Information Risk Insights Study

RANSOMWARE

A Detailed Analysis of the Frequency and Impact of Ransomware Events

Introduction

> “Is it not a strange fate that we should suffer so much fear and doubt for so small a thing?”
>
> ~Boromir, *Lord of the Rings*

Few cyber threats have inspired more fear, uncertainty, and doubt (FUD) in recent years than ransomware. Organizations fear they’ll be the subject of headlines detailing the latest crippling wave of attacks. They’re uncertain about the likelihood and impact of such an event and doubt whether current defenses adequately mitigate the risk. While ransomware FUD-mongering abounds, the risk is real, and concerns justified. It is no “small thing.”

Like its predecessors, this latest edition of the Information Risk Insights Study (IRIS) is about FUD-managing rather than FUD-mongering. Our goal is to shine the clarifying light of rigorous analysis to dispel the fog of FUD enshrouding ransomware that inhibits organizations’ ability to make informed, rational decisions to manage risk.

To enable that, we leverage a massive dataset containing over 14,000 ransomware events that compromised over a billion data records and led to projected financial losses topping $270 billion over the last five years. We highlight key findings from our analysis of that data on the next page and dive right into the detailed insights straightaway after that.

THANKS FOR READING!

Acknowledgements

The Cyentia Institute wishes to thank and acknowledge 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 lead to better cyber risk assessments and cyber risk reduction decisions for years to come.

The Cyentia Institute is a research and data science firm working to advance cybersecurity knowledge and practice. We pursue this goal through our data-driven products and joint research publications like this study.

Table of Contents

INFORMATION RISK INSIGHTS STUDY

RANSOMWARE Key Findings

Ransomware was behind 32% of all security incidents and 38% of financial losses from cyber events reported over the last five years.

In 2015, ransomware accounted for <1% of all incidents. 2023 culminated with ransomware averaging over half (52%) of monthly reported cyber events.

There's about a 10% upper bound chance that any given organization will experience at least one ransomware incident in the next 12 months.

51%

Ransomware claims 51% of incidents in the Manufacturing sector. Compare that with 15% for Financial Services.

Less than 10% of all incidents affecting $100B+ enterprises is attributed to ransomware. But among <$100M companies, that ratio jumps into the 30% to 40% range!

$50 MILLION

The typical ransomware event costs $1.4M—almost twelve times that of other types of incidents! The 95th percentile loss balloons to about $50M.

$3.7 MILLION

In just the last five years, the typical financial loss from ransomware incidents has grown from $686K to $3.7M.

Ransomware accounts for about 80% of all reported cyber losses in the Transportation, Education, and Manufacturing sectors.

$276 BILLION LOST

We estimate the total financial losses of ransomware incidents to be about $276B over the last five years.

140x

Total estimated losses from ransomware have grown 140X over the last 10 years!

Like what you see? *Join the vision!*

We intend to continue the IRIS in the future to discover even more insights for managing cyber risk. If you'd like to join in that effort by contributing relevant data or sponsoring, please reach out to us at research@cyentia.com.

119 Cy cyentia

DATA USED IN THIS STUDY

The bulk of **INCIDENTS** and **LOSSES** analyzed in this study originate from Advisen's Cyber Loss Data, a repository containing ~150,000 security incidents spanning decades. They compile this data through publicly available sources, such as breach disclosures, company filings, litigation details, and Freedom of Information Act requests.

This dataset is widely used, with three features that make it ideal for this research:

1 It is the most comprehensive list of historical cyber incidents we've found.

2 It tracks losses publicly disclosed in the wake of those incidents.

3 It includes supplemental firmographic information on organizations affected by cyber events and the broader economy.

Our analysis focuses primarily on ransomware incidents that occurred from 2019 through 2023, unless otherwise indicated. This five-year timeframe encompasses over 14,000 ransomware events that compromised more than 1.1B data records and led to projected financial losses topping $270B.

It's important to note that we're not claiming the data used in this report reflects all ransomware incidents that occurred during this timeframe. We can only analyze those that make their way into the public record, through outward signs or impacts, mandatory reporting, voluntary disclosure, etc. Advisen and Cyentia closely monitor such events and have high confidence that this dataset is representative of significant ransomware events.

In addition to Advisen's standard fields, we further enriched the dataset through a combination of natural language processing, classification models, and manual analysis. We also incorporate datasets from Fortinet, Ransomwatch, and Tidal Cyber into our analysis where appropriate.

Advisen
A ZYWAVE COMPANY

Find out more at Advisen.

TERMS USED IN THIS STUDY

INCIDENT: We use the terms incident, cyber event, and loss event interchangeably in this study to generally refer to adverse events that impact the confidentiality, integrity, or availability of a firm's information assets. Specifically for this report, these terms refer to instances in which organizations were actually compromises by ransomware.

LOSSES: We use losses to refer to the reported financial consequences of incidents. This can include costs associated with disrupted services, incident response, ransom payments, fines and judgements, etc.

Ransomware Ranks Somewhere...But *Where?*

LET'S START WITH THE BIG PICTURE RIGHT UP FRONT: WHERE DOES RANSOMWARE RANK IN RELATION TO OTHER TYPES OF CYBER EVENTS?

Roughly 30% of all publicly reported security incidents from 2019 through 2023 involved ransomware, earning it a #2 ranking for frequency. Ransomware also ranks #1 for financial impact, accounting for over one-third of all reported losses. Finally, this threat is no slouch when it comes to data exposure, ranking 3rd for the total number of records affected.

If we rewind the clock to the prior five-year period from 2014–2018, ransomware had barely begun its climb up the rankings (<2% of events; <4% of losses). The next five years would see a 17X surge in prevalence and ransomware's share of losses balloon sevenfold!

This ranking, of course, begs the question of what other threats we're using for comparison. The incident patterns we traditionally use in the IRIS series are comparable in scope/nature to ransomware and suit this purpose well. All incidents[1] in our dataset are categorized according to common patterns of threat actors, techniques, vectors, and technical impacts. These incident patterns are listed in Table 1 for context and described more fully in Appendix A.

It should be noted that there is some degree of overlap among these patterns. For example, many ransomware events involve attacks intruding into the target organization's network and systems. In such cases, we classified those events as ransomware rather than system intrusion. When events exhibit characteristics that span multiple patterns, we go with the one that fits best.

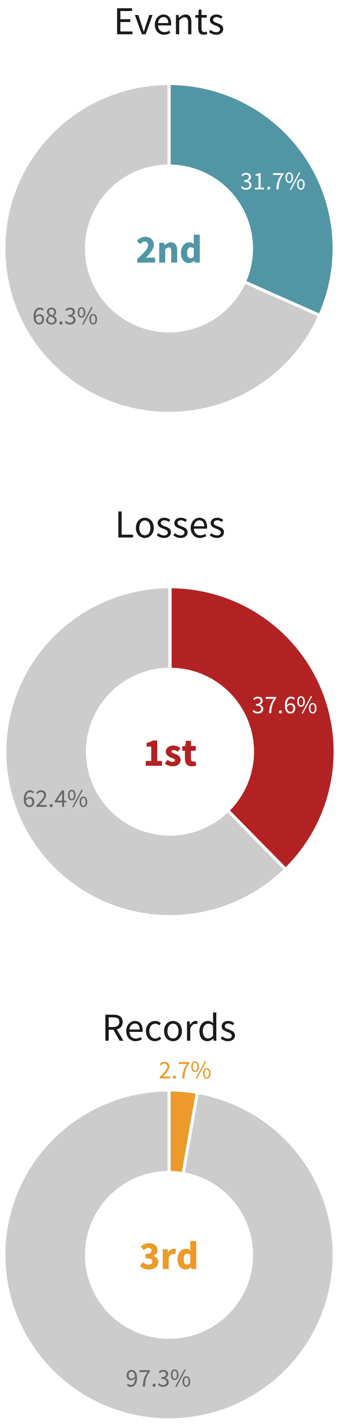

Figure 1 (Right): Ransomware's share of cyber events, losses, and data records

[1] Refer to the Terms Used in this Study for more information on how we define incidents and losses throughout this report.

	Event frequency		Financial impact		Records affected	
	Percentage	Rank	Percentage	Rank	Percentage	Rank
Accidental disclosure	15.9%	3rd	36.1%	2nd	48.67%	1st
Defacement	0.3%	8th	0.0%	9th	0.00%	9th
DoS attack	1.8%	6th	0.0%	8th	0.12%	5th
Insider misuse	2.5%	5th	4.4%	4th	1.08%	4th
Physical threat	3.3%	4th	0.1%	6th	0.01%	8th
Ransomware	**31.7%**	**2nd**	**37.6%**	**1st**	**2.75%**	**3rd**
Scam or fraud	1.1%	7th	0.2%	5th	0.02%	6th
System failure	0.1%	9th	0.0%	7th	0.01%	7th
System intrusion	43.3%	1st	21.4%	3rd	47.35%	2nd

Table 1: Summary of key risk statistics for common incident patterns

The message from the data is clear: ransomware is a top cyber risk for organizations today. The rest of this report is dedicated to digging into the details surrounding key ransomware frequency and impact trends so that organizations are better prepared to manage this risk.

This report is the first in the IRIS series to focus on a single incident pattern. If your agency or organization would like to see a similar analysis applied to other incident patterns or another aspect of cyber risk, we'd be glad to explore sponsorship opportunities.

Visit www.cyentia.com/sponsor to start that discussion.

Ransomware Event Frequency

"May the odds be ever in your favor!"

Effie Trinket, Hunger Games

In our journey to better assess the risk posed by ransomware events, we first explore how often they occur. Our initial step is to examine high-level trends, and then we'll establish an annualized probability of a given organization experiencing a ransomware event. Our ultimate goal is to develop an event frequency model along with the associated parameters to support risk analysis focused on ransomware.

Historical Ransomware Events

History doesn't always repeat itself, but studying past events is usually a better predictor of future trends than blind predictions. With that in mind, Figure 2 tallies the monthly count of all public security incidents (gray dashed line) and splits that into ransomware (light blue line) and non-ransomware (dark blue line) events. We've opted for a 10-year window here to grant a wider view of historical trends.

> While the landscape of security incidents fluctuates, ransomware emerges as the primary driver behind a recent uptick in frequency, overshadowing the slight decline in non-ransomware events.

Keep in mind that public incident reporting often lags months (even years) behind as events progress from discovery to disclosure, which explains the apparent falloff of the overall and non-ransomware trendlines toward the end of the period.

The overall frequency of security incidents fluctuates with a slight rising trend over the last few years. Breaking that out between ransomware and non-ransomware events reveals ransomware to be the primary cause of that rise. Non-ransomware events actually show a slightly downward trajectory.

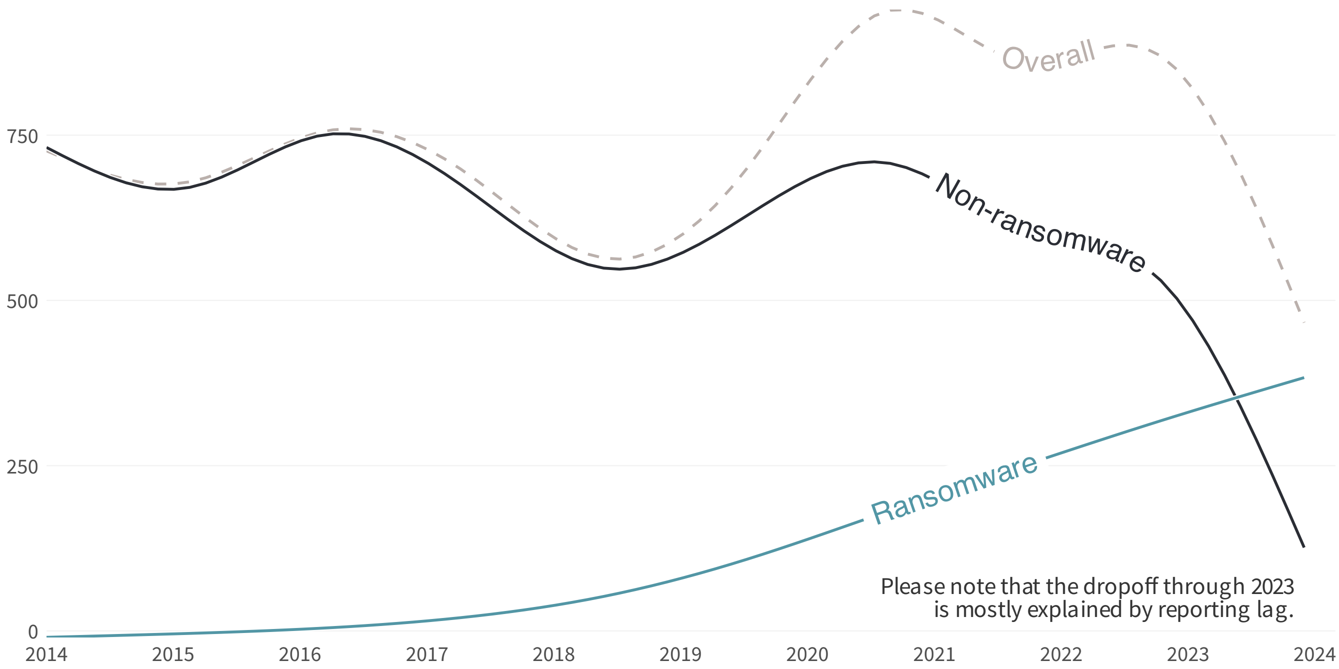

Figure 2: Monthly count of all cyber loss events (top) and ransomware events (bottom)

So, ransomware's clearly running up. But the event counts depicted in Figure 2 actually downplay the rise of ransomware relative to other types of security incidents. Figure 3 makes that more apparent by showing ransomware as a proportion of all recorded events.

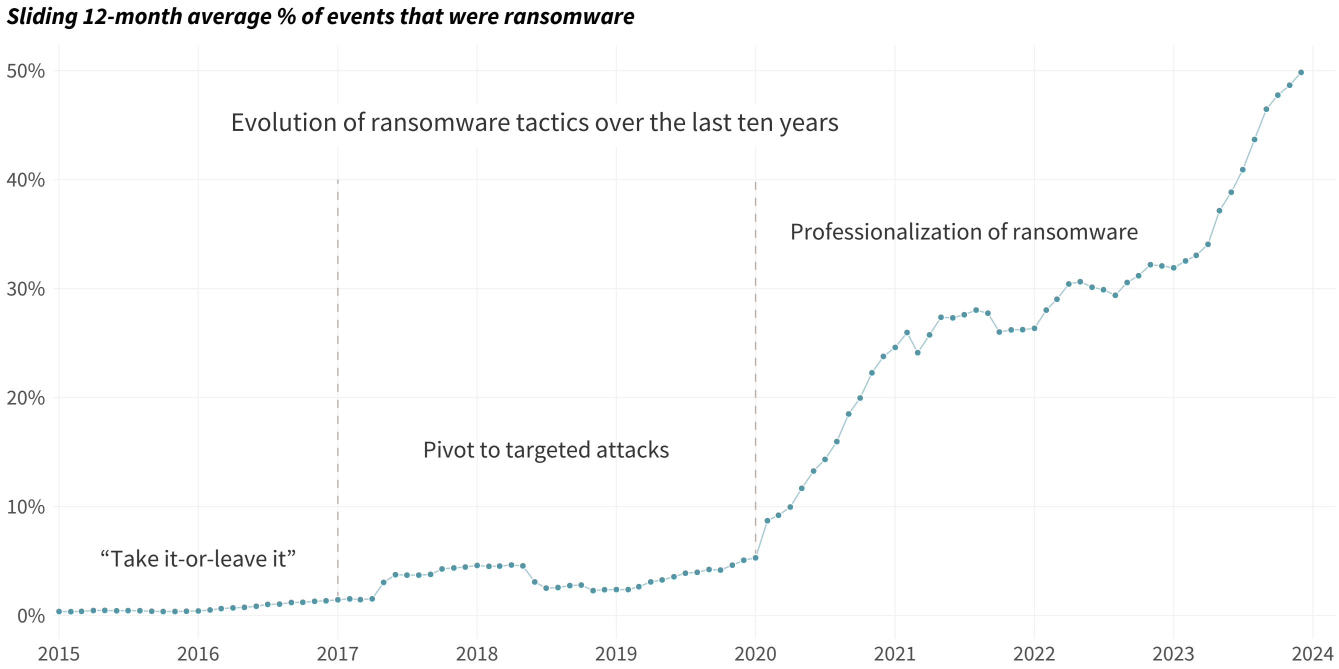

Figure 3: Monthly percentage of all security incidents categorized as ransomware

Ransomware was merely an "also ran" among contenders for the top cyber event category back in 2015, accounting for ~1% of all incidents. As described in this Atlantic Council brief[2], ransomware gangs during this time primarily engaged in "take-it-or-leave it" extortion schemes that involved low-end ransoms.

That same Atlantic Council brief identifies a new phase of ransomware evolution starting around 2016, when it was increasingly used in targeted and destructive attacks. Prominent ransomware strains during this period include Ryuk and REvil. The one-two punch of WannaCry and NotPetya captured global attention and signaled more disruptive use cases.

While everyone else hunkered down during the COVID-19 pandemic, ransomware gangs seized the opportunity to ramp up their operations. They enjoyed higher leverage over victims, more sophisticated capabilities, and huge profits. This era of professionalization fueled the surge seen in Figure 3, culminating in ransomware averaging 52% of monthly reported cyber events throughout 2023!

We suspect that ransomware's dominance among recent cyber events could be a tad inflated. That's because ransomware events tend to enter the dataset faster than other types of incidents because their effects are often immediately noticeable through outages or extortion. By comparison, more surreptitious threats take months or years to detect. So as reporting of other events catches up (see apparent drop-off in Figure 2), ransomware's share in Figures 2 and 3 may wane if we updated it a year from now.

SECTOR AND SIZE TRENDS

Thus far, we've looked at the historical frequency of ransomware events without distinction for the entities experiencing them. But intuition—and perhaps experience—suggests that ransomware plagues certain types of organizations more than others. Let's investigate that now.

Per Figure 4, more ransomware incidents hit the Education sector than any other. The many cybersecurity challenges faced by educational institutions are well documented, and it appears ransomware has amplified those challenges. An environment filled with a diverse array of devices that often aren't centrally controlled creates a large attack surface vulnerable to ransomware infections.

The Professional sector brings up a close second. Most entities in that industry exist to provide services to others, hinting at the disruptive ripple effects of ransomware that spread beyond organizational and sector boundaries.

[2] Behind the Rise of Ransomware, *Atlantic Council*, 2022. https://www.atlanticcouncil.org/in-depth-research-reports/issue-brief/behind-the-rise-of-ransomware

Headlines routinely feature the ravages of ransomware on Manufacturing production lines and Healthcare providers, so it's not surprising to see those industries high on the list in Figure 4. Retail/Trade (a combo of Retail and Wholesale Trade sectors in NAICS) rounds out the top five sectors most frequently affected by ransomware incidents.

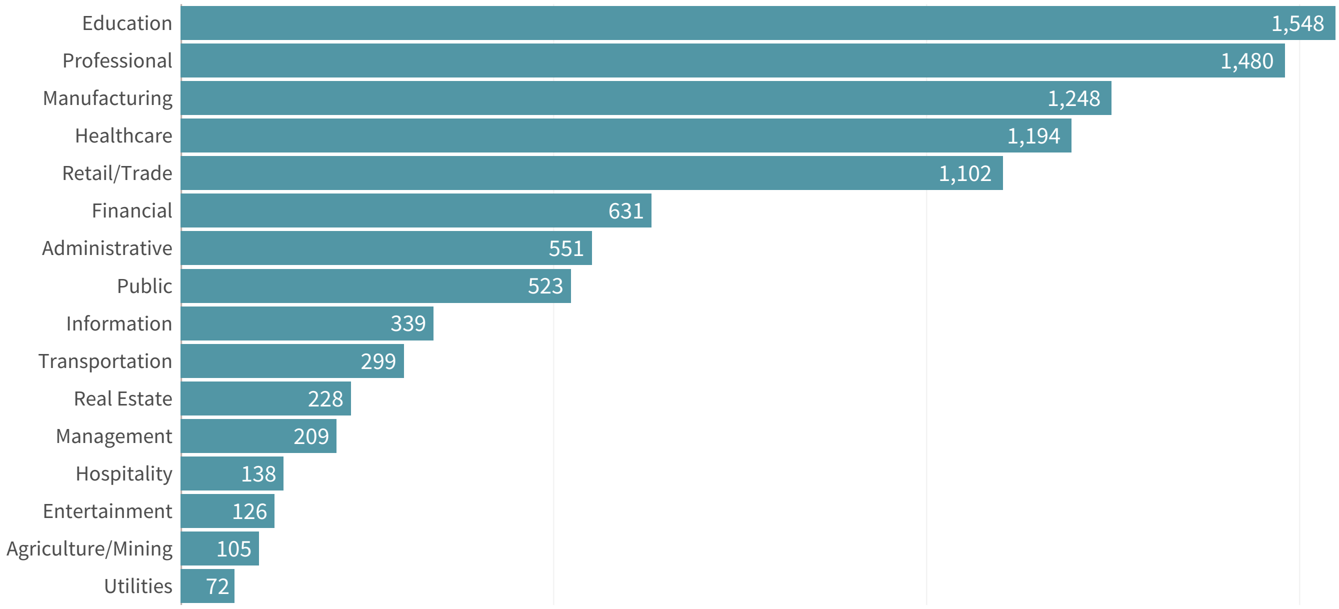

Figure 4: Total number of ransomware events per sector

If we examine the share of all cyber events classified as ransomware in each sector, a different picture emerges from Figure 5. The previously second-ranked Professional sector drops to the middle of the pack, Utilities jumps from last place into the top six, and healthcare plummets to the bottom three. Ransomware claims the highest proportion of security incidents for the Manufacturing, Agriculture/Mining, and Management sectors, echoing the supply chain theme that we're well familiar with from ransomware headlines.

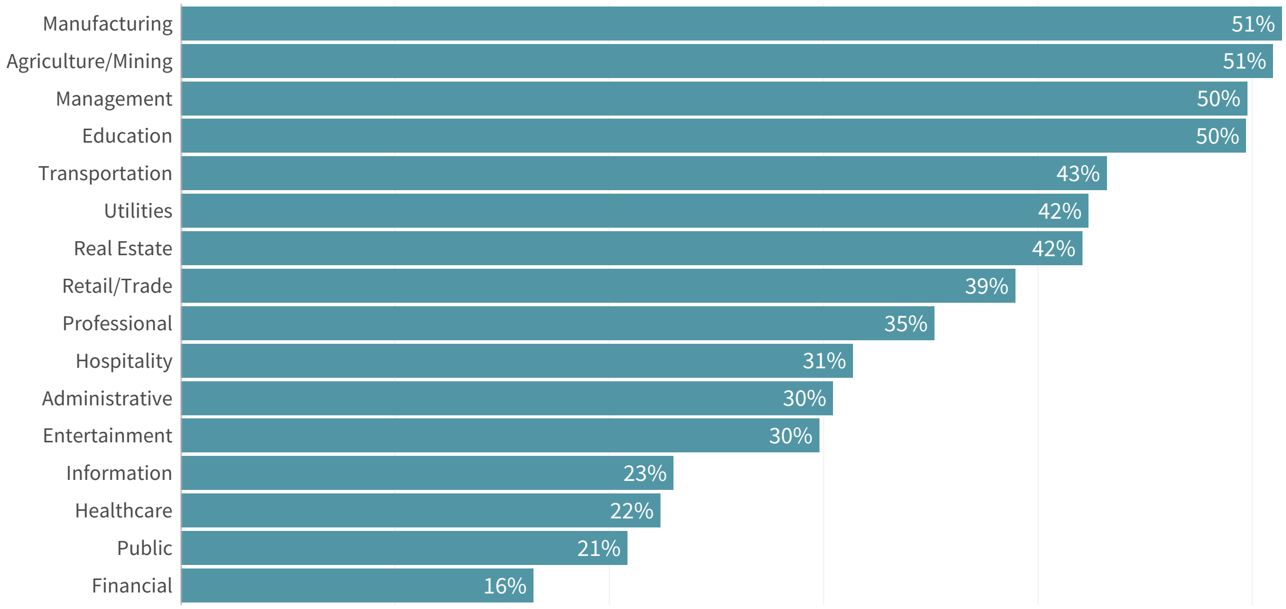

Figure 5: Ransomware percentage of all security incidents, by sector

We can also apply this same analysis to organizations of different sizes based on annual revenue. Doing so reveals that ransomware accounts for less than 10% of all incidents experienced by the largest $100B+ enterprises. But among <$100M companies, that ratio jumps into the 30% to 40% range! In case you didn't know it already, ransomware presents big problems for small businesses.

> Ransomware disproportionately affects small businesses, with incidents comprising 30% to 40% of all security breaches in companies earning less than $100M annually.

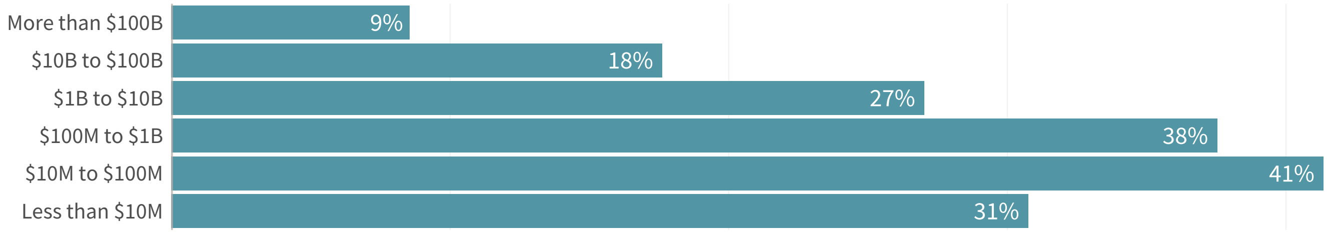

Figure 6: Ransomware percentage of all security incidents, by firm revenue

From this, it's obvious that ransomware gangs have no interest in playing fair. To help level the playing field, the Institute for Security + Technology created a Blueprint[3] for building ransomware defenses in small and medium-sized businesses. CISA also has a hub of resources focused specifically on helping small businesses meet cybersecurity challenges.

Modeling Ransomware Event Frequency

It's enlightening to examine ransomware frequency trends overall or within groups of similar organizations, but the most pressing question for cyber risk management teams is "What's the likelihood of us getting hit?" Answering that requires a model—and you know we love building those!

> Over the last 5 years 93% of affected firms experienced a single ransomware event.

Our dataset contains the names of victim organizations, enabling us to determine how many ransomware incidents they have on the public record. We give that breakdown in Figure 7. Among firms that had at least one ransomware event in the last five years[4], 93% had only the one. A much smaller 6% of organizations experienced two, and thankfully, very few suffered more than that. Good to know, but it's still a ways away from the probability-based questions we ultimately want to answer.

[3] Blueprint for Ransomware Defense, *Institute for Security + Technology*, 2022. https://securityandtechnology.org/wp-content/uploads/2022/08/IST-Blueprint-for-Ransomware-Defense.pdf

[4] Only firms with at least one publicly known ransomware event are in our dataset. There is no explicit, confirmed record of firms with zero events.

Figure 7 (Right): Number of ransomware incidents per organization (that had at least one)

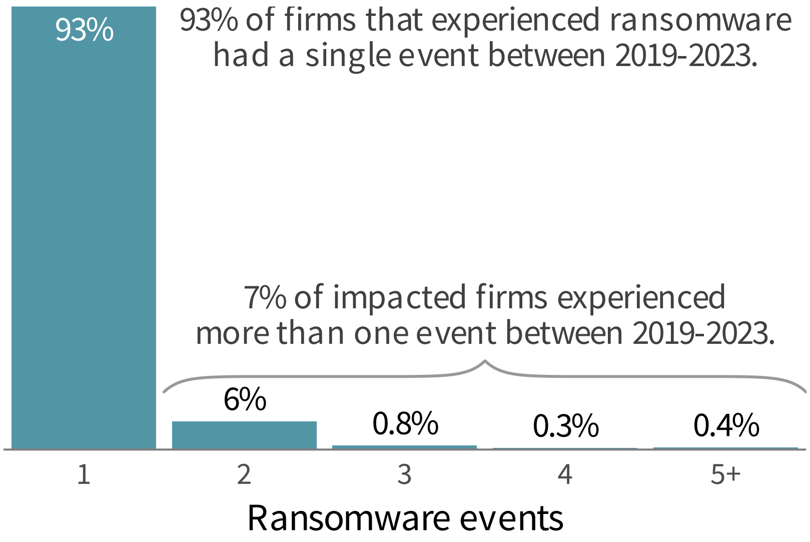

93% of firms that experienced ransomware had a single event between 2019-2023.

7% of impacted firms experienced more than one event between 2019-2023.

To model yearly loss event frequency on a per-firm basis, we could tally the number of ransomware events each year for each organization in the data. But the small number of samples would result in very erratic measurements. Instead, we divided our dataset into 12-month rolling windows and counted the events for each organization. This gave us a larger number of observations that we could employ more confidently to model the annualized loss event frequency.

We then treated these observations as samples from an underlying probability distribution and used maximum likelihood estimation to find the distributions[5] and parameters that best fit the data. We then ran some tests[6] to validate we weren't doing anything that would make our statistics forefathers turn over in their graves.

The benefit of doing all this is that we generate a closed-form representation of the probability that an organization will experience a certain number of ransomware events in a one-year time span. For the DIY risk modelers out there, we include the necessary parameters for ransomware frequency in Table 2. Have fun!

What does the output of this model look like and how does it fit the observed data? Figure 8 presents observed values (gray) and modeled estimates (blue) for annualized loss event frequency using our upper and lower bound approaches. For the most part, the observed and modeled values align, which points to a good model that fits the data.

Frequency parameters: Poisson log-normal

Upper bound

Revenue category	Mean (μ)	Standard deviation (σ)
More than $100B	-2.568299	1.2583515
$10B to $100B	-2.501164	0.7328914
$1B to $10B	-2.431373	0.5410008
$100M to $1B	-2.308153	0.3026587
$10M to $100M	-2.154609	0.2144022

Frequency parameters: Negative binomial

Lower bound

Revenue category	Size (s)	Probability (p)
More than $100B	0.16517935	0.5469493
$10B to $100B	0.31581330	0.8612243
$1B to $10B	0.28175948	0.9085491
$100M to $1B	0.16653069	0.9294919
$10M to $100M	0.04606583	0.9307407

Table 2: Annualized ransomware event frequency model parameters

[5] We tested different distributions for each revenue group and selected the best fit. For most, the Poisson log-normal or negative binomial provided the best results.

[6] Specifically, Kolmogorov-Smirnov and Cramér-von Mises tests.

Lower and Upper Bound Estimates

The reference to upper and lower bound in Table 2 demands an explanation. In a nutshell, the difference between these stems from the count of "all organizations" used as the denominator for the calculation. The dataset records how many firms had a ransomware event, but we don't know how many didn't have one. So we developed two ways of approximating this, one that yields high estimates (upper bound) and one that gives a lower bound. These are described in more detail in Appendix B. For now, just know that the upper bound gives a more risk-averse view that we believe is generally better suited to managing risk.

According to the upper bound estimate in Figure 8, there's about a 10% chance that any given organization will experience at least one ransomware incident in the next 12 months. That works out to getting hit once every 10 years. Now we're much closer to being able to answer that "how likely" question.

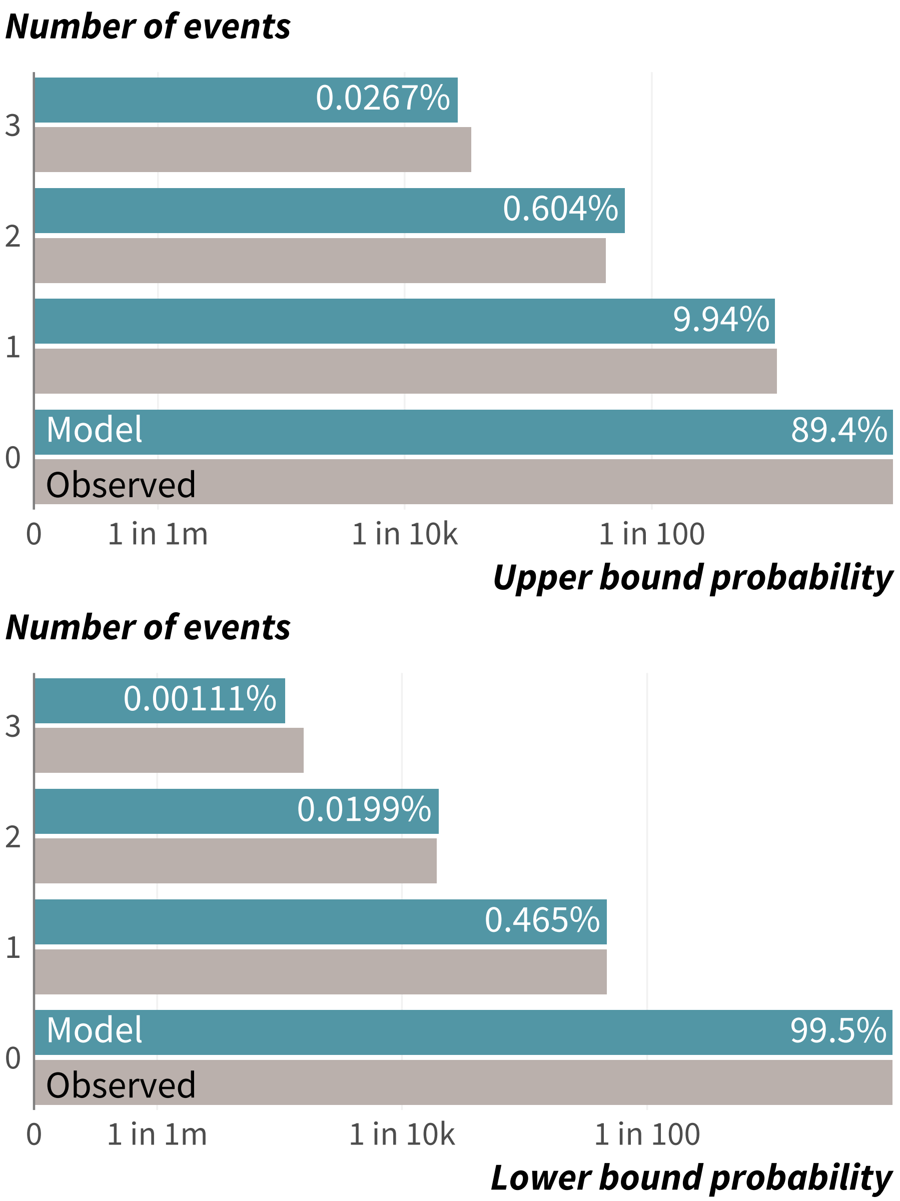

Figure 8: Upper and lower bound model for annual ransomware event frequency

"Wait a tick," one may argue. "That doesn't distinguish between companies that have weak vs. strong ransomware defenses. Wouldn't that affect likelihood?" This is a valid point. It's true that we don't know anything about the security posture of the victim organizations in our sample. Chances are it's a mixed bag. Maybe your firm's defenses are comparatively stronger and warrant a lower likelihood. Maybe the opposite is true. We have no way of knowing.

While we don't modify event frequency estimates by security posture, we can rightsize the output to be more in line with organizations similar to yours. Prior research in the IRIS series shows that incident frequency differs substantially by firm size, so we repeated the process outlined above to create separate models for the different revenue brackets[7] shown in Table 3.

[7] We decided not to develop a model for organizations under $10M in revenue because of insufficient data coverage and model reliability. Anytime we show stats for loss event frequency, they refer to organizations over $10M in revenue.

Revenue category	One or more	Two or more	Three or more
Upper Bound			
More than $100B	12.89%	2.50%	0.70%
$10B to $100B	9.70%	0.82%	0.07%
$1B to $10B	9.48%	0.65%	0.04%
$100M to $1B	9.83%	0.52%	0.02%
$10M to $100M	11.32%	0.69%	0.02%
Lower Bound			
More than $100B	9.37%	2.69%	0.90%
$10B to $100B	4.55%	0.42%	0.05%
$1B to $10B	2.73%	0.13%	0.01%
$100M to $1B	1.20%	0.06%	0.00%
$10M to $100M	0.36%	0.02%	0.00%

Table 3: Annual probability of a firm experiencing a given number of ransomware events

When reviewing Table 3, the first thing that jumped out to us is the relatively low variation among the different revenue tiers for having a single event (at least for the upper bound). Companies from $100M to $100B are pretty much all equally likely to be hit by ransomware. Interestingly, the smallest (<$100M) and largest ($100B+) organizations exhibit somewhat elevated probabilities that are surprisingly similar. To put that in perspective, the IRIS 2022 showed a roughly 250% difference in likelihood between those groups across all types of security incidents.

We attribute this to the nature of ransomware and the criminals behind it. Ransomware gangs are known to tailor their demands to fit the pocketbooks of the victim organization. Combine that with highly scalable distribution mechanisms, and ransomware is a threat that cybercriminals can adapt to any target irrespective of size. Table 3 suggests they're doing that well.

> The likelihood of experiencing multiple ransomware incidents escalates significantly with an organizations size, underscoring a drastic vulnerability among $100B+ corporations, which are 35 times more at risk of multiple events within one year.

This pattern shifts when comparing probabilities for more than one event per annum. The probability of two or three events increases at an increasing rate from the smallest to the largest organization's. Here we see that those $100B+ mega corporations are 35 times more likely to suffer three or more events in a single year than companies under $100M!

We'll go out on a limb and assume that all organizations—regardless of size—would like to minimize their odds of becoming a ransomware victim. Here's a few resources from CISA to help with that:

Stay ahead of the next outbreak with Pre-Ransomware Notifications.

Learn how to bolster your defenses using the vast resources at StopRansomware.gov.

Attack Frequency vs. Incident Frequency

After reading this section, some may be thinking something like "These numbers are way too low—we see a lot more ransomware than that!" First, it's entirely possible that these numbers are low relative to your organization's frequency of ransomware incidents. But it's also possible you're comparing apples and oranges.

If by "see," you're referring to ransomware attacks or attempts that are detected by your organization's defenses, this is not equivalent to the incidents or loss events studied in this report. Because the concepts are often conflated, we thought it would be helpful to demonstrate how they differ. To do that, we'll use a one-year snapshot of ransomware detections shared by Fortinet from their expansive array of sensors across the globe.

On average, each organization encountered dozens to well over 100 ransomware attempts per month during 2023. Since these events were detected and blocked by Fortinet devices, they constitute attacks rather than incidents. Had they gotten through to infect systems, we'd be analyzing them in the rest of this report rather than featuring them here.

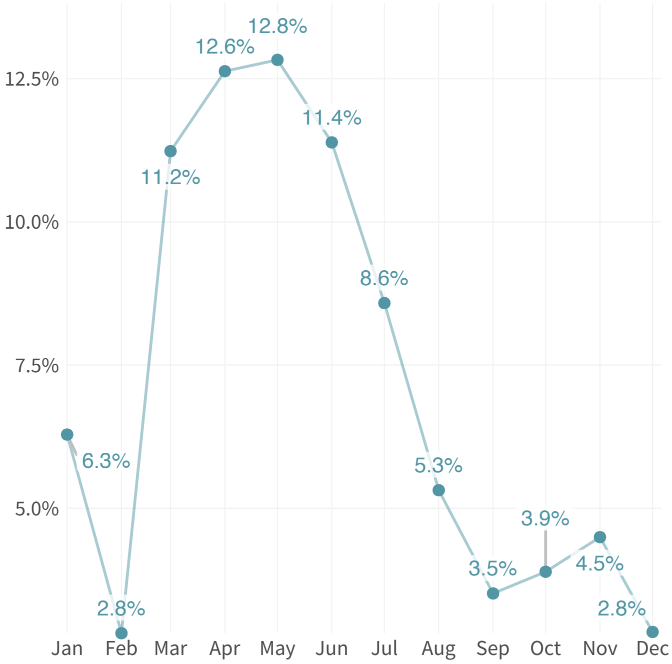

Figure 9: The average volume of monthly ransomware detections per organization

Figure 9 will cause some to wonder what's behind the mid-year rise and subsequent drop in that rate, but digging into such things is not our purpose here. Suffice it to say that attack frequency will fluctuate based on adversary campaigns, capabilities, and myriad other shifts across the threat landscape[8].

[8] Fortinet's Outbreak Alerts offer information on trending ransomware (or other) campaigns.

Ransomware Losses

"Boy that escalated quickly."

~Ron Burgundy, *Anchorman*

Now that we know how often ransomware events occur, it's time to evaluate how much they cost. There's been quite a bit of information shared on ransom demands, but that's only part of the total impact to organizations hit by these incidents. We'll start with observed losses from our historical dataset and then fit a distribution to those values to support ransomware risk models.

While reading this section, keep in mind that not all losses for all incidents become public. Certain types of losses are easier to identify from public records, such as class action suits and SEC Filings. Other forms of loss get absorbed internally with no outward expenditures and/or are simply difficult to quantify. We suspect the losses from major ransomware events are more complete than other types of incidents due to their disruptive and often public nature. Thus, we hold that our recorded losses suitably reflect known financial losses from publicly visible ransomware incidents.

Historical Loss Events

Financial losses tend to be less reported than other data points for cyber events—though the new SEC ruling requiring the disclosure of material incidents may change that. For now, we'll have to work with what's come to light about the impact of prior events. Thankfully, we have enough of those to establish a range of historical financial losses triggered by ransomware incidents.

Inflation-adjusted ransomware losses are shown in red in Figure 10 amid those recorded across all other types of incidents for comparison (in gray). Even on a log scale that diminishes the apparent length of the tail, it's easy to see that ransomware losses cluster toward the upper end of the range. The annotated statistics reinforce that point.

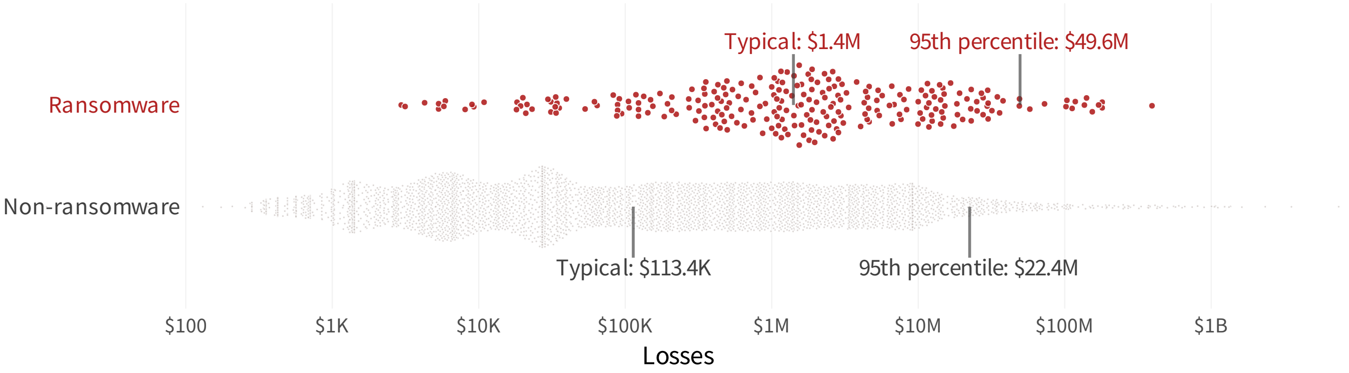

Figure 10: Distribution of reported losses from ransomware incidents (all-time)

> Dramatically, losses once considered "typical" just four years ago would be considered "small" by more recent standards, with 2023's 25th percentile loss matching 2019's "typical".

The typical ransomware event as measured by the geometric mean, costs $1.4M—over 12 times the typical financial impact for other types of cyber events! Extreme ransomware events are also substantially more impactful. The 95th percentile loss balloons to about $50M, compared with $22M for non-ransomware incidents.

These loss statistics cover the last five years as a whole, which begs the question of how ransomware costs are changing over time. Figure 11 plots losses reported each year, and we've overlaid violin plots to make it easier to eyeball the trend. The typical loss (annotated) clearly trends up, as does the 75th percentile (top horizontal line).

The most concerning trend, however, is that even the "smaller" losses are getting substantially bigger. Note how the 25th percentile loss (bottom horizontal line) in 2023 now lines up with the typical loss back in 2019. Not the trend we'd like to see, obviously.

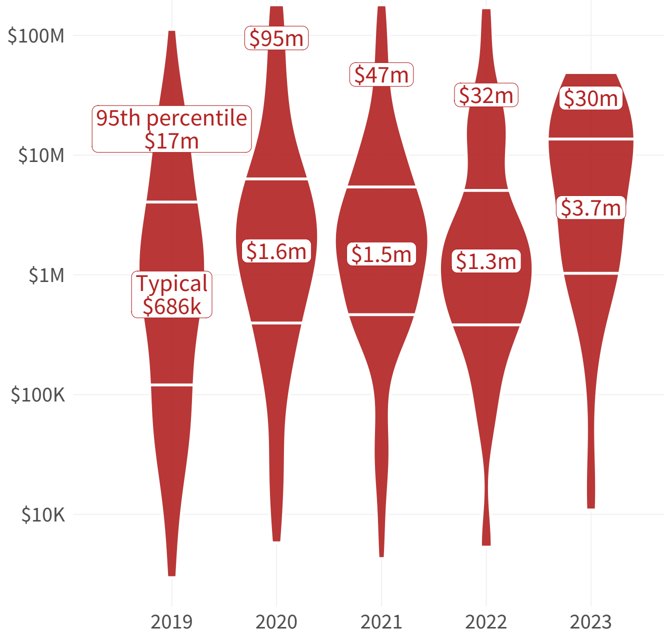

Figure 11: Annual distribution of reported losses from ransomware events

While the per-event costs of ransomware are trending up in recent years, its overall impact relative to other types of incidents has inflated even more (see Figure 12). Rewinding the clock a decade, ransomware accounted for less than 0.5% of recorded costs across all security incidents in 2014.

Based on losses tallied over the past few years, ransomware's share of financial impact regularly hovers around a third of all cyber events! It's no wonder the ransomware epidemic has attracted the attention of criminals and law enforcement alike.

Figure 12 (Right): The annual percentage of all reported cyber losses associated with ransomware

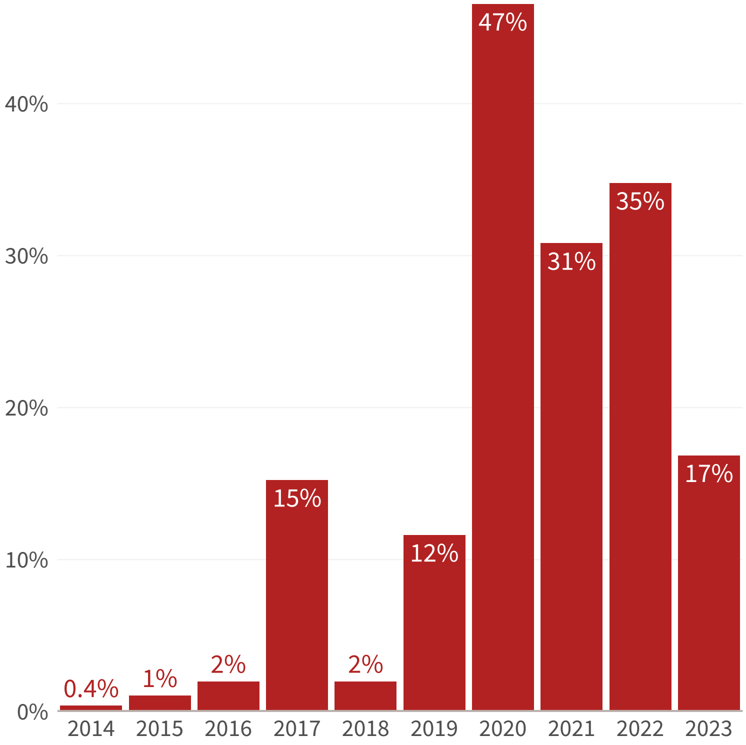

LOSSES BY SECTOR AND SIZE

We'd be remiss if we didn't provide some analysis of how ransomware impacts different types of organizations. As was done with frequency, we'll start with a comparison across industries. Figure 13 tees that up with ransomware's share of reported losses across all cyber events in each sector.[9]

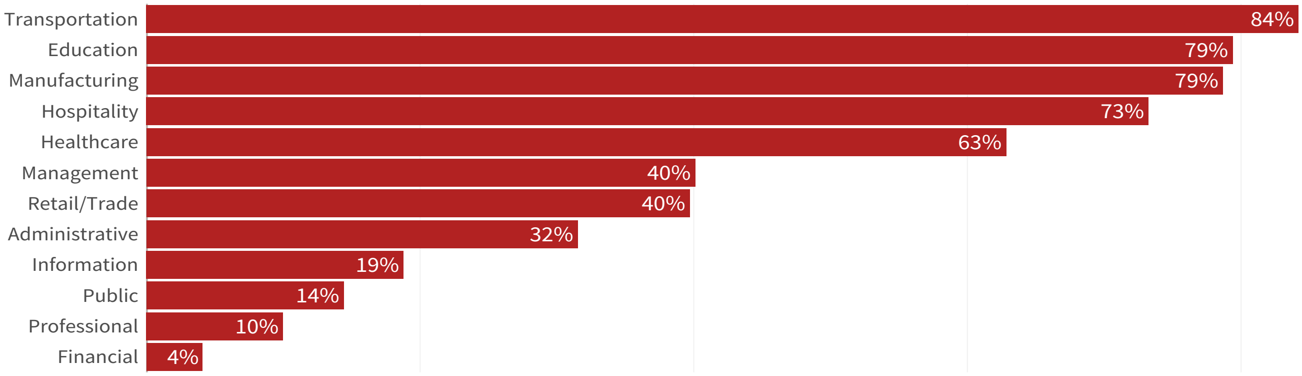

Figure 13: Percentage of all reported cyber losses from ransomware events by sector

YES—YOU'RE READING THE CHART CORRECTLY.

“ 80% of all known losses from all security incidents in the top four sectors come from ransomware.

[9] We removed sectors with low numbers of events with recorded financial losses.

It shows that ~80% of all known losses from all security incidents in the Transportation, Education, and Manufacturing sectors come from ransomware. Furthermore, ransomware claims roughly two-thirds of all costs recorded for the Hospitality and Healthcare industries. On the flip side, only a small proportion of all reported cyber losses for the Professional and Finance sectors is attributed to ransomware.

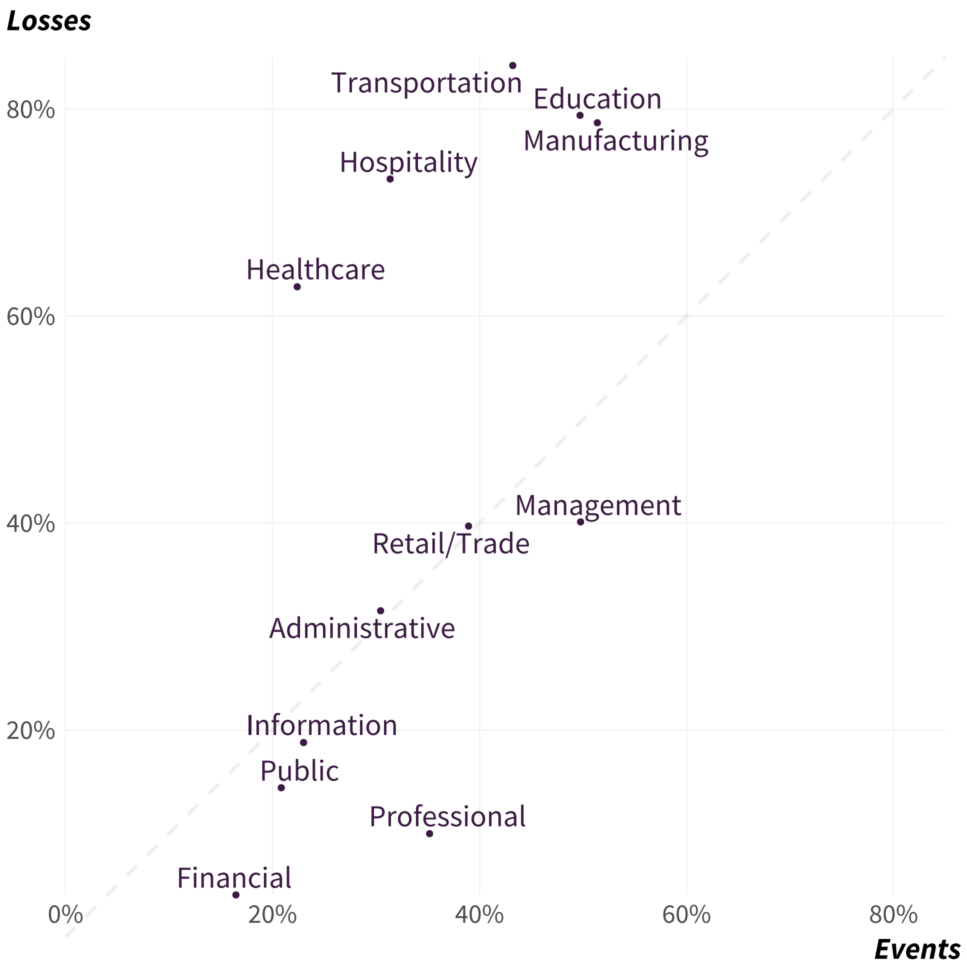

Figure 14: Percentage of cyber events and losses tied to ransomware by sector

If you're like us, you're now tempted to scroll up to see how these industries stack up in terms of event frequency. We'll save you (and us) the effort. Figure 14 plots each sector according to the share of events and losses over the last five years attributed to ransomware.

If frequency and losses were perfectly correlated, sectors would lie on or near the dashed line. In general, that's not the pattern we see here. Instead, we see industries that are very disproportionately impacted by ransomware relative to event frequency (e.g., Healthcare, Hospitality), while the opposite is true for others (e.g., Financial, Professional). A myriad of factors contribute to the placement of sectors in Figure 14, but the targeting strategy of ransomware gangs is likely a major driver among them.

> Ransomware creates an uneven battlefield, sparing $100B+ giants with less than 1% of losses, while midsize firms face catastrophic costs, shouldering up to 50% of cyber-related losses, particularly in vulnerable sectors like Healthcare.

We can also compare ransomware's proportional losses by organization size. Ransomware accounts for less than 1% of cyber-related losses recorded for $100B+ mega corporations. Contrast that with 50% of all costs tallied for midsize organizations spanning $100M to $1B in annual revenue!

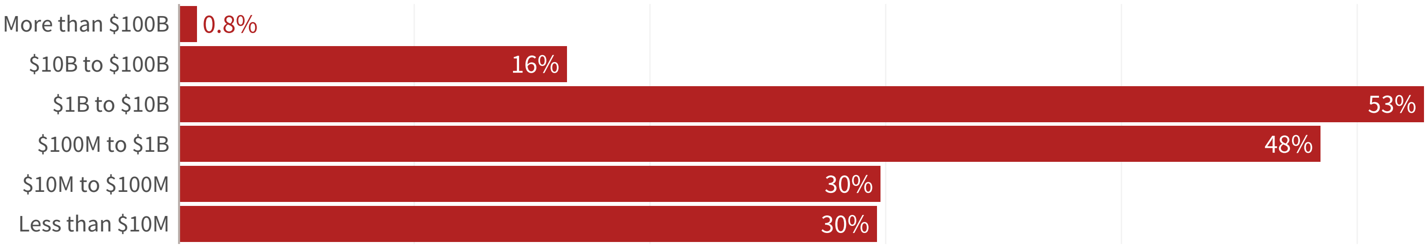

Figure 15: Percentage of all reported cyber losses from ransomware events by organization revenue

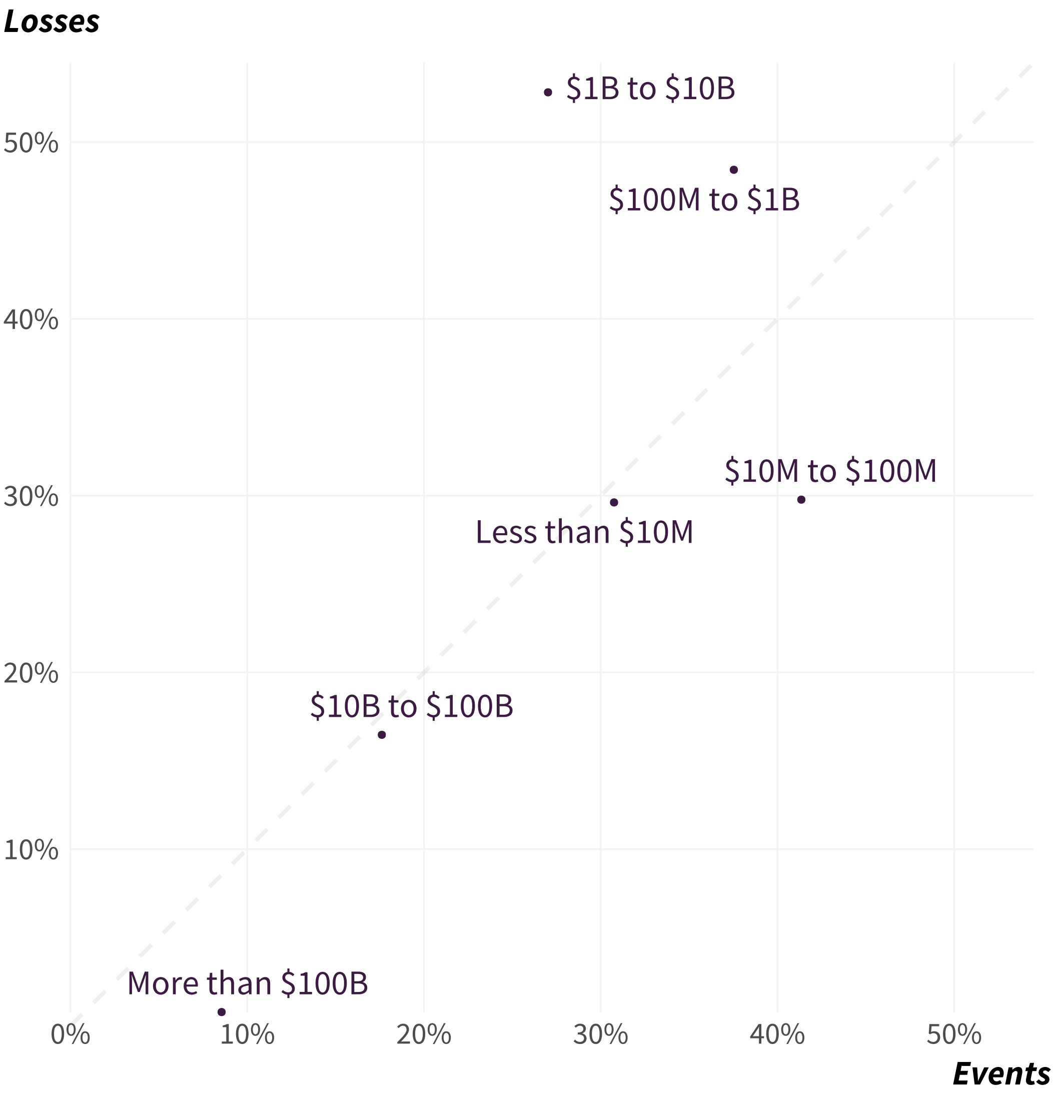

Since these findings also prompt the urge for back scrolling, we'll do another scatterplot. Proportional frequency and losses are a bit more correlated here than in the sector-based version. But the revenue tiers definitely don't follow in order from lower left to upper right. On the whole, this view strengthens the perception that midsize firms are particularly at risk from ransomware.

Figure 16 (Left): Percentage of cyber events and losses from ransomware by organization revenue

Ransomware Loss Exceedance

Given what we've learned thus far about ransomware frequency and losses, is it possible to answer questions like "What's the probability that we'll lose $10M or more over the next year from ransomware incidents?" Yes, it is!

One way of answering such questions is to create an exceedance probability curve (EP Curve), more commonly known as a loss exceedance curve (LEC) among cyber risk professionals. The purpose of LECs is to demonstrate the probability of experiencing a minimum amount of loss in a given time period. This can be very useful for supporting risk decisions and mitigations.

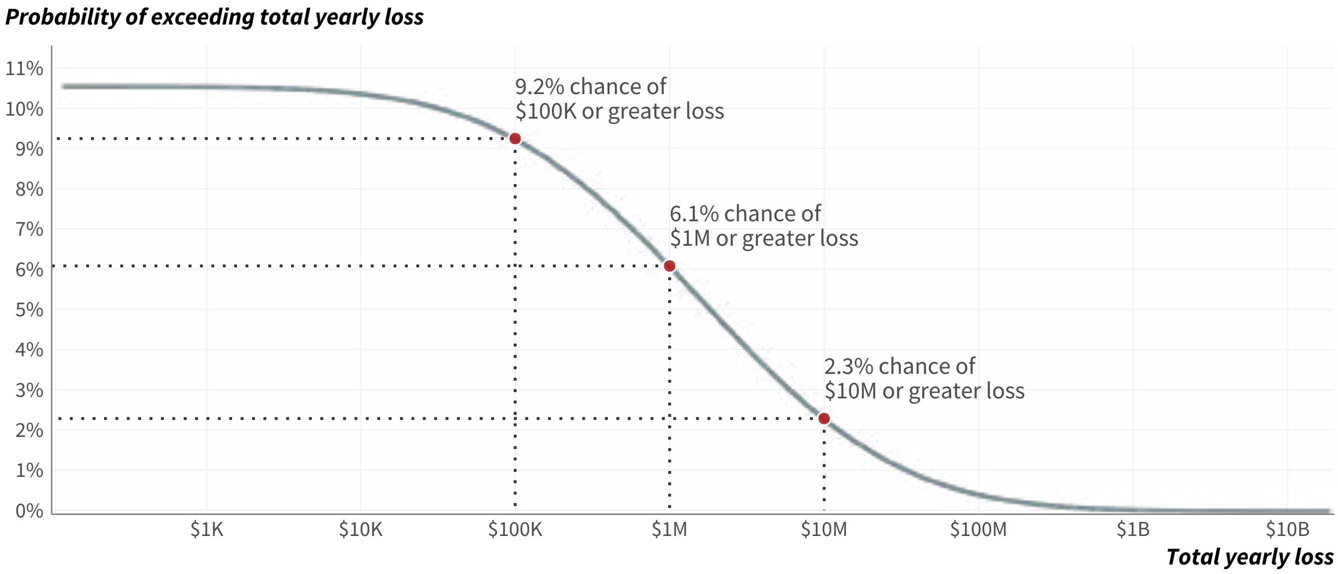

Figure 17: Ransomware Loss Exceedance Curve

Figure 17 presents the results from a simulation of frequency and loss data to produce a ransomware LEC for a typical organization. Trace any point on the curve to the x and y intercepts to determine exposure. For example, there's a 2.3% chance that a firm's ransomware-related losses will exceed $10M in a year.

We emphasize "typical" in the preceding paragraph to remind readers that this LEC represents a composite view across numerous firms of various types, sizes, and security postures. A LEC for your organization or its peers would certainly look different.

TOTAL COST OF RANSOMWARE

FAIR WARNING: WE'RE ABOUT TO DO SOMETHING WE'VE NEVER DONE BEFORE (BUT IT'S NOT *TOO* CRAZY).

As mentioned previously, verifiable financial losses are recorded for only a subset of the more than 14,000 ransomware incidents in our dataset. Adding up only those known losses would be trivial but would vastly underestimate ransomware's total impact over the last five years. So, we won't even put that number out into the eternal memory of the internet. But still—it sure would be nice to have such a number, huh?

We think so too, which is why we're going to break our tradition of sticking rigidly to hard data on prior events. But don't worry—we're not going to abandon our core principles and ride the trolley into the Neighborhood of Make-Believe. We won't force the data anywhere it doesn't lead.

There's no strong reason to believe that the ransomware incidents for which we have recorded losses are significantly different than those for which we do not[10]. That means it is reasonable to apply the distribution of known losses to the events for which we have no known losses. Note that this doesn't mean simply adding the average loss or always assuming the worst case. It means properly sampling estimated values based on the distribution parameters—standard stats stuff.

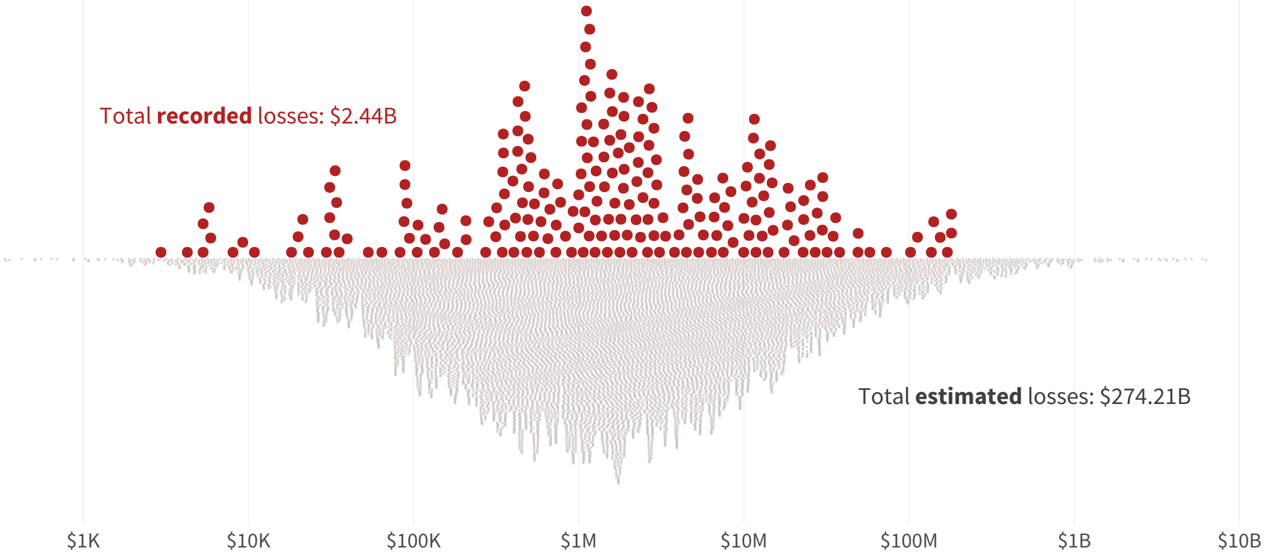

Figure 18: Projected total cost of ransomware events based on known losses

> In the last decade there has been a 140X increase in projected losses, with 2023's damage nearing $95B—an alarming trend that's only expected to rise.

[10] Though some of the most common ransomware strains in our incident dataset do tend to be those with higher ransom demands (e.g., Clop). It's possible the costlier strains are overrepresented in our dataset, which would create an upward bias for our total loss estimates. Analysis of ransomware strains follows later in this report.

Using this method, we calculate the total financial losses from publicly known ransomware incidents over the last five years to be about $276B dollars. We suspect that that estimate will be perceived by some as outrageously high and scandalously low by others. The lack of consensus is understandable; it's difficult to find a comparable statistic to determine how this estimate aligns (or not) with that of other sources.

Figure 19 breaks down that five-year total on an annual basis and stretches back another five years before that. This supplies the astounding observation that the total estimated losses from ransomware events increased around 140X over the last 10 years! 2018, 2020, and 2021 stand out in terms of major leaps in aggregate loss. While not quite the high-water mark set in 2021, ransomware's projected impact for 2023 stands at almost $95B (a figure we expect to climb still higher as new data emerges).

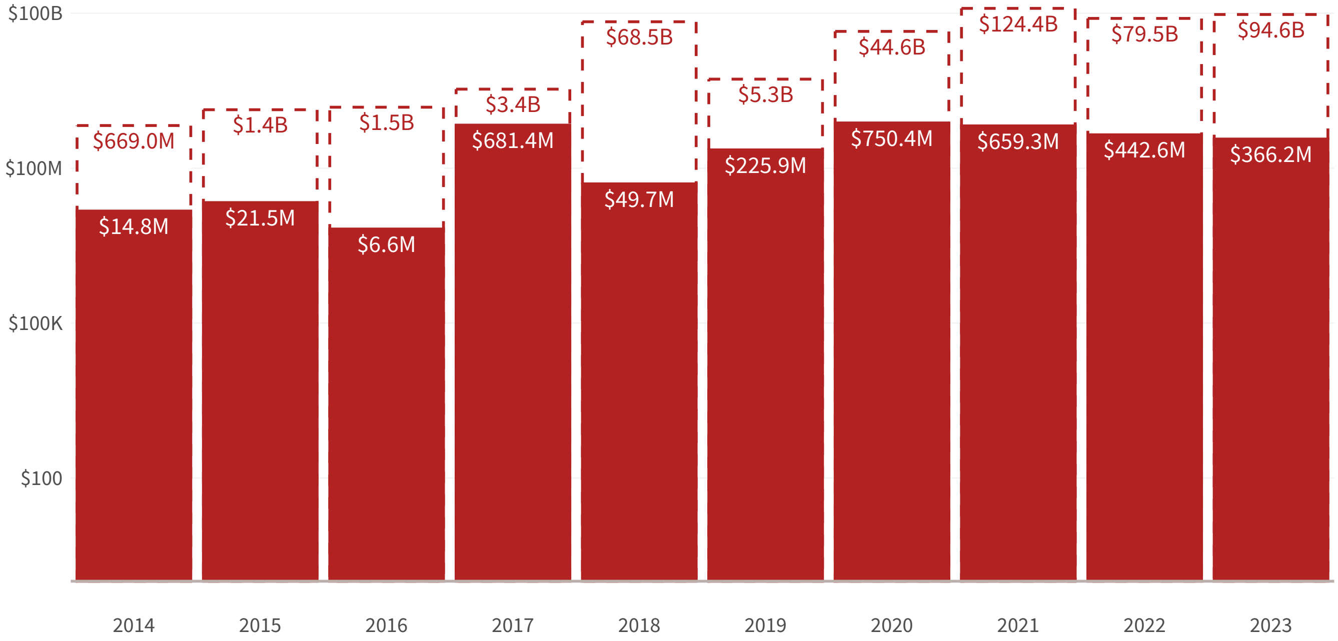

Figure 19: Annual projected total cost of ransomware events based on known losses

As large as these projections are, they likely represent a conservative estimate from a geographic standpoint. The dataset on which we're basing that projection is global, but coverage is most comprehensive for incidents involving organizations with a presence in the United States. Plus, the losses reported for those events tend to be the direct financial costs that often don't capture the full extent of impact on the affected organization.

Caveats aside, however, the overall takeaway still stands. Ransomware is a major financial drain on industry and the overall economy.

Ransomware Campaigns & TTPs

"If you find yourself in a fair fight, your tactics suck."

John Steinbeck *(we imagine this posted prominently in cybercriminal office breakrooms)*

Aside from focusing exclusively on ransomware, everything we've served up to this point has been standard "bread and butter" IRIS fare. A great deal of far more technical analysis exists out there regarding the actors behind ransomware, the most prevalent variants, and the common tactics, techniques, and procedures (TTPs) they employ. Much of that is deeper than we care to go for our risk-oriented purposes, but we would like to at least create a bridge to that domain for those who want to cross over to it.

We'll build that bridge using three components:

- Open-source intelligence on ransomware campaigns collected by Tidal Cyber
- Most prevalent ransomware strains from our core dataset and Ransomwatch
- Top MITRE ATT&CK techniques based on ransomware in our core dataset

Top Ransomware Strains

Figure 20 lists the most common ransomware identified in events from our core incident dataset. There's no shortage of resources and intelligence on these ransomware, so we're not even going to attempt to summarize all that here. Instead, we'll just add a few general notes to guide your interpretation of these results and external research.

First, check the timeframe before you start analyst-splaining about how some of the ransomware gangs in scope here are now defunct. This list is based on incidents that occurred from 2019–2023 and contains a mix of gangs and strains. Some are no longer active and some have become more/less active over time.

Second, remember the nature of the dataset. It's based on events impacting individual organizations. Not campaigns, malware detections, number of variants, ransom payments, or even simple infections that were handled internally. These led to real incidents that, for one reason or another, became publicly known.

Third, keep in mind that many factors contribute to the prevalence of ransomware shown here. For example, the dominance of the Cl0P (aka CLOP, TA505) ransomware gang is largely due to its exploitation of the infamous "MOVEit" vulnerability in 2023. Such attacks are far more scalable to a large population than more targeted/bespoke campaigns.

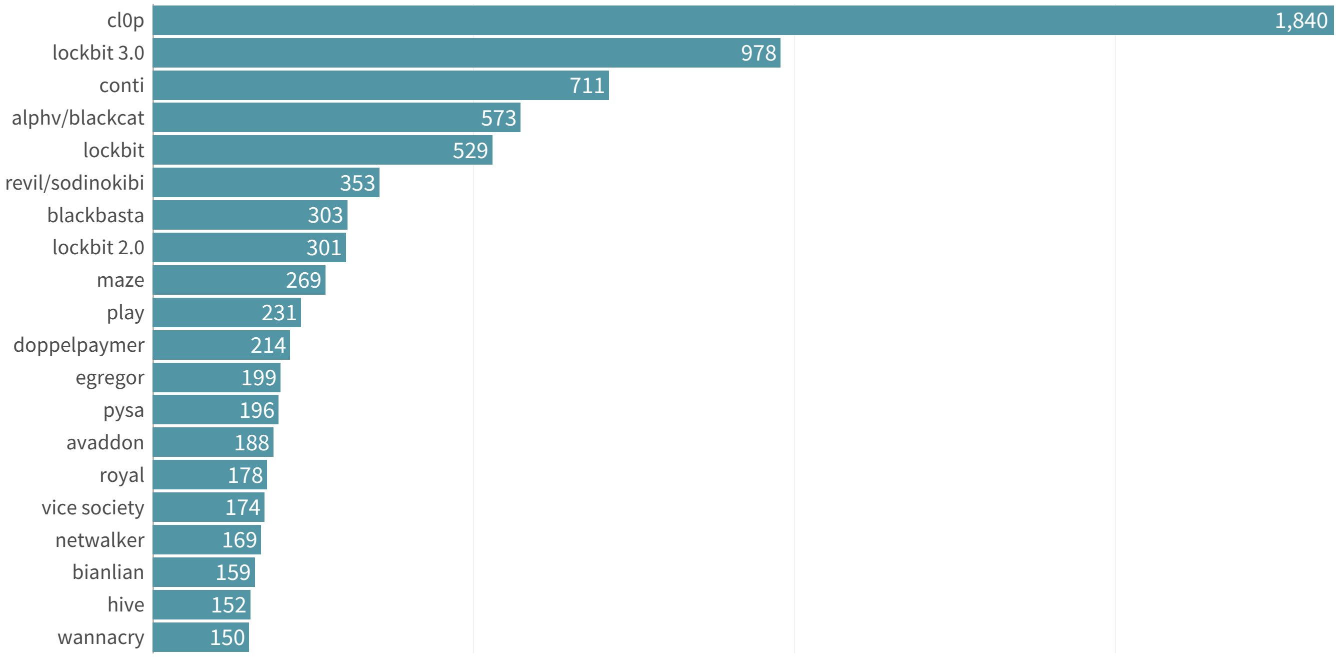

Figure 20: Most common ransomware observed among incidents (2019–2023)

To offer a more temporal perspective on ransomware strains, we'll turn to Ransomwatch. This project trails the extortion sites used by ransomware groups and surfaces an aggregated feed of claims (as in claiming responsibility; not insurance claims). This makes it an inherently more timely accounting of ransomware activity than our core incident dataset.

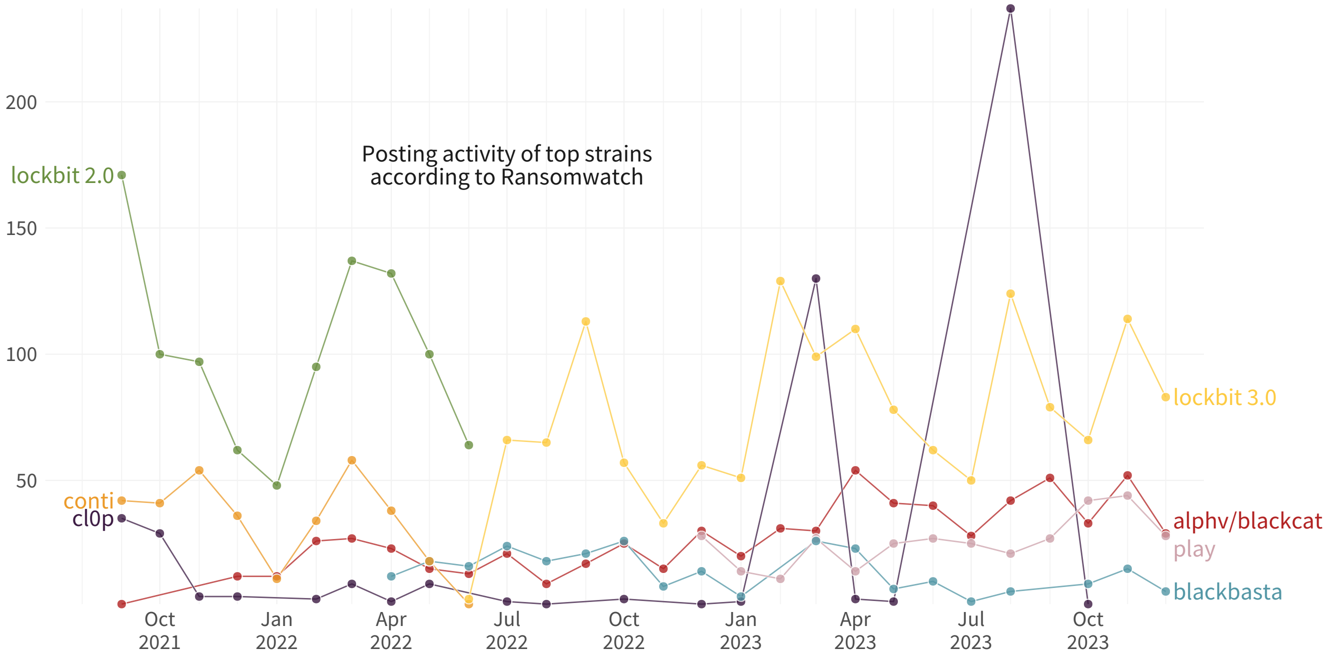

Figure 21: Trending of top ransomware strains as tracked by Ransomwatch

Figure 21 trends the activity associated with the top ransomware from Figure 20 through the lens of Ransomwatch. Here, we see the death of LockBit 2.0, the reincarnation of LockBit 3.0, the 2023 exploits of Cl0P, etc. Again, we won't go into details on these campaigns in this report, but you're more than welcome (even encouraged!) to do some homework.

Top Ransomware Techniques

MITRE ATT&CK is quickly becoming the common language of adversary TTPs used across the cybersecurity industry. A major benefit of ATT&CK is that it enables readers to easily find definitions and examples of each technique referenced, as well as explore a wealth of information on associated threat groups, malware, mitigations, attack simulations, etc.

Unfortunately, public disclosures or media coverage of security incidents rarely come with a detailed list of the ATT&CK techniques involved. Evidence collected via a digital forensics investigation is generally needed for that. However, through a combination of analytical techniques, we've managed to do some level of ATT&CK-ification on over 40% of the ransomware events in our dataset.

We've organized identified techniques into three key stages of an incident: initial access, post-compromise (execution through lateral movement), and exfiltration and impact. In this section, we'll present the most frequent and impactful techniques for each of these stages.

INITIAL ACCESS

The Initial Access tactic describes techniques used by adversaries to gain an initial foothold within a target victim environment. Of all techniques in ATT&CK, these are probably the most familiar (who hasn't heard of phishing?). Nevertheless, it's important to understand these trends because repulsing attacks at this stage avoids the many problems and costs that ensue once they gain access.

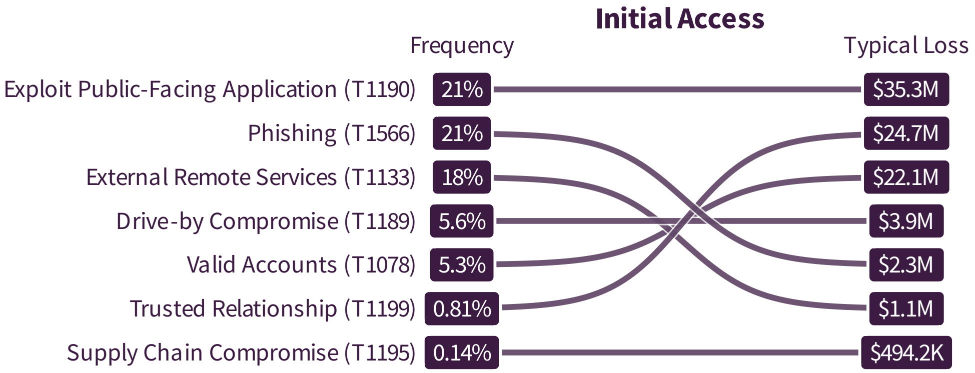

Figure 22: Relative frequency and losses for observed initial access techniques

Exploiting external, public-facing, applications is tied with phishing for the #1 spot on the frequency side of Figure 22. But the former typically carries substantially higher losses. Attacks that exploit trust relationships with third parties occur less often but punch well above their weight in terms of average impact. This is one of the few lists we've seen where the exploitation of valid accounts isn't near the top of the list, though it does round out third place in the losses column.

POST-COMPROMISE

Let's go beyond the initial intrusion and take a look at the TTPs utilized by adversaries to maintain presence, escalate privileges, spread across the internal network, evade security defenses, establish command and control channels, and other nefarious activities. Figure 23 provides a breakdown of the top post-compromise techniques we were able to identify.

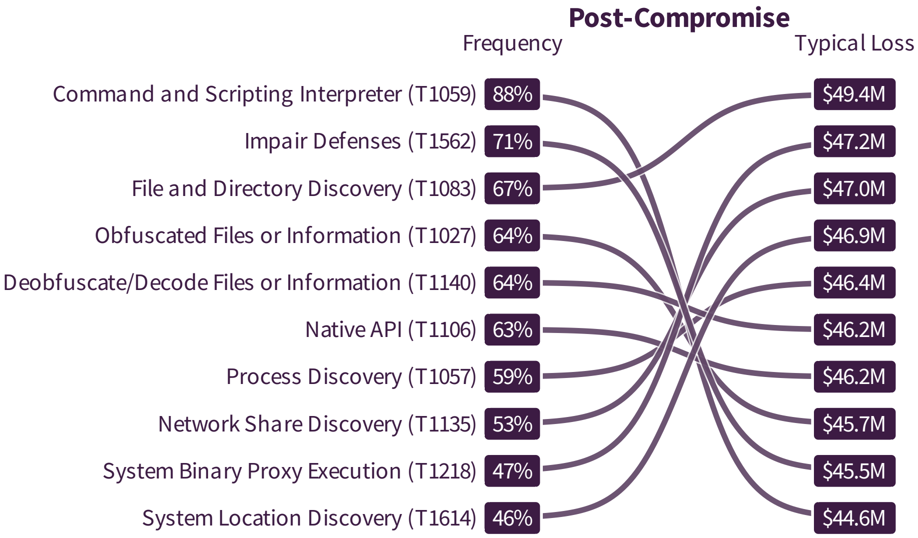

Figure 23: Relative frequency and losses for observed post-compromise techniques

Compared with the initial access group, there's decidedly less variation in frequency and losses among the top 10 post-compromise techniques. This is largely because many of the major ransomware groups incorporate many of the same basic functions, being derived from or patterned after successful strains that came before. All these techniques won't necessarily be used in every incident, but the capabilities are there when needed.

EXFILTRATION AND IMPACT

Adversaries generally have some ultimate goal in mind when carrying out an attack, and those behind ransomware schemes are no different. While some stick to the classic playbook of infect > encrypt > extort, others prefer to siphon off data for double/triple extortion or disrupt the entire network. Whatever their ends, the means are captured under ATT&CK's exfiltration and impact tactics.

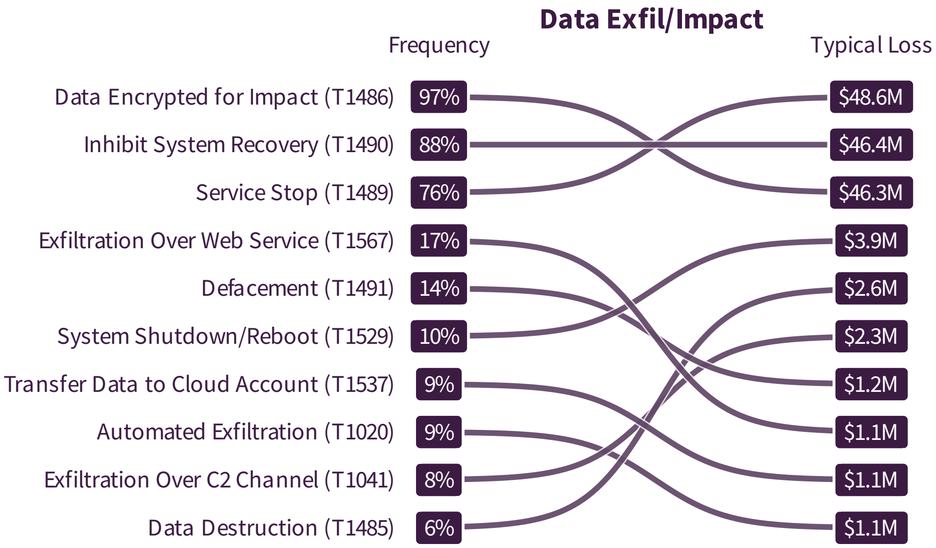

Figure 24: Relative frequency and losses for observed Exfiltration & Impact access techniques

In an unforeseen twist, the most common impact associated with ransomware is . . . encryption. We're obviously kidding; it goes without saying. The two biggies beyond that both seek to undermine the defensive services and recovery capabilities of the infected system so ransomware can sink its hooks deep. The presence of several data exfiltration and transfer techniques reflects the popularity of double/triple extortion schemes.

TTP Time Travel with Tidal Cyber

As part of its mission to enable businesses to implement a threat-informed defense, Tidal Cyber collects open-source intelligence on adversary campaigns and TTPs. They were kind enough to share some of the fruit of those efforts with us specific to ransomware groups to include in this study.

Figure 25 ranks the top ATT&CK techniques attributed to ransomware campaigns over the last four years. We've removed T1486 (Data Encrypted for Impact) since it's pretty much a given for ransomware and always on top.

Beyond encryption capabilities, T1059 (Valid Accounts) is the undisputed champion among ransomware techniques, according to OSINT. That's noticeably different from what we see in Figure 22 and a good reminder of the value of using multiple sources for this kind of analysis. Speaking of, see our report Multi-Source Analysis of Top MITRE ATT&CK Techniques for an in-depth review of observed TTPs from 20+ sources.

We find it interesting how techniques appear to consolidate over time in Figure 25. That timeframe roughly overlaps the "professionalism of ransomware phase from way back in Figure 3. It seems to tell a story of ransomware gangs honing what works and incorporating or purchasing successful capabilities for their own campaigns.

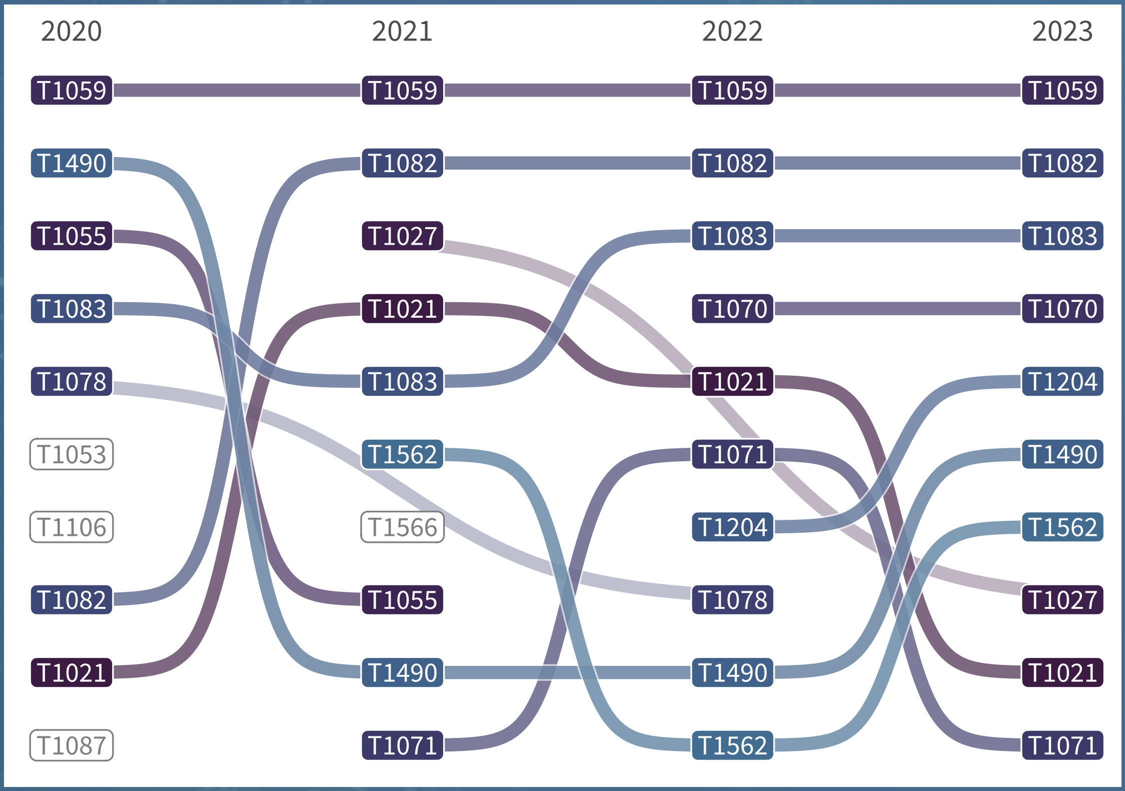

Figure 25: Ranking of ATT&CK techniques over time based on adversary campaigns tracked by Tidal Cyber

Conclusion

A principal aim of this IRIS report is to bring data and sound analysis to bear on organizations' cyber risk questions. To that end, insights from this special edition on ransomware support the following observations:

Without accounting for the (potentially large) effect of one's sector, cyber risk managers should start their cost-benefit analysis around preventing and mitigating ransomware at a 1-in-10 chance of experiencing at least $100K in losses from ransomware over the next year and a 1-in-20 chance of experiencing at least 10x that.

Larger organizations are more likely to experience one or more ransomware events in a given year than mid-size and smaller firms—probably because they're more likely to experience security incidents in general. However, events experienced by mid-size and smaller firms are more likely to be ransomware, and ransomware represents a markedly higher share of their cyber losses than those of larger organizations.

When sufficiently incentivized, malicious actors are willing to shift tools and tactics in a manner that can completely upend an otherwise relatively stable threat landscape. Such incentives appear to be financial gain, ease of monetization, and the perception of little risk of repercussion. The same incentives promote standardization and consolidation of techniques over time, potentially simplifying—if not easing—the job of defenders over time.

Software weaknesses lead to large losses. Specifically, ransomware actors exploit public-facing applications in 21% of the cases, leading to the highest typical per-event losses at $35.3M. They exploit external remote services 18% of the time, leading to $1.1M in losses per event.

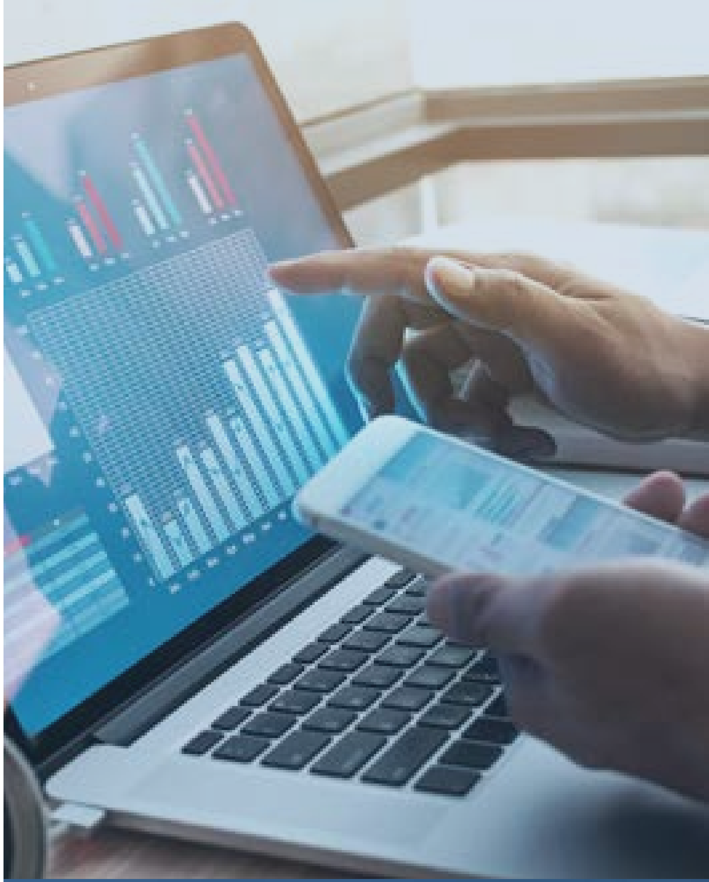

What specific actions should organizations take to prevent and mitigate ransomware today? The analysis of top TTPs may be of some use when considering specific scenarios or controls, but allow us to direct you to stopransomware.gov for a more complete answer. There you'll find ample resources to help you in preventing ransomware while also preparing for the worst with mitigation, response, and reporting guidance. You'll also find information about major actors, sector-specific guidance, relevant agency alerts, and no-cost services to support anyone in guarding against ransomware.

But we'd be remiss if we stopped at only advocating for updated practices among would-be victims. We need to look more closely at the software the malicious actors exploit. While many ransomware attacks take advantage of exposures specific to the affected firm, some of the most impactful attacks have taken advantage of defects in widespread software packages. Thus, a more complete response to ransomware requires action from the software manufacturers who build and maintain our digital world as well. CISA's Secure by Design initiative is one such effort, seeking to rebalance the burden of staying cyber safe to those most capable of doing so, namely the software manufacturers. Products that were secure by design would not be vulnerable to the most common recurring classes of defect that we see today. The result would be safer products, higher costs and lower profits to ransomware actors, and fewer successful ransomware incidents.

Appendix A: Incident Pattern Descriptions

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

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 online brand.

ACCIDENTAL DISCLOSURE: Information inadvertently made accessible to unauthorized parties by exposing data stores, publishing private info, emailing wrong recipients, etc.

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

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

Appendix B: Upper and Lower Bound Models

As mentioned earlier, estimating the probability or expected frequency of security incidents requires a known sample of organizations on which to base calculations. Unfortunately, we don't have a reliable count of active firms relevant to this dataset around the world. But we have a couple proxies that can be used as a basis of reasonable lower and upper bound estimates.

> **LOWER BOUND:** This includes all registered organizations in the United States according to Dun & Bradstreet (because we don't have numbers for the whole world). This assumes that incident frequency among the U.S. firms is similar to that everywhere else, which is certainly not the case. But it's a good starting point, even if you don't work for a U.S. firm. We call this the lower bound because it assumes that all registered firms engage in activities that subject them equally to the kinds of incidents found in this dataset. We don't believe that to be the case.

> **UPPER BOUND:** This includes all organizations recorded in our dataset, which means these organizations have experienced a known incident at some point in the past. While that's clearly not the case for all organizations, this upper bound approach is based on the premise that not all firms are equally subject to the kinds of incidents contained in this dataset (i.e., perhaps they don't use IT or aren't subject to incident disclosure regulations). This assumes that all firms prone to incidents have already had one incident, thus likely resulting in overestimation.

The "just right" (Goldilocks) zone is, of course, somewhere in the middle. It's impossible for us to know exactly where your organization falls between the lower and upper bounds, so we've opted to share both to support your assessment. In general, the upper-bound offers a more risk-averse view with higher values. Choose one or fuse both to suit your organization's risk posture and tolerance.

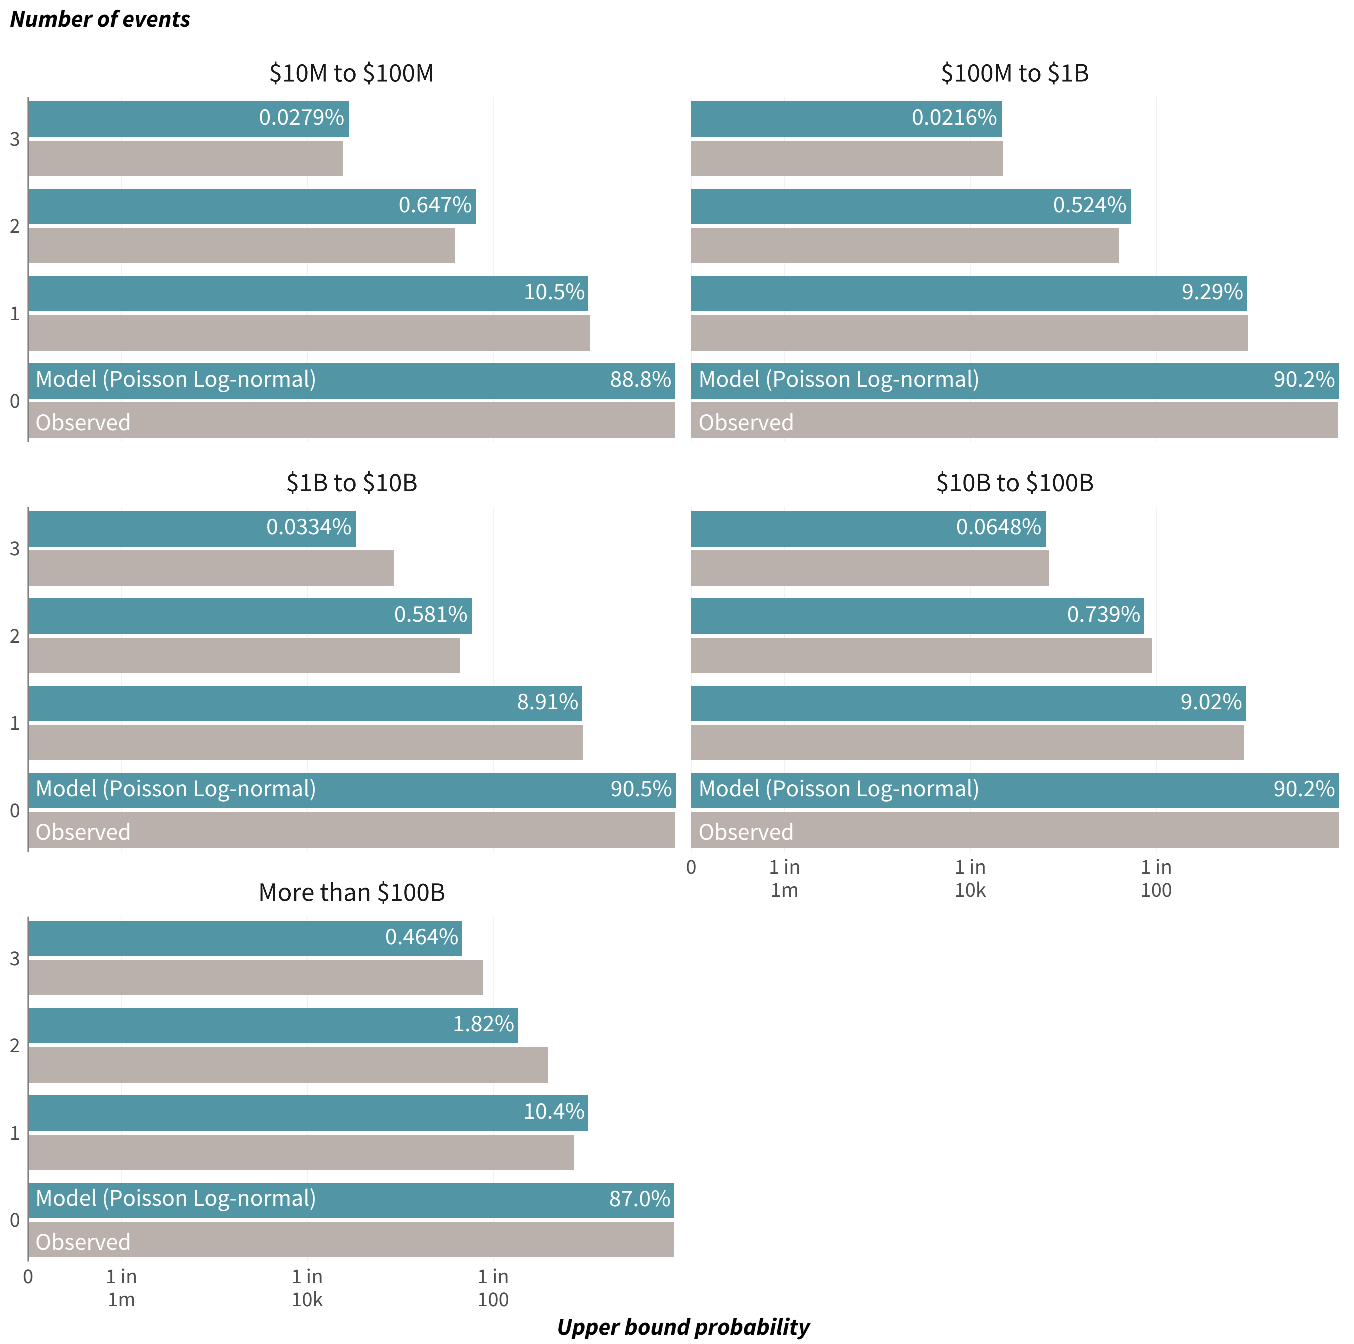

Figure B1: Upper bound model for annual ransomware event frequency by revenue category

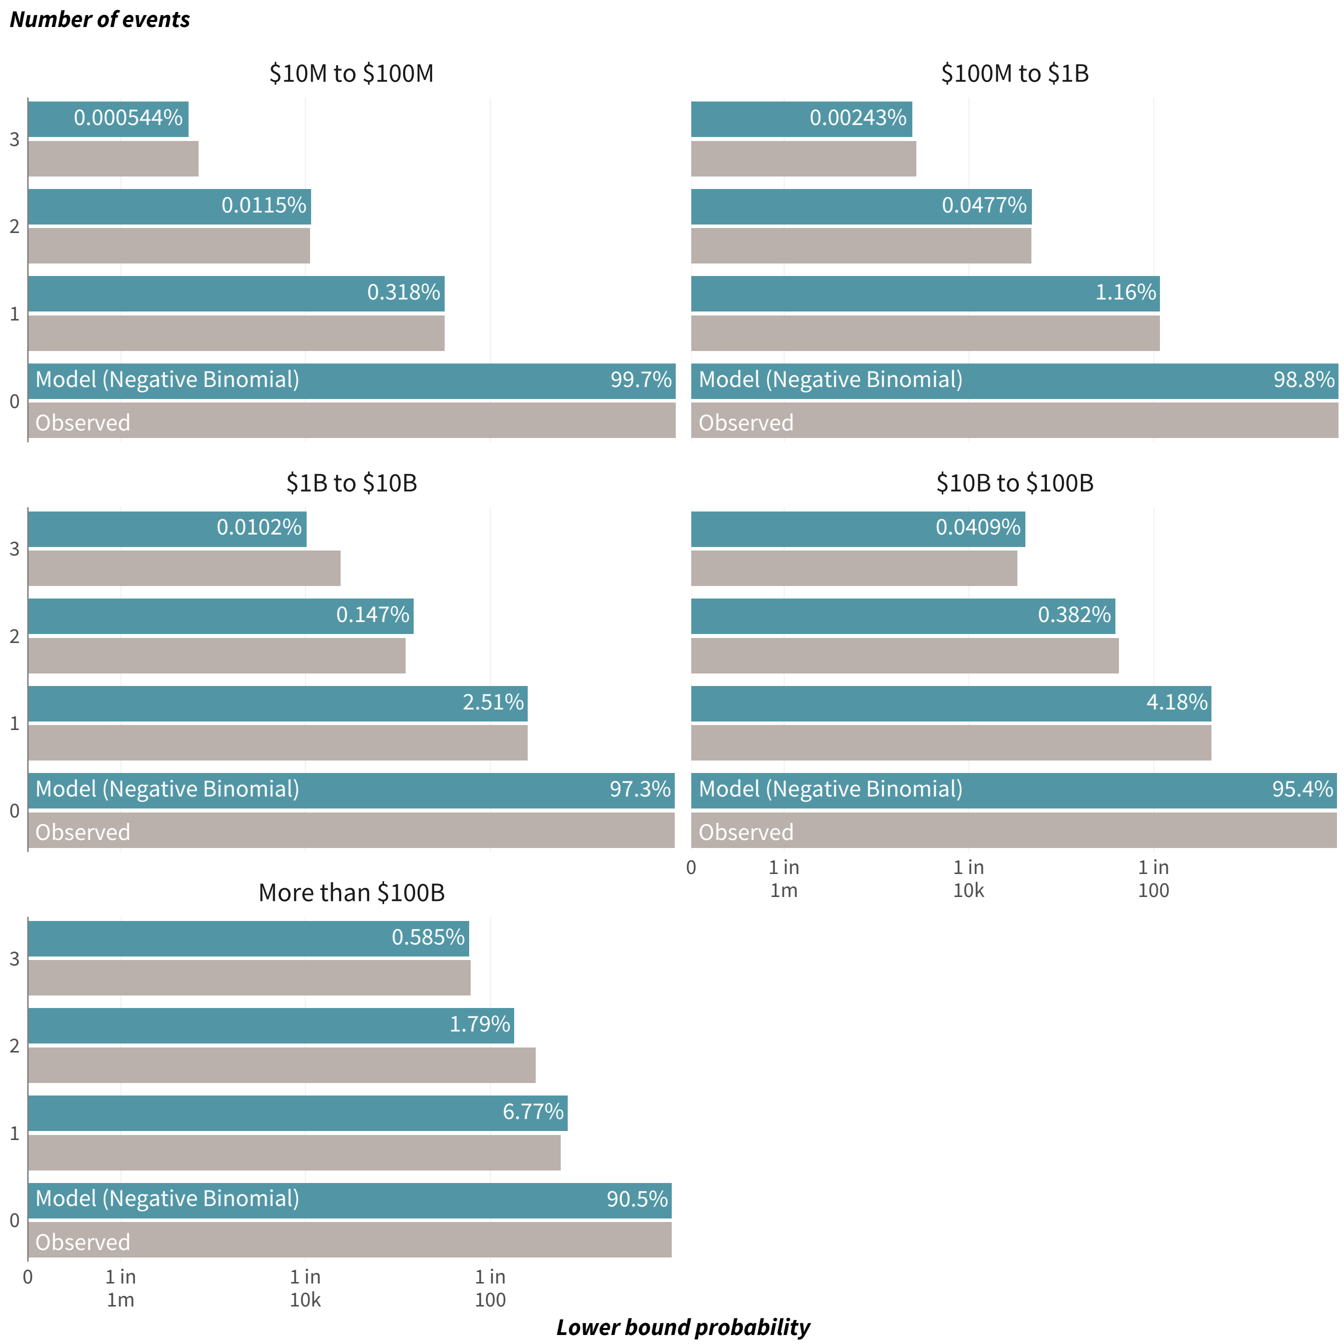

Figure B2: Lower bound model for annual ransomware event frequency by revenue category