

Nonprofit

A BASELINE ANALYSIS OF
CYBER RISK FACTORS IN
THE NONPROFIT INDUSTRY

Mar 2022



IRIS Risk Retina Nonprofit by Cyentia

A Clearer Focus on Cyber Risk

Welcome to this special [IRIS Risk Retina](#)® focused on the nonprofit industry! IRIS Risk Retina is a service from the [Cyentia Institute](#) built on the highly-regarded [Information Risk Insights Study](#) (IRIS) series. The goal of Risk Retina is the same as that of the IRIS research—to offer a clearer focus on cyber risk through real-world data and rigorous analysis.

To succeed in this mission, we've partnered with [Advisen](#) to leverage their [Cyber Loss Data](#) containing ~100,000 historical events collected from publicly verifiable sources. It's the most comprehensive incident dataset we've seen and is used by many cyber insurers and reinsurers for that reason. Upon that solid foundation, we add our own supplemental research, data science techniques, and security expertise to provide the analysis presented in this IRIS Risk Retina.

Since most cyber risk frameworks revolve around the frequency and financial impact of security incidents—but don't provide actual data—we direct much of our attention there. You'll find distributions and key parameters to support any cyber risk quantification (CRQ) process or platform. We also examine common types of incidents that affect nonprofit organizations along with the industry's propensity for extreme cyber loss events. Let's get to it!

About Cyentia

The Cyentia Institute is a research and data science firm working to advance cybersecurity knowledge and practice. We pursue this goal through our research publications and analytic services like [IRIS Risk Retina](#).

If you would like to see the analysis provided in this IRIS Risk Retina focused on your own sector(s), please contact retina@cyentia.com.



A Glimpse Through This Risk Retina:

- 3** At a Glance
- 4** Loss Event Frequency
- 8** Loss Magnitude
- 12** Quantifying Risk Exposure
- 14** Tail Risk Analysis
- 15** Common Incident Patterns
- 17** Viewing Other Dimensions
- 19** BetaPERT Parameters

Risk at a Glance

Cyber Risk Quantification Parameters for the Nonprofit Industry

This IRIS Risk Retina® focuses on incidents experienced by nonprofit organizations over a ten year period from January 2012 through December 2021. This date range gives a sufficiently-sized sample of 1,994 loss events for analysis while also remaining relevant for organizations managing present day risks. Here's the key stats to support cyber risk quantification that we'll expand and explore in the sections that follow.

The upper-bound average annual probability of a nonprofit organization experiencing one or more cybersecurity incidents that become publicly known is 12.5%.



Annualized probability

12.5%

of one or more events

Reported financial losses stemming from incidents impacting nonprofits vary widely around a geometric mean of \$145,000. The 95th percentile impact runs \$9M.



Typical loss magnitude

\$145K

in nonprofit incidents

Based on frequency and loss models for the nonprofit industry, a typical organization has less than a 1% chance of losing more than \$10M in a single year.



<1% chance of losing

\$10M

or more in 1 year

But cyber risk has a very long tail of rare but highly damaging events that demand attention. The 95% Tail Value at Risk for a large nonprofit exceeds \$350M.



Tail Value at Risk

\$350M

at the 95% level

System intrusions are far and away the most common and costly type of incident in nonprofits. They account for over half of all events and two-thirds of total losses.



Intrusions behind

55/66

% of events / losses

Loss Event Frequency

In our journey to better assess the risk posed by cybersecurity incidents, we first explore how often they occur. Our initial step is to examine historical loss events affecting the nonprofit industry. We then evaluate the number of events expected for a given organization over a period of 12 months. Our ultimate goal is to develop a model along with the associated parameters for estimating loss event frequency in nonprofit organizations.

Standard IRIS Risk Retinas generally focus on a top-level sector as defined by the [North American Industry Classification System](#) (NAICS). According to NAICS, however, nonprofit organizations are a subsector of [Other Services](#) that includes religious organizations, social advocacy and human rights organizations, professional associations as well as other civic and social organizations. ([NAICS code 813](#)).

