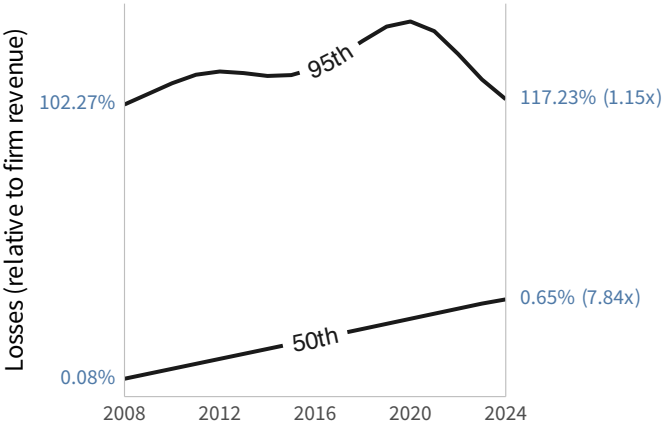


*Source: IRIS 2025 (Cyentia Institute)*

Figure 11: Trend analysis of median and 95th percentile losses as a percent of revenue

| | Loss percentile | |
|---|---|---|
| **Revenue** | **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)*

Table 1: Loss statistics by revenue tier

An overall distribution like Figure 9 is useful for illuminating the big picture, but it obscures any shifts in losses that may be happening over time. To visualize how the costs of cyber events have evolved, we'll track two reference points—the median and 90th percentile—during the 15-year timeframe. The results are shown in Figure 10.

Indeed, moving from a static distribution to a more dynamic view reveals some major shifts. Median losses from a security incident have absolutely exploded over the last 15 years, rising 15-fold from $190K to almost $3 million! The cost of extreme events[16] at the upper end of the distribution has also risen substantially (~5x). So, yeah—the financial toll of cyber events is definitely getting bigger.

Having seen that, we imagine you're anxious to see how these costs are trending for firms like yours. While we can't drill down quite that far, we can show these loss distribution parameters for organizations of different types and sizes. We'll start with the latter.

It's a no-brainer that larger firms would experience larger losses, and we've shown that in prior studies. We include Table 1 to reconfirm that fact and arm you with the most recent stats to inform loss estimates that are better sized to your organization. Note that the size-based differences in loss magnitude are especially prominent for extreme (95th percentile) events. There's a longer tail for larger organizations.

That being said, a million-dollar loss is different for a mom-and-pop shop versus a multinational megacorporation. We need a way to normalize the relative impact to the firm. Measuring losses as a percentage of the victim firm's annual revenue works well for this purpose.

Median losses fall well below 1% of revenue for the majority of incidents. But as Figure 11 attests, that ratio has grown by nearly 8x over the last 15 years. The top 5% of loss events continue to exceed the annual revenue of affected firms.

16 You may notice that we switch between the 95th and 90th percentile for losses to represent the concept of extreme events. We generally use the 95th for long, fixed time periods. However, we found that in the specific context of time series models, data availability in the tails wax & wane such that estimates of the 95th percentile became unreliable. Using the 90th percentile is merely a small concession in consistency to ensure that we're confident that the percentile is being estimated reliably.

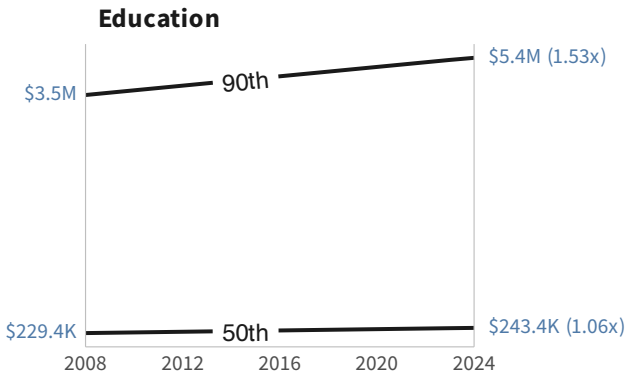


*Source: IRIS 2025 (Cyentia Institute)*

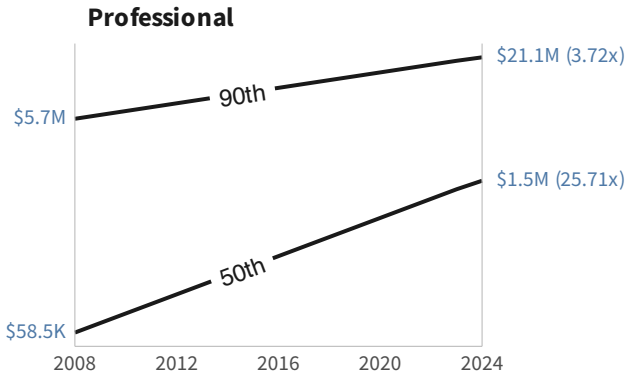


*Source: IRIS 2025 (Cyentia Institute)*

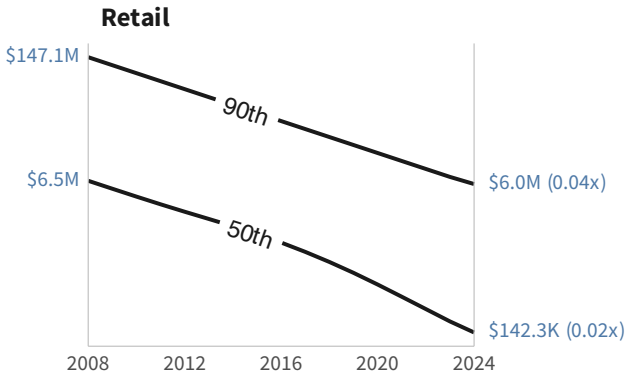


*Source: IRIS 2025 (Cyentia Institute)*

Figure 12: Trend analysis of median and 90th percentile event losses for example sectors.

We'll now briefly demonstrate that industry makes a difference too. Rather than a plot crammed with all 20 NAICS sectors, we've done some pre-screening to highlight three industries that represent distinct rising, flat, and falling trends that we see across all sectors.[17]

When it comes to the escalating costs of security incidents, no industry has been hit as hard as Professional Services. Median losses for firms in that sector are up 25x over the last 15 years! Top end (90th percentile) costs have seen a nearly four-fold increase too. The Administrative, Financial, Healthcare, Manufacturing, and Public sectors also exhibit increasing trends for loss magnitude.

Event losses in the Education sector show a fairly steady, albeit increasing, trajectory over the years. Perhaps a side effect of the "ivory tower" insulating it from cyber risk trends experienced by the rest of the world?

And then there's the Retail sector, which seems to have figured out a way to cut the price tag of a security incident. Current losses for both typical and extreme events are a fraction of their starting points in 2008.

That deserves an entire study of its own, but we suspect the push for Payment Card Industry (PCI) compliance and the rollout of Chip-and-Pin technology may have helped to limit the amount of data that can be easily exfiltrated from retailers. Since larger breaches generally (but not linearly[18]) correspond with higher losses, this would help cap potential losses. But Retail is not alone; Information Services and Management firms also show declining loss distributions.

17 Don't fret if your favorite sector is missing. Appendix 2 includes a table that gives median and extreme loss values for all sectors similar to the IRIS 2022.

18 Prior IRIS have conclusively demonstrated that losses do not follow a flat cost-per-record formula.

For a seemingly simple leading question, this section contains many "Yes, if..." and "No, but..." answers. A quick recap of what we've learned would be helpful:

- Overall, security incidents have indeed gotten more costly.

- Losses relative to the affected firm's annual revenue have also grown.

- The absolute cost of incidents is higher in large organizations, but the relative impact is worse for smaller businesses.

The magnitude and directionality of loss trends differ substantially by sector.

## Key Risk Insight

Not only do typical security incidents cost more these days (up 15x since 2008)—they hurt a lot more too (up 8x relative to annual revenue).

*Don't see your sector listed here among the three examples we chose?*

***Don't despair because we've done the analysis!***

*We just couldn't squeeze all the charts into the pages of this study.*

*You can get a version of these cyber loss trendlines for your sector from the IRIS page on our website at www.cyentia.com/iris.*

***Sectors Include:***

- *Financial*
- *Healthcare*
- *Hospitality*
- *Transportation*
- *And more!*

# Q5 DO LOSS TRENDS DIFFER AMONG EVENT TYPES?

Remember how we said to pay attention to the overall shape of the distribution in Figure 9 at the beginning of the last section? If not, feel free to jump back and refresh your memory. We'll wait.

We weren't pulling one of your lower extremities when we said that would be important. There is a pronounced "shoulder" in the lower half of the distribution. Whenever a bimodal tendency like this exists in a distribution, it's worth investigating what's behind it. So, we did—and discovered that the shoulder has grown over time, as evidenced by Figure 13.

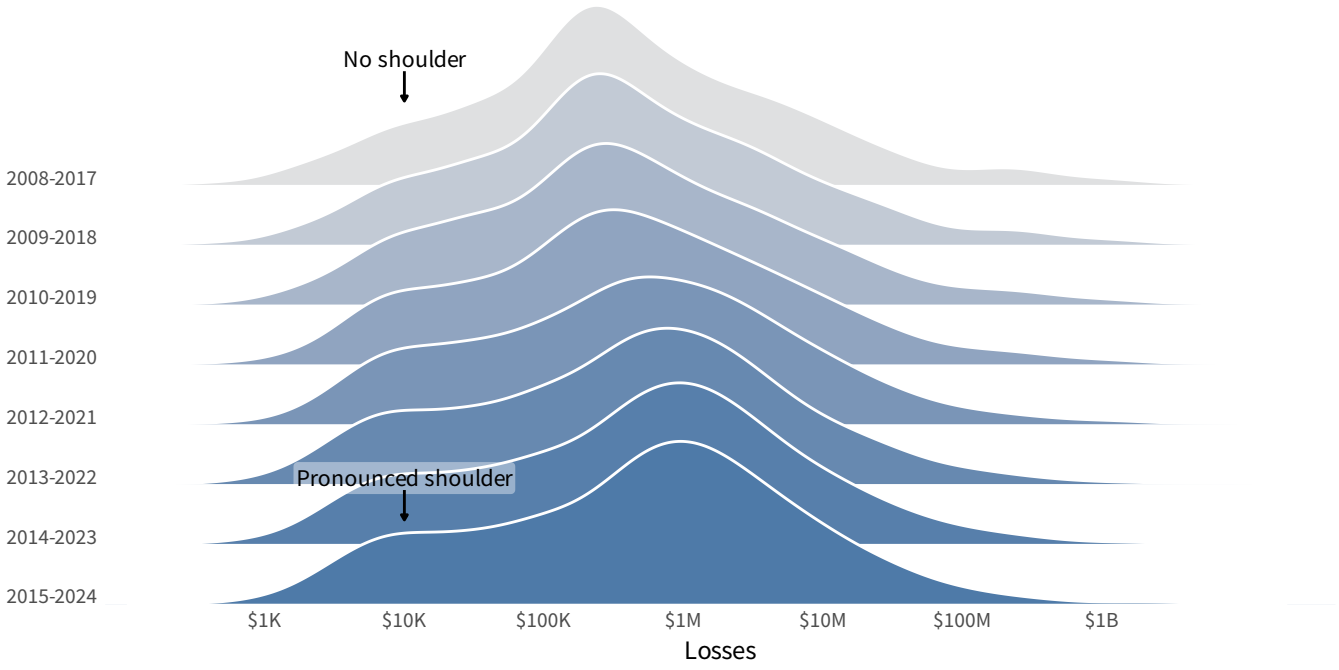


*Source: IRIS 2025 (Cyentia Institute)*

Figure 13: Ridge plot showing loss distribution shape in successive 10-year windows

The reason we're belaboring this technical detail is that it turns out the underlying cause highlights the importance of distinguishing different types of incidents when assessing loss magnitude. Note what happens when we plot separate loss distributions for each of our incident patterns in Figure 14.

We can now see that the shoulder in the overall loss distribution actually results from a proliferation of fairly small losses from accidental disclosure events. Incidents tied to physical threats and insider misuse also contribute to lower losses, but those patterns are considerably less common.

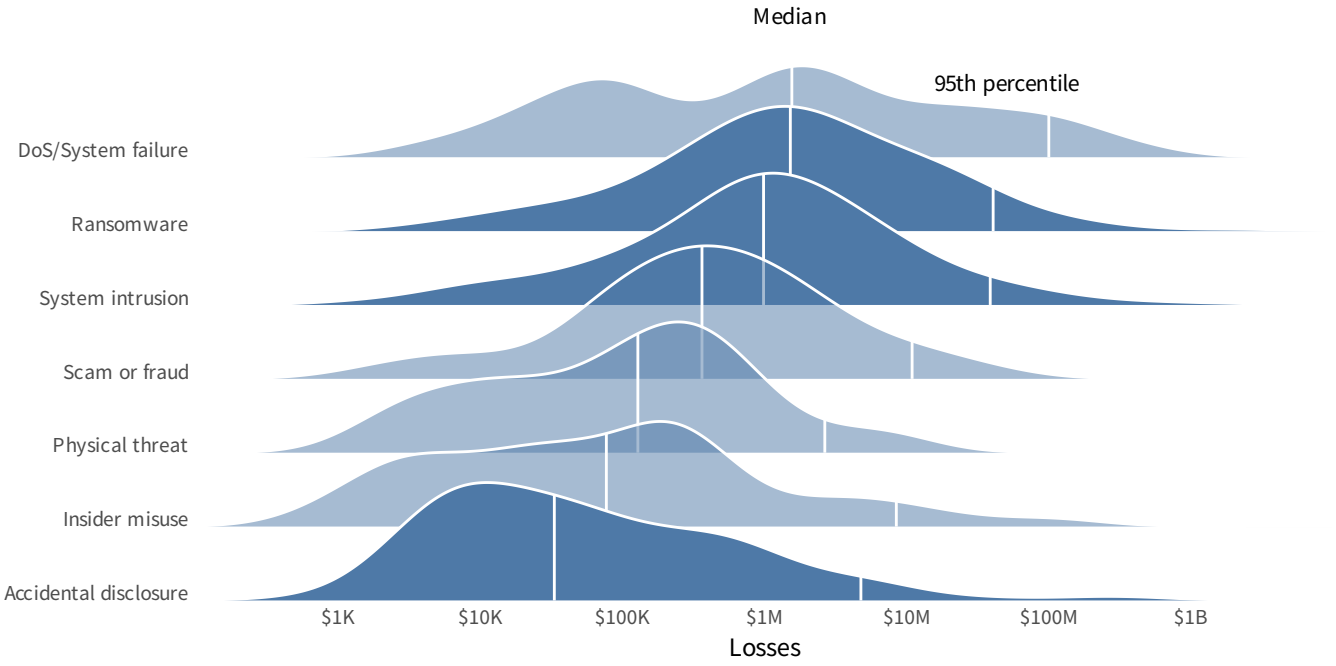


*Source: IRIS 2025 (Cyentia Institute)*

Figure 14: Distribution of reported losses by incident pattern (2008 to 2024)

A possible explanation is that the growing shoulder reflects the rollout of regulations over the years that require public disclosure of even relatively minor data loss events. Thus, comparing loss trends specific to these different incident patterns could provide helpful context. We do just that in Figure 15.

While ransomware losses surge past $27M at the high end, top-tier system intrusions have dropped sharply—down to $7.4M from over $200M.

The trends seen here support our hypothesis. Losses stemming from accidental disclosure and insider misuse have indeed declined over time, particularly for run-of-the-mill events (now just 5% of median cost in 2008).

Conversely, the median loss magnitude for system intrusion and ransomware incidents has risen, with the latter ballooning 20-fold!

The decline in costs associated with system intrusions at the top end of the distribution is both dramatic and curious. The interpretation is clear, though: The biggest intrusions nowadays are significantly less expensive than they used to be.

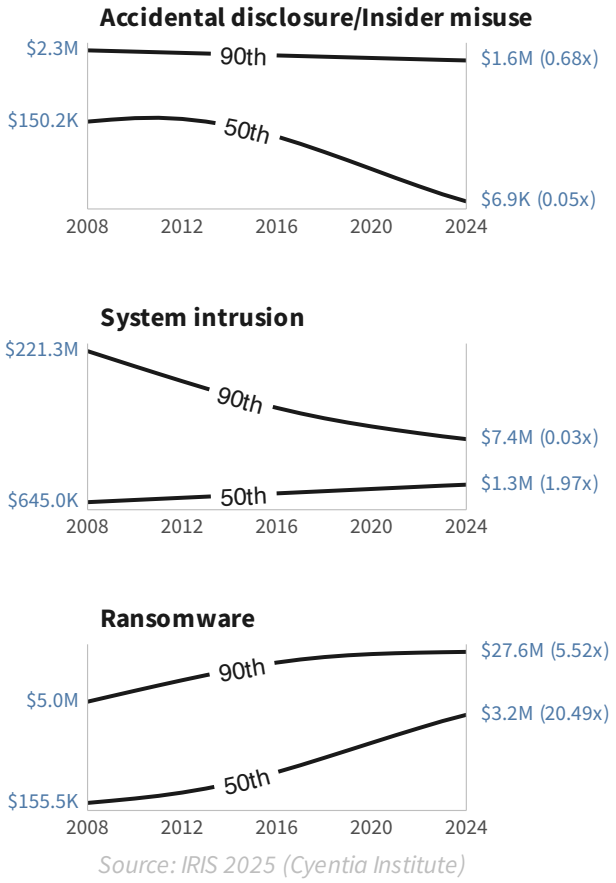


*Source: IRIS 2025 (Cyentia Institute)*

Figure 15: Trend analysis of median and 90th percentile losses by incident pattern

## Key Risk Insight

Both the magnitude and trend of losses from security incidents depend heavily on the type of event. This supports the need for risk scenarios to specify threats

# Q6

# ARE INTRUSION METHODS CHANGING OVER TIME?

Let's say for a moment that the probability and loss estimates for your own organization mirror some of the upswings observed in prior sections. Assuming that trend is climbing above your risk tolerance, you'll want to do something to flatten the curve. But what?

Choosing the best risk treatment strategy has never been easy and likely never will be. But organizations make those decisions even harder when they attempt to jump all the way from high-level assessments down to specific measures to mitigate risk. Discerning whether BlinkyBox1 vs. BestPractice2 vs. the latest [enter acronym] solution is the most effective option strongly depends on the maturity of your security program and the particular threats driving your risk exposure upward.

That's why tracking common adversary tactics, techniques, and procedures (TTPs) behind cyber events can bridge the divide between risk assessments and risk treatment. MITRE ATT&CK offers a knowledge base of TTPs that is convenient for this purpose. Cyentia uses a combination of methods to identify ATT&CK techniques associated with incidents, and we'll use those capabilities here to explore the titular question.

Tracking adversary tactics, techniques and proceduers (TTPs) is the missing link that connects high-level risk assessments to effective action.

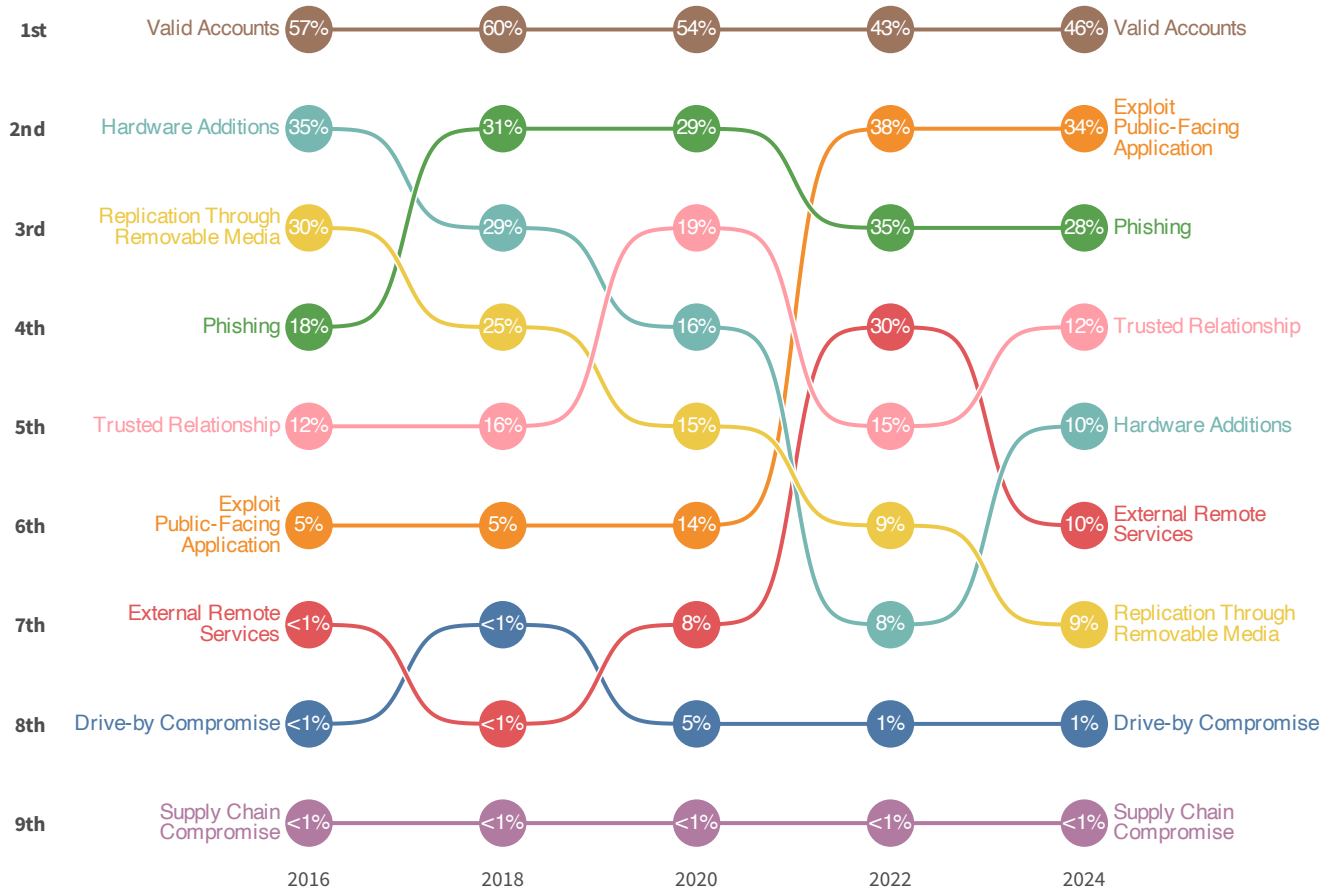


*Source: IRIS 2025 (Cyentia Institute)*

Figure 16: Prevalence of ATT&CK Initial Access techniques observed in incidents over time

Figure 16 presents trends in ATT&CK Initial Access techniques observed over the last nine years, according to the percentage of incidents[19] associated with each. Overall, the shifts seen here largely reflect the never-ending cat-and-mouse game played between attackers and defenders. A few trends are particularly noteworthy.

Cyber threats are ever-changing, which makes it even more remarkable when TTPs don't change all that much. Using Valid Accounts to gain illicit access (e.g., by compromising user credentials) has held the pole position for the entire time period. Phishing—a popular means of obtaining those credentials—has also consistently ranked among the top techniques. To be fair, there are many different schemes by which attackers abuse user accounts as well as many sub-techniques for phishing. But that doesn't change the overall lesson here regarding the longevity of these methods.

The "moving-up-the-charts" award among intrusion methods goes to Exploit Public-Facing Applications (i.e., web application attacks) and External Remote Services (i.e., misconfigured remote access tools). Both techniques have surged from single-digit percentages to heights of 38% and 30%, respectively. This likely reflects the expansion of enterprise attack surfaces over the last decade. These external points of presence are intended to serve customers, third parties, and remote employees, but attackers increasingly take advantage of them as well.

19 Percentages are based on incidents for which at least one ATT&CK technique was identified.

Recent Cyentia Institute studies have found that 99% of Global 2000 companies are connected to vendors that have had recent breaches[20] and that 90% of organizations consider third-party risk management a growing priority.[21] Those concerns appear to have merit, based on the persistence of actors leveraging Trusted Relationships with external service providers to compromise target organizations. Though third-party risk and Supply Chain Compromise events are often referred to interchangeably, MITRE has a more narrow definition for the latter that's relatively rare among publicly known incidents.

Web app exploits and remote access misconfigurations rise from single digits to 38% and 30% of intrusions.

On the topic of evolving threat actor strategies, we noticed an interesting trend when analyzing initial access techniques among organizations of varying sizes. At a high level, one could rightly infer that attackers use similar techniques regardless of the size of the target entity. Valid Accounts is firmly on top for all three size tiers, and the others are similarly clustered below that.

But upon closer inspection, interesting variations become apparent. Abusing Valid Accounts is trending down over the last few years for organizations below $10B in annual revenue, but rising sharply for the largest corporations. Phishing and exploiting applications are most prevalent among incidents affecting smaller firms (<$100M) and progressively less in higher revenue tiers. Attacks targeting Trusted Relationships disproportionately affect larger organizations, which makes sense given their extensive portfolio of third-party vendors.

## Want more ATT&CKification of incidents and losses?

### Here are two Cyentia resources:

**MULTI-SOURCE ANALYSIS OF TOP MITRE ATT&CK TECHNIQUES (WITH TIDAL CYBER)**

**INFORMATION RISK INSIGHTS STUDY RANSOMWARE EDITION (SPONSORED BY CISA)**

**Cyentia is also currently working on an update to the IRIS 20/20 "Xtreme" that analyzed the largest incidents from 2015-2019. The prior report did not incorporate ATT&CK, but the next iteration most definitely will. So, if you want to learn more about TTPs behind the biggest loss events of the last five years, keep an eye on the IRIS site for updates on that study.**

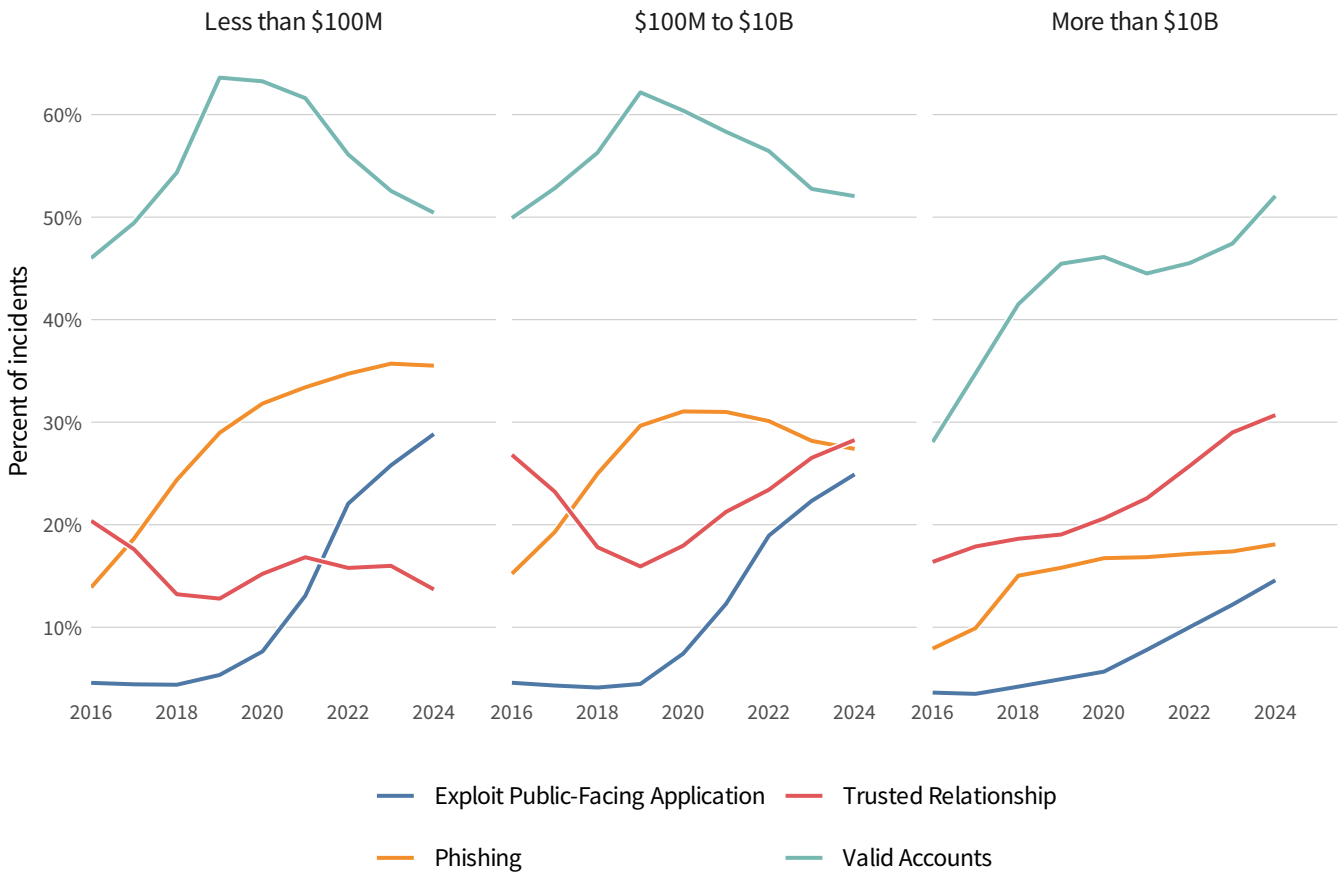


*Source: IRIS 2025 (Cyentia Institute)*

Figure 17: Prevalence of ATT&CK Initial Access techniques observed by revenue tier

These seemingly contradictory takeaways are reasonable based on the historical attack trends specific to each class of organizations. The point is that the adversary TTPs trending for *waves hand* them aren't necessarily the ones shaping your risk posture.

We'll close with a reminder that Initial Access is only one of more than a dozen tactics in MITRE ATT&CK. While trending techniques within each tactic are possible, this is already a long report, and we have one last question we'd like to explore.

20 Global 2000: Industry Titans Battle the Beast of Supply Chain Cyber Risk (with SecurityScorecard).

21 The State of Third-Party Risk Management (with RiskRecon).

# Q7

# WHAT ARE WE MISSING FROM CURRENT EVENTS?

Since the beginning of the IRIS series, we've tried to be open about the shortcomings and potential biases in our dataset of publicly reported incidents. One point often mentioned is the reporting lag that stems from the time it takes for details to make their way into the public record.[22] Another is that some types of cyber events (or attack details) are underrepresented because they don't have immediate visible impacts or don't trigger mandatory disclosure laws. That's why we're especially grateful to Feedly, a real-time threat graph, for allowing us to analyze cyberattacks collected via their intelligence capabilities for this section.

To derive the Feedly data analyzed in this study, we started with events identified by Feedly's Cyber Attacks AI model from 2024. The model was developed by Feedly to discover and research emerging threats, which makes it more of an "ear-to-the-ground" signal of current cyber events than our historical loss database.

We'll state up front that this isn't an attempt to determine who's right and who's wrong. As the lead question implies, we're concerned with what our core incident dataset might be missing from recent media coverage that hasn't yet made (and may never make) it into the official public record.

With that in mind, let's get to it!

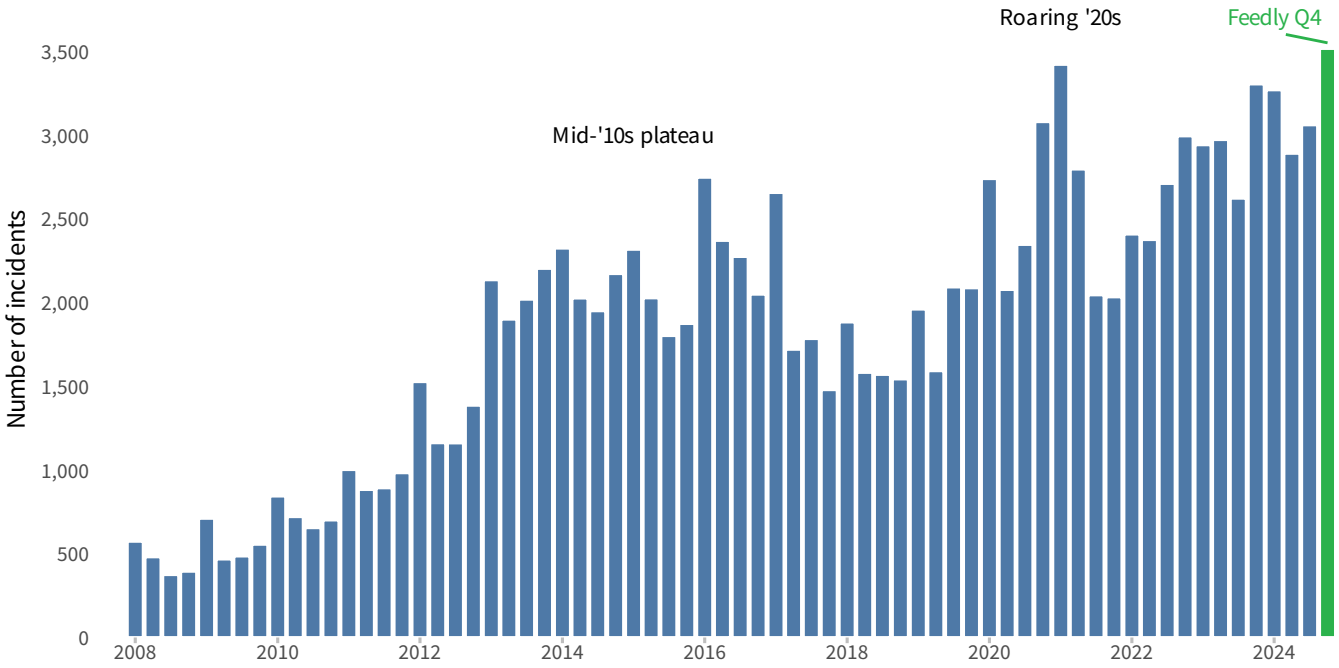


*Source: IRIS 2025 (Cyentia Institute)*

Figure 18: Number of security incidents publicly reported or discovered each quarter. Feedly was used for incident discovery in Q4-2024 to compensate for the lag in the historical dataset.

22 This is why many charts in this report show an apparent downturn in 2024; we'll still be learning of 2024 incidents long after this report is published in 2025.

Remember that apparent Q4-2024 dip in security incidents back in Figure 1? It's not real. Figure 18 plainly shows that the 3,500+ incidents observed by Feedly during that timeframe fill the hole in our base historical dataset. This corroborates our suspicion of a reporting time lag rather than a drop in frequency. It also emphasizes the need for more real-time tracking of incidents if your risk analysis relies on near-term trends.

It's worth noting that the incidents in both datasets show a similar representation of affected industries. Each attributes the highest number of events to Healthcare and Finance. Feedly rounds out the top three with the Public sector, while our data has Professional Services in third place (Public is #4). That's a good indication that the comparisons in this section are based on samples that are reasonably similar in nature.

Speaking of near-term trends, you may have noticed that we rarely include analysis of the latest threat actor campaigns in the IRIS series. That's partially due to our focus on risk management vs. threat intelligence. But it's also because public disclosures and filings that comprise our dataset usually don't delve into attribution. Media outlets, on the other hand, love a good "whodunnit?" story, and Feedly… well… feeds off those stories.

Table 2 lists the top three threat actors behind the most incidents affecting the Finance, Healthcare, and Public sectors, according to Feedly's collections during 2024. Since this isn't a threat intel report, we won't dive into the backstory of these groups—other resources are better suited to that. We include this simply to make the point that adversaries often have unique goals and targets. Keep that in mind if you're looking to incorporate specific threat actors into your risk scenarios and assessments.

| FINANCE | HEALTHCARE | PUBLIC |
|---|---|---|
| Lazarus Group | Alpha Spider | GhostEmperor |
| Shiny Hunters | RansomHub | Volt Typhoon |
| RansomHub | Vanilla Tempest | Flax Typhoon |

Table 2: Top threat actors associated with 2024 security incidents by sector (via Feedly)

It's hard to talk about threat actors without the conversation turning to the TTPs they use. So, let's go there next. Table 3 lists the top five ATT&CK techniques observed by Feedly and two different date ranges for our historical incident dataset. We do that to enable comparisons based on both time period and source.

All sources place Valid Accounts and Phishing in the top three spots, albeit in varied order. This further substantiates these techniques as primary attack vectors for treating risk exposure. There's also relative agreement for several techniques that land in the middle of the pack. Beyond that, there are some notable differences between the two sources.

We chalk the disparity around the ranking of Trusted Relationship primarily up to Cyentia's classification choices. MITRE's definition strictly refers to external third parties for that technique, while we traditionally broaden it to also apply to certain types of insider and contractor misuse. We plan to revisit this in the future in light of MITRE's Insider Threat TTP Knowledge Base project.

Frequency of MITRE ATT&CK Initital Access Techniques

Top ATT&CK initial access techniques identified by Cyentia (2014-2023 & 2024) and Feedly (2024)

| Name | Cyentia (2014-2023) | Cyentia (2024) | Feedly |
|---|---|---|---|
| Phishing (T1566) | 2nd | 3rd | 1st |
| Valid Accounts (T1078) | 1st | 1st | 2nd |
| Exploit Public-Facing Application (T1190) | 6th | 2nd | 6th |
| **Supply Chain Compromise (T1195)** | 9th | 9th | **3rd** |
| **Hardware Additions (T1200)** | **3rd** | **5th** | - |
| Drive-by Compromise (T1189) | 8th | 8th | 4th |
| Replication Through Removable Media (T1091) | 4th | 7th | 7th |
| **Trusted Relationship (T1199)** | **5th** | **4th** | - |
| External Remote Services (T1133) | 7th | 6th | 5th |

*Source: IRIS 2025 (Cyentia Institute)*

Table 3: Comparison of top ATT&CK Initial Access techniques observed in incidents by source

The disparate ranking of Hardware Additions is threefold. First, that technique was much more prevalent among incidents a decade ago, but has been declining ever since (see Figure 16). Second, even in its heyday, this technique was mostly related to skimmers added to payment terminals rather than network taps and other more advanced threat scenarios that ATT&CK seems to have in view. Third, this technique tends to be a bit "in the weeds" for media coverage of events and is described in ways that maket echnique extraction difficult.

The relative prominence of Supply Chain Compromise in Feedly's collections is very interesting to us because very few public incident disclosures list this as an initial access technique (see Figure 16, where it's always in a distant last place).

We suspect one of the key reasons for this disparity is that supply chain security is hot of late, and media outlets are more motivated to dig for such details than companies are to include them in official reports. The aforementioned time lag of incident disclosure could be another factor; perhaps details will soon emerge that retroactively bump this technique higher in recent years. Thus, Feedly's collections may grant a forward-looking, headline-driven view of the most common techniques, while our legacy data offers more of an actuarial perspective. Both views are informative for assessing risk and developing mitigation strategies.

Last but not least, we'll examine a comparison of reported financial losses in Figure 19. The statistics for the 10-year and 2024 views of losses in our dataset are relatively comparable. The loss magnitude of events identified by Feedly in 2024, however, reveals major differences.

Note first that Feedly contains more 2024 events with identified financial losses than we found in our core data source. This further corroborates the utility of monitoring current events to help compensate for the reporting lag in risk data.

Feedly-sourced cyber events show a median loss of $28.5M—30 times higher than the historical median.

The median loss for Feedly-sourced cyber events stands at $28.5M, which is 30x higher than that of our historical dataset ($603K). The 95th percentile for the two sources shows a disparity of almost $750M!

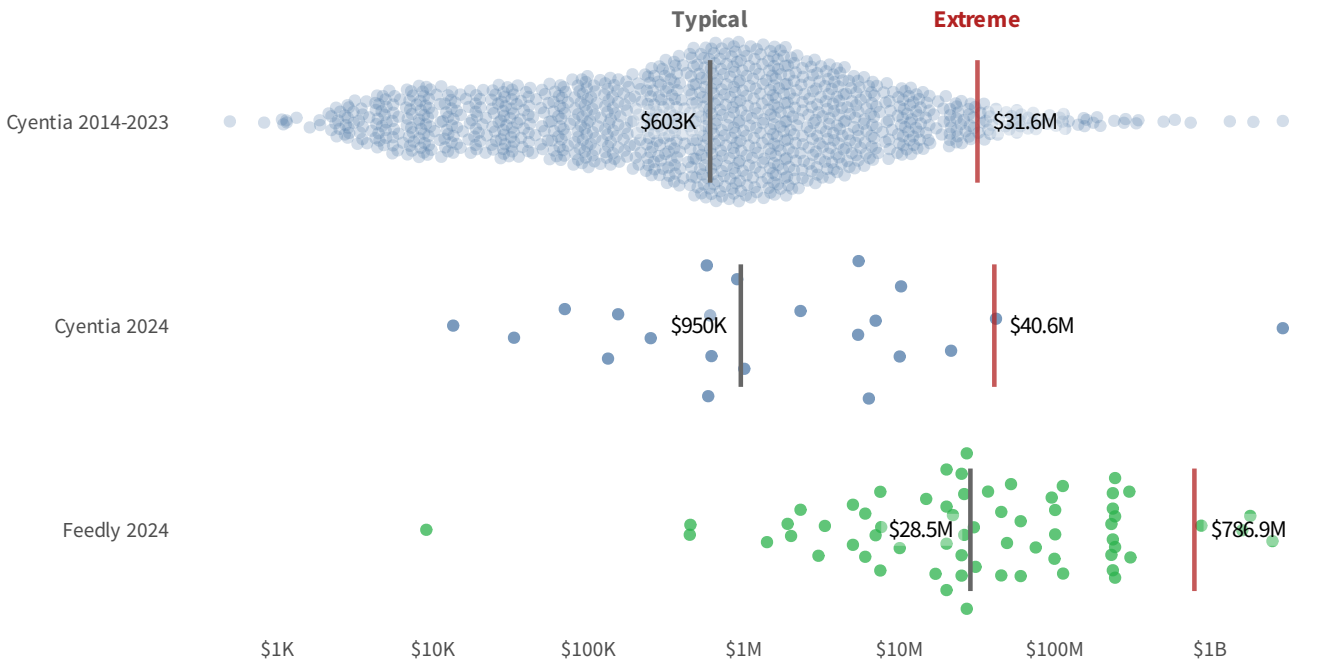


*Source: IRIS 2025 (Cyentia Institute)*

Figure 19: Comparison of reported financial losses from incidents by source

Much of what we see (or don't see) here goes back to sourcing methods. The primary source of our historical loss data, Zywave, has a strong focus on the insurance and reinsurance market. They're diligent in gathering datapoints on all aspects of losses, both large and small. Minor loss events might be interesting to insurers managing risk across a large portfolio of organizations, but media outlets tend to focus on major breaches or disruptions. Since such events rarely go unnoticed, they're ripe for open-source intel collection. If you're focused on tail risk—which is what many execs are most concerned with—closely tracking these mega-loss events is essential.

## Key Risk Insight

If your cyber risk analysis relies on near-term trends, consider incorporating sources that emphasize current events to supplement historical event data.

> *One of my favorite things about my time on Verizon's DBIR team was working with the scores of external organizations that contribute data for analysis in that report. It's something I'm glad we've been able to continue in our research at Cyentia.*
>
> *I'm grateful to all our data partners and sponsors who make it possible for us to do impactful research for the community. Reach out if you have data that would unlock new avenues of cyber risk analysis!*

***~ Wade Baker PhD***
*Co-Founder | The Cyentia Insitute*

# A1 METHODOLOGY

## Data Collection

The IRIS research draws heavily upon Zywave's Cyber Loss Data, which contains over 150,000 security incidents and associated losses spanning decades. The data is compiled from publicly available sources, such as breach disclosures, company filings, litigation details, and Freedom of Information Act requests. It is the most comprehensive source of cybersecurity incidents and losses available.

That said, we're not claiming this dataset 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 are well represented.

Additionally, Cyentia does extensive processing of Zywave's base dataset to extend and enrich it for cyber risk analysis use cases. This is done using a combination of classification models, natural language processing (NLP), taxonomy mapping, malware behavioral analysis, and manual tagging by our analysts.

Incident and loss data collected by Feedly is used for the last section to study trends we might be missing from current events. We started with 2024 events identified by Feedly's Cyber Attacks AI model. We further refined this by focusing on "memes," which is Feedly's method of clustering articles and information on a trending topic. The point is that we're not simply counting articles in this analysis. Finally, we reviewed the identified events to train a classification model to distinguish successful incidents affecting organizations from other threats and trends that aren't comparable to our core dataset.

### GOT DATA FOR THE NEXT IRIS?

If you have information on security incidents and losses and might be willing to contribute anonymized data for analysis in a future IRIS, please reach out! We're especially keen to incorporate insights from cyber insurance claims and incident response investigations.

## Incident Likelihood

Though we don't delve into it in this edition, readers of prior IRIS may recall that we have modeled the frequency of security incidents over a fixed 10 year time period, allowing for the fact that organizations can have more than one in a single year. We took the same approach in this edition, expanding the model to include a time component. Ultimately, we present estimates as the annual probability of an organization experiencing at least one event.

To do this, we divide our historical dataset into 12-month rolling windows and count the number of incidents for each organization. This gives us a large number of observations that allow us to more confidently model the annualized loss event frequency.

We then treat these observations as samples from an underlying probability distribution and use random effects models to estimate the parameters both overall and within specific slices like industry and revenue bands over time. The result is a closed-form representation of the probability that an organization will experience a certain number of incidents in a given year that can change over time.

Additionally, in prior IRIS reports we reported both upper and lower bound estimates for incident likelihood. We've dropped that distinction in this edition in favor of exclusively using the more risk-averse upper bound estimate.

In a nutshell, the difference between these approaches stems from the count of organizations used as the denominator for the calculation. We don't know how many exist throughout the world, so the upper bound uses the total number of organizations that exist in our historical incident database. While it's true that this approach excludes some extremely secure or lucky firms, the fact is that those prone to incidents in the future have probably had one at some point in the past. The result is a more conservative estimate that we believe is more suitable for risk management.

## Financial Losses

Financial losses tend to be less reported than other data points for cyber events. There are many reasons for this, but the result is that the majority of incidents in our dataset do not include anything about losses. Those that do tend to reflect direct losses that are easier to quantify (e.g., response costs or lost revenue) and/or identify from public records (e.g., class action suits or SEC filings). Indirect and intangible impacts usually aren't captured.

The good news, from a data standpoint, is that the record of losses from major security incidents—like those we analyze in this study—is more complete than for minor events due to increased visibility and reporting. Thus, we hold that our loss dataset is sufficient to form a well-supported model of cyber events over the last 15 years.

Note that all financial loss values presented in this report have been adjusted for inflation.

# A2

# INCIDENT PATTERN DEFINITIONS

All security incidents in our historical dataset are assigned one of these mutually exclusive[23] patterns using a combination of natural language processing techniques and human expert assessment.

**ACCIDENTAL DISCLOSURE:** Data stores that are inadvertently left accessible to unauthorized parties, typically through misconfigurations on the part of the data custodian.

**DOS ATTACK:** Any attack intended to render online systems, applications, or networks unavailable, typically by consuming processing or bandwidth resources.

**DEFACEMENT:** Any unauthorized content modification to an organization's website or other online assets.

**FRAUD OR SCAM:** Any incident that primarily employs various forms of deception to defraud the victim of money, property, identity, information, and so on.

**INSIDER MISUSE:** Inappropriate use of privileged access, either by an organization's own employees and contractors or a trusted third party.

**PHYSICAL THREATS:** Threats that occur via a physical vector, such as device tampering, snooping, theft, loss, sabotage, and assault.

**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.

**SYSTEM FAILURE:** All unintentional service disruptions resulting from system, application, or network malfunctions or environmental hazards.

**SYSTEM INTRUSION:** All attempts to compromise systems, applications, or networks by subverting logical access controls, elevating privileges, deploying malware, and so on.

23 Yes, it's true that an incident could involve more than one of these (e.g., system intrusion and ransomware). However, the purpose of these patterns is to represent the primary nature of the event.

# A3

# CHARTS & TABLES FROM PRIOR IRIS STUDIES

This appendix contains up-to-date versions of selected figures from IRIS 2022 that provide probability and loss comparisons among sectors and revenue bands. If there are other figures you'd like to see from an IRIS of yesteryear, let us know!

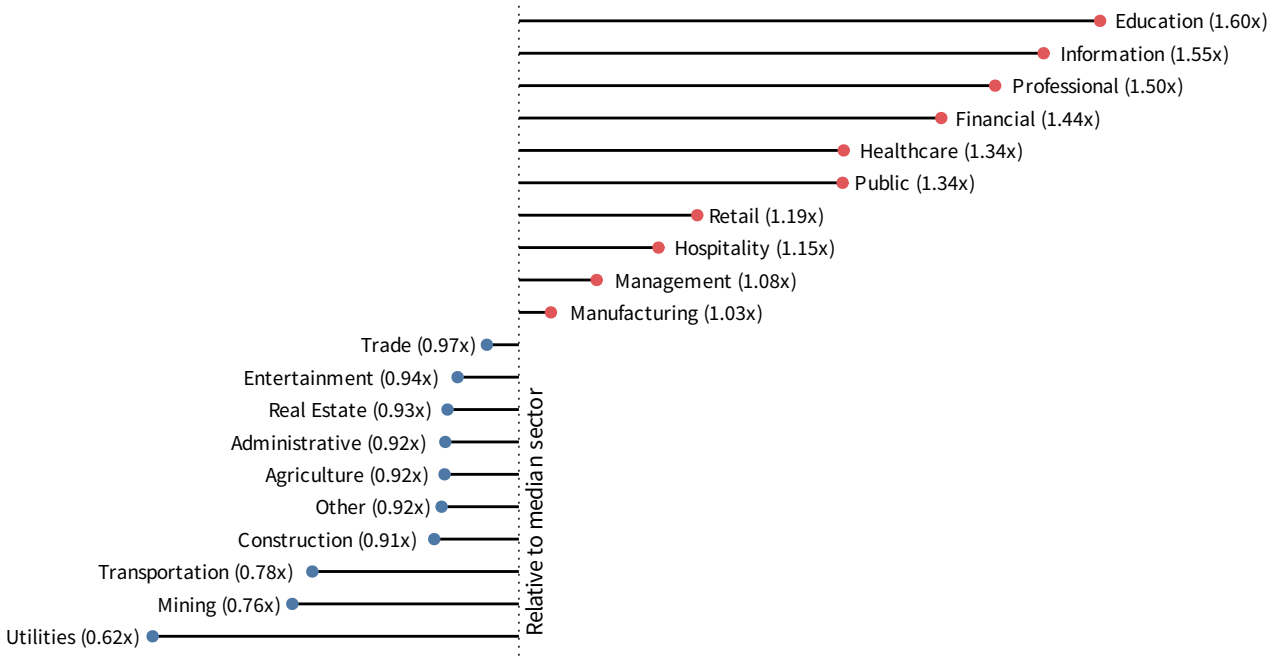


*Source: IRIS 2025 (Cyentia Institute)*

Figure A1: Relative probability of one or more loss events among sectors (Figure 5 in IRIS 2022). The point of comparison is the overall median across all organizations.

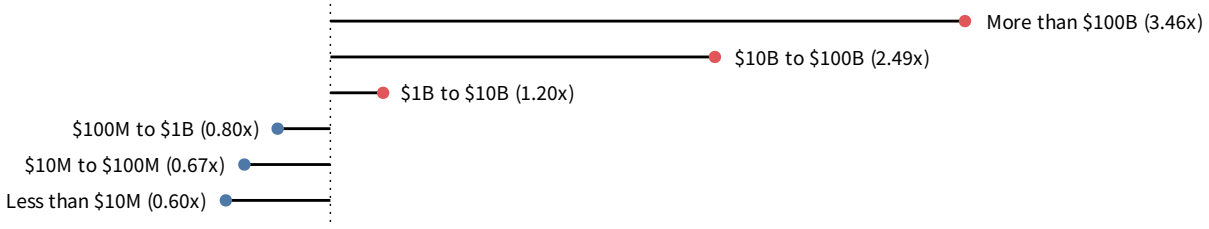


*Source: IRIS 2025 (Cyentia Institute)*

Figure A2: Relative probability of one or more loss events among annual revenue tiers. The point of comparison is the overall median across all organizations.

| Losses observed per sector | | | |
|---|---|---|---|
| **Sector** | **Geometric mean** | **Median** | **95th percentile** |
| Administrative | $318K | $529K | $31M |
| Agriculture | $1M | $2M | $3M |
| Construction | $164K | $189K | $5M |
| Education | $226K | $249K | $6M |
| Entertainment | $147K | $282K | $12M |
| Financial | $951K | $1M | $194M |
| Healthcare | $524K | $557K | $14M |
| Hospitality | $687K | $600K | $62M |
| Information | $783K | $718K | $217M |
| Management | $343K | $332K | $140M |

| Losses observed per sector | | | |
|---|---|---|---|
| **Sector** | **Geometric mean** | **Median** | **95th percentile** |
| Manufacturing | $1M | $1M | $42M |
| Mining | $1M | $1M | $2M |
| Other services | $262K | $348K | $41M |
| Professional | $400K | $736K | $17M |
| Public | $234K | $214K | $18M |
| Real Estate | $244K | $236K | $2M |
| Retail | $872K | $746K | $45M |
| Trade | $902K | $1M | $23M |
| Transportation | $286K | $490K | $23M |
| Utilities | $113K | $146K | $3M |

*Source: IRIS 2025 (Cyentia Institute)*

Figure A3: Loss magnitude summary statistics by sector (Table 4 in IRIS 2022)

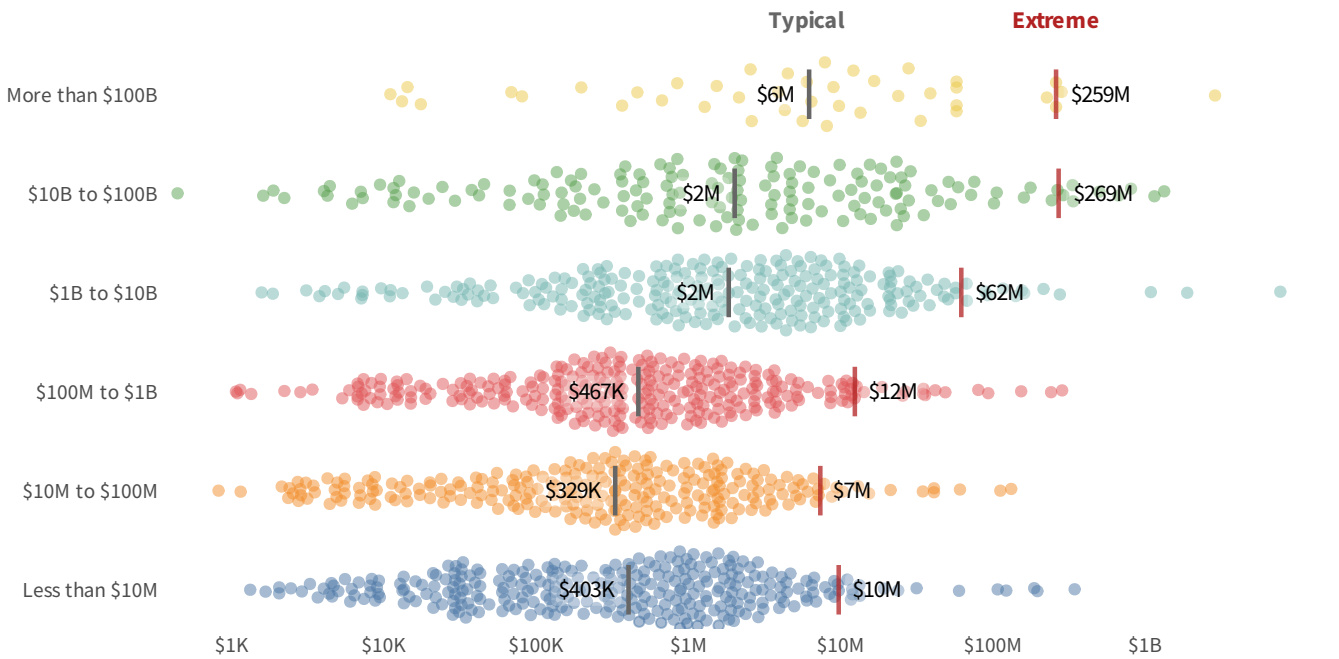


*Source: IRIS 2025 (Cyentia Institute)*

Figure A4: Distribution of reported cyber event losses by annual revenue of affected firms (Figure 7 in IRIS 2022)


**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.