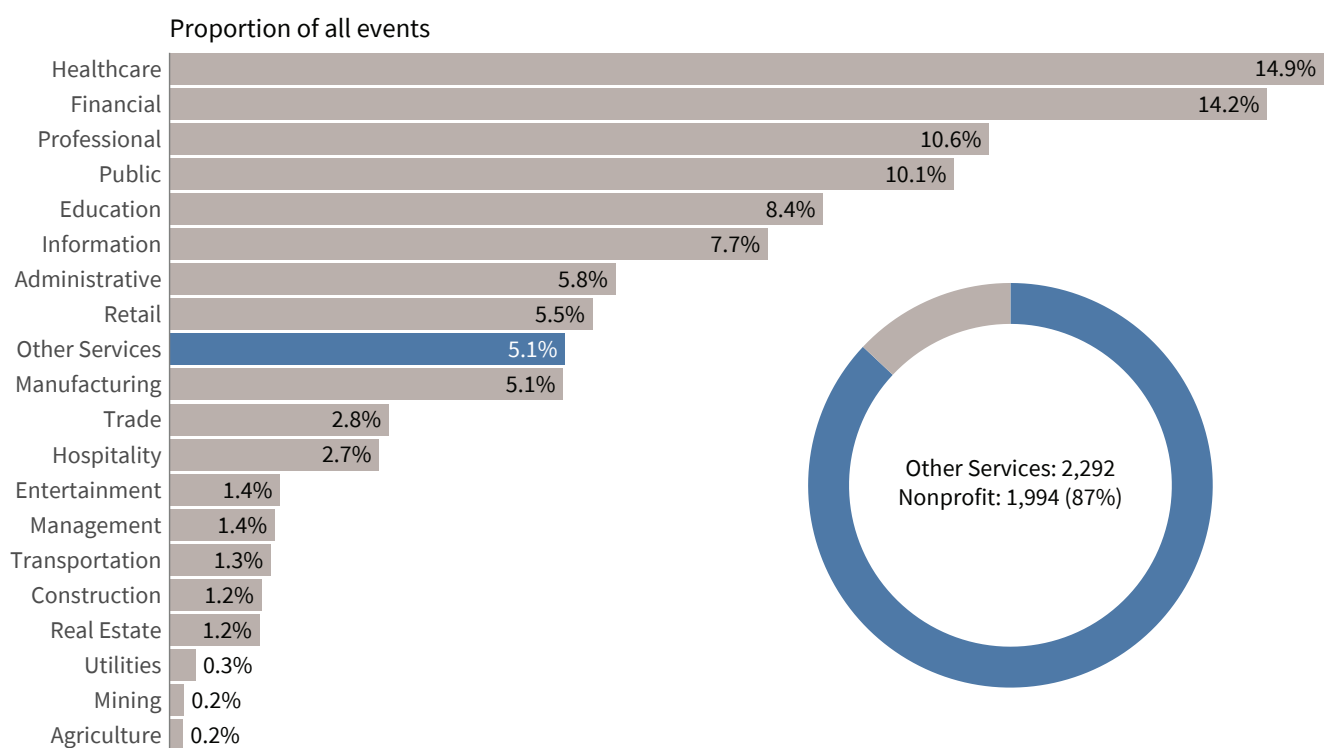
Historical Loss Events

Surprising no one, some industries experience more security incidents than others. But where does the nonprofit subsector—and the Other Services sector—stand relative to others? We briefly explore that here.

Comparing the total number of incidents attributed to each of the top-level NAICS sectors gives us Figure 1. We see more incidents in our dataset for healthcare and financial sectors than any other, while agriculture and mining record the fewest. The Other Services sector falls right smack in the middle of the pack.

Take care not to draw any hasty conclusions from Figure 1 about which industries are more/less secure or risky than others. Each sector differs in the number of active firms, regulatory obligations to report incidents, business models, technology portfolios, distribution of organization sizes, etc. All that to say, many things other than cyber risk posture contribute to the number of publicly-reported incidents shown in Figure 1.


Figure 1: Proportion of publicly-known incidents attributed to each sector



Now that we know where the Other Services sector sits, let's narrow in on the nonprofit subsector. The donut chart in Figure 1 illustrates that nonprofits account for the vast majority (87%) of incidents attributed to Other Services. The remainder of this Risk Retina narrows the scope to just those 1,994 incidents impacting nonprofit organizations.

Past events aren't necessarily a predictor of future trends, but ignoring them certainly doesn't help. In that vein, Figure 2 tallies incidents recorded for the nonprofit industry each month over the last decade. Keep in mind that incident reporting often lags months (even years) behind as events progress from discovery to disclosure, which explains the apparent falloff over 2021. Overall, the monthly incident count fluctuates but doesn't show a strong or sustained trend up or down.

Figure 2: Number of incidents reported each year in the nonprofit subsector




The giant spike in early 2020 stems from the Blackbaud ransomware breach that impacted over 800 organizations, including more than 100 identified nonprofits. It's a sobering reminder of the downstream impacts of large-scale [multi-party cyber incidents](#).

Although incident frequency at the industry level is interesting, enterprise cyber risk models generally focus on individual firms. We begin that transition by investigating how often particular organizations appear in our dataset.

Per Figure 3, 85% of nonprofit organizations with known public loss events over the last decade experienced only one event. The other 15% recorded multiple incidents, with just 1% suffering five or more. This is not the first chart you'll see depicting the long tail of cyber risk.

Studying incident frequency across 10 years provides useful perspective but most cyber risk managers seek to answer forward-looking questions like "what's the likelihood we'll suffer a loss event in the next 12 months?" The next section develops an answer via a well-fit statistical model.

Figure 3: Nonprofit event frequency over last 10 years



Modeling Loss Event Frequency

To derive event frequency, we could tally the number of incidents each year for each organization in the data. But that yields just 10 observed periods for each firm and results in very erratic measurements. Instead, we slice our dataset into rolling 12-month windows and count the events for each organization. This gives us up to 107 observations per firm,¹ a larger sample that we can employ to more confidently model annualized event frequency at an individual firm level.

Because the IRIS 20/20 demonstrated that incident frequency differs by organization size, we developed models for nonprofits ranging from \$10M to \$100M, \$100M to \$1B, and \$1B to \$10B in revenue.² For each revenue slice, and the nonprofit subsector overall, we treated the sliding sample described above as samples from an underlying probability distribution. We used maximum likelihood estimation to find the parameters that best fit the data for a number of candidate distributions.³

Using the Kolomogorov-Smirnov test and the Cramer von-Mises statistical tests, we then examined whether we could reject the null hypothesis that samples were drawn from the fitted distribution. We found that one particular distribution, the Poisson log-normal distribution, ‘passed’⁴ both statistical tests for all revenue slices and gave realistic estimates of multi-event years.

What’s the upshot of this statistical pedantry? We have a nice, closed-form representation of the probability of seeing a certain number of events in a one year time span for a non-profit organization. Adventurous souls who would like to implement their own version of this model will find the requisite parameters in Table 1. Plug and chug in the risk modeling tool of your choice.

Table 1: Event frequency model parameters for nonprofits

Frequency parameters - Poisson log-normal		
Upper and lower bounds		
Revenue category	Mean (μ)	Standard deviation (σ)
Upper Bound		
\$1B to \$10B	-2.291528	0.6778705
\$100M to \$1B	-2.193283	0.7072029
\$10M to \$100M	-2.237543	0.7326220
Lower Bound		
\$1B to \$10B	-3.238799	1.0055574
\$100M to \$1B	-4.847886	1.4259573
\$10M to \$100M	-6.247705	1.6921987

Lower and Upper Bound Estimates

Estimating the probability of incidents requires a known sample of firms on which to base calculations. Unfortunately, we don’t have a reliable count of relevant, active nonprofits around the world. But we have a couple proxies that can be used as a basis for reasonable lower and upper bound estimates. See Appendix F for more information about how we derived lower and upper bound models for loss event frequency.

But what does the output of this model look like and how does it fit the observed data? Figure 4 presents upper bound observed values (gray) and modeled estimates (blue) for annualized loss event frequency broken down by organization size (in annual revenue). For the most part, the observed and modeled values align, which points to a good model that fits the data. And the model has the added advantage of offering estimates for higher incident frequencies that were never observed in the historical data.

Risk Pro Tip: If you are interested in how this loss event frequency model fits the data—and we hope you are—but aren’t sure how to implement it, take heart! In a full Risk Retina we supply code to implement our suggested distributions in Python, R, and Excel!

¹ Subject to the firm being in operation over the entire 10 year period, a fact we account for in our data preparation.

² We decided not to develop a model for organizations under \$10M in revenue for various reasons related to data quality and model reliability. Anytime we show results for the Overall industry, it refers to organizations over \$10M in revenue. It’s standard practice for IRIS Risk Retinas for size categories to go beyond \$10B in annual revenue, but no registered nonprofits exceeded that threshold according to our dataset.

³ Distributions we tried: Poisson, Negative Binomial, Geometric, zero-inflated versions of those and the Poisson log-normal.

⁴ “Failed to reject,” but we don’t want silly things like weird statistical language to get in the way of the point we are trying to make.

Figure 4: Observed and modeled annualized loss event frequency for nonprofits (upper bound)

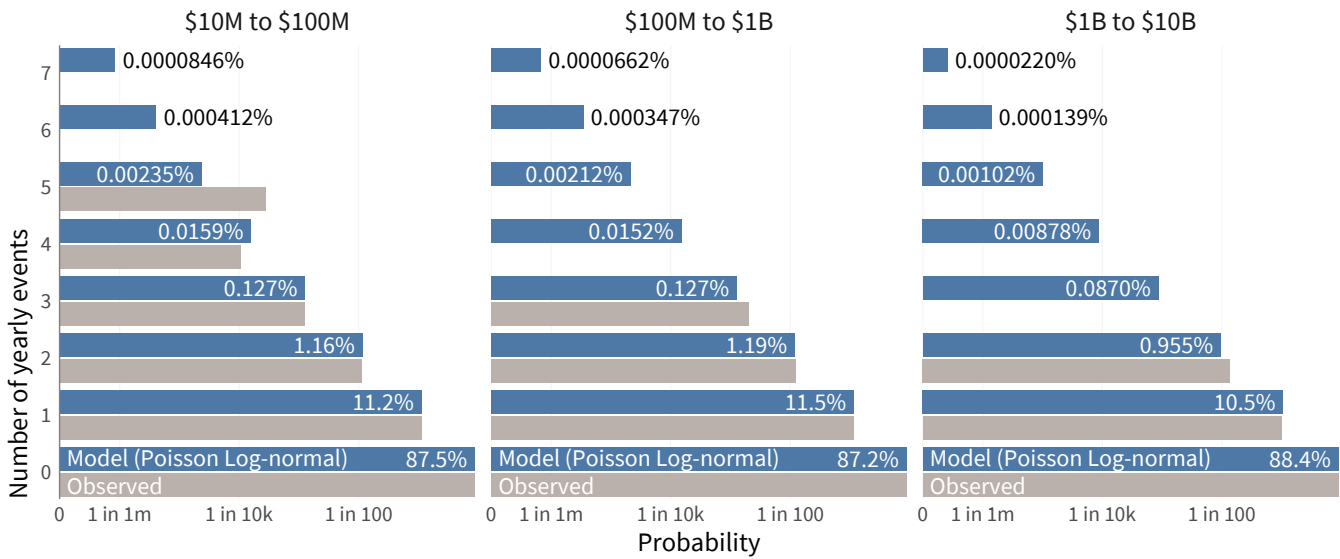


Table 2 is either a “check your work” for the DIY modelers out there or a handy pocket reference guide for nonprofit cyber risk managers who need a quick answer when asked about loss event frequency. Overall, the most likely outcome is that a nonprofit won’t experience any incidents. But risk management is all about understanding and managing unlikely outcomes. What’s the chance that your organization might suffer multiple incidents in a single year? We wouldn’t wish that on anyone, but Figure 4 and Table 2 should help you better plan for that unlikely scenario.

Table 2: Quick reference for loss event frequency estimates for nonprofits

Probability of a firm experiencing a given number of events			
Revenue category	One or more	Two or more	Three or more
\$1B to \$10B	11.53%	1.09%	0.10%
\$100M to \$1B	12.74%	1.33%	0.15%
\$10M to \$100M	12.57%	1.34%	0.14%

Where are the BetaPERT parameters?

Short answer: In [Appendix B](#).

Longer answer: While BetaPERT distributions are commonly used in cyber risk quantification, they’re not appropriate for estimating the expected number of events. If you’re modeling the probability of an event, the BetaPERT will work fine because it’s a continuous distribution. But as seen above, organizations can and do experience more than one incident in a year. Thus, a discrete distribution must be used to estimate frequency as we’ve done here.

Still, we realize that many solutions use BetaPERT distributions for event frequency. And so we’ve added model parameters for the probability of at least one event to [Appendix B](#).

Probable Loss Magnitude

Armed with estimates and distribution parameters for incident frequency, we now turn our attention to the financial losses incurred when an organization suffers cyber events. We'll start with observed losses from our historical dataset and then construct a model that best fits those values. We also include analysis of data losses and how to leverage this oft-misused measure to estimate overall financial impact.


Risk Pro Tip: 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, SEC filings, etc. Other forms of loss can be difficult to quantify and/or get absorbed internally rather than resulting in outward expenditures. We suspect the losses from highly public, major incidents are more complete than those from minor events due to increased scrutiny and public records. Thus, we hold that our recorded losses suitably reflect known financial losses from publicly visible cyber incidents.

Historical Event Losses

Financial losses tend to be less reported than other data points for cyber events, and this becomes a challenge when zooming in for specific sectors—especially a subsector like nonprofit organizations. The question becomes whether we have enough examples of historical losses to represent the range of probable future losses. Not an easy question to answer, but let's take stock of what we have.

We have a total of 64 nonprofit events with loss magnitude recorded in our dataset going back to 2000. Most of those are within the 10-year timeframe for this Retina, but we decided to extend that window to allow for additional (and still relevant) observations. Losses for those 64 events are shown in blue in Figure 5 amid losses recorded across all other industries in gray for comparison.

Figure 5: Distribution of reported losses for nonprofit incidents between 2000 and 2021



Overall, loss magnitude for the nonprofit industry follows the overall distribution in that values cluster in the low hundreds of thousands USD with a long tail. The main difference being the typical loss (geometric mean) is less (\$145K vs \$231K) and the 95th percentile of \$9M is about ¼ that of the overall distribution.

Those who appreciate the succinct conveyance of information should find Table 3 to their liking. It nets out key stats for loss magnitude in the nonprofit industry. Some may wonder about the likelihood of a nonprofit incurring more than our maximum observed loss (\$31M) from a cyber event. Unfortunately, historical data can't fill in gaps or project beyond its bounds to answer questions like that. But that's what models are for.

Table 3: Loss magnitude summary statistics for the nonprofit subsector


Loss summary								
	Minimum	1st quartile	Geometric mean	Median	3rd quartile	95th percentile	Maximum	Total events
Sector	\$462	\$22K	\$145K	\$236K	\$665K	\$9M	\$31M	64
Overall	\$27	\$33K	\$231K	\$206K	\$1M	\$34M	\$11B	2,747

Modeling Loss Magnitude

In this section we identify a distribution and associated parameters for cyber event losses in the nonprofit industry based on the historical data. Our first order of business is to determine whether we need different distributions based on organization size as we did for event frequency. The short answer is that we do not need separate loss models.

Because that's counterintuitive, we offer Figure 6 to show our work. It compares annual revenues and total losses for each nonprofit that had an incident with a recorded loss amount. As indicated by the nearly flat regression line with endpoints that lie within the gray confidence interval, there is no significant relationship between annual revenue and losses. That means a single distribution to describe losses for the whole nonprofit industry will suffice.

Figure 6: Correlation of annual revenue and loss magnitude among nonprofit incidents



While frequency involves a discrete number of events (with a probability of occurrence), losses can be fractions of dollars. This shift from a discrete to a continuous space broadens the options for applicable distributions. We tested several continuous distributions to determine which achieved the best fit for observed losses among incidents in the nonprofit industry. We landed on log-normal. That's convenient because it means that we can apply all the same techniques from a normal distribution by taking the log of every point in our loss data.

Figure 7: Log-normal distribution fit to historical cyber event losses in the nonprofit industry




Figure 7 fits the log-normal distribution to historical losses reported by nonprofits. The model fits pretty well until around the \$10M mark, where it appears to underpredict observed losses. But don't fret; that's a visual effect of the log scaling. Plus, we examine those tail losses more closely in a later section and offer another approach to anticipating such events.

Table 4 provides the relevant parameters associated with that distribution to plug into your cyber loss modeling tool of choice. As with frequency, we include BetaPERT parameters in the appendix for those who can't support a log-normal distribution.

Table 4: Loss magnitude model parameters for nonprofits


Log-normal parameters for loss	
Mean (μ)	Standard deviation (σ)
11.88129	2.553358

Estimating losses based on records compromised

We live in an age where information arguably has more value than any other non-human asset. It's intuitive, then, to assume that financial losses from an event would correlate with the number of data records compromised. Unfortunately, that intuition leads many down the path of using a fixed cost-per-record calculation. It just doesn't work that way. But record count can be used to estimate losses if done correctly.

Before we go there, however, let's set the stage by viewing the number of records compromised in nonprofit data breaches. Figure 8 shows the distribution along with highlighting key statistics. Interestingly, the typical exposure in a nonprofit breach is larger than the overall average and the 95th percentile is pretty much dead on at 80,000 records.

Figure 8: Number of data records affected by incidents in the nonprofit industry




Knowing that both financial losses and records affected are both log-normally distributed in cyber events enables us to better understand the relationship between these two variables. We're on the record⁵ showing that a flat cost per record is a terrible way to estimate loss magnitude. But applying just a little bit of math creates a sound model⁶ and produces the handy record-to-losses conversion guide in Table 5. Use it with our (and more importantly, the data's) blessing to replace the records × average cost calculation that you may have seen promulgated by others.

Table 5: Probable losses based on number of records affected in a nonprofit incident

Records	Probability of at least this much loss					
	\$10K	\$100K	\$1M	\$10M	\$100M	\$1B
10K	90.9%	66.4%	31.3%	8.1%	1.0%	0.1%
100K	93.0%	71.4%	36.5%	10.4%	1.5%	0.1%
1M	94.8%	76.1%	42.0%	13.2%	2.1%	0.2%
10M	96.1%	80.3%	47.6%	16.5%	3.0%	0.3%
100M	97.2%	84.0%	53.3%	20.3%	4.1%	0.4%
1B	98.0%	87.2%	58.9%	24.6%	5.5%	0.6%
10B	98.6%	90.0%	64.4%	29.3%	7.3%	0.9%

Risk Pro Tip: We realize converting data records affected to probable financial losses in the manner espoused above is more complicated than a flat cost-per-record approach. It may even garner pushback from executives that “just want a number.” But Table 5 is a far more honest representation of the large degree of uncertainty involved in the records-to-dollars conversion. Helping decision makers to understand that reality and incorporate it into their planning might be a short-term battle, but will be a long-term win. And it’s perfectly fine to start with simple statements like “A breach of 100M records has a median loss of about \$1M but there’s a small chance (4%) it could cost 100 times that amount.”



“

A SINGLE COST-PER-RECORD METRIC SIMPLY DOESN'T WORK AND SHOULDN'T BE USED. IT UNDERESTIMATES THE COST OF SMALLER EVENTS AND (VASTLY) OVERESTIMATES LARGE EVENTS.

“

THE TOTAL ERROR FROM THOSE ESTIMATES IS MORE THAN \$1.7 TRILLION DOLLARS. WE HOPE THIS EXPOSES THE FOLLY (AND PUTS THE LAST NAIL IN THE COFFIN) OF LOSS ESTIMATES BASED ON A SIMPLE AVERAGE COST PER RECORD DERIVED FROM A LIMITED RANGE OF DATA.

⁵ See the “Straight Talk on Cost-Per-Record Estimates” section in the IRIS 20/20.

⁶ A log-log linear regression.

Quantifying Risk Exposure

Event frequency and loss magnitude are good things to know in and of themselves, but many risk managers have questions like “what’s the likelihood we’ll lose \$X over the next year?” 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. Combined with an understanding of a firm’s risk appetite, LECs are great for exploring whether additional mitigation efforts are warranted.

In Figure 9, we combine frequency and loss parameters into a simulation to produce an overall LEC for a nonprofit organization. Trace any point on the curve to the x and y intercepts to determine exposure. For example, there’s less than a 1% chance that any given nonprofit organization will suffer cyber event losses exceeding \$10M in a year. We’ve added annotations to make those lookups easier and Table 6 should help too.

Figure 9: Loss Exceedance Curve for a typical nonprofit organization (Upper Bound)



The phrase “any given nonprofit” is important because this curve makes no distinction for the particulars of your organization. Your organization’s risk exposure might be higher or lower due to any number of factors, including external profile, location of operation, business model, IT environment, security posture, etc. We encourage you to consider the information presented in Figure 9 and Table 6 in light of such things.

We suspect the prevailing response after reviewing Table 6 is something akin to “where’s the risk!?” The relatively low probability of high annual losses certainly surprised us initially. But thinking it through given what we learned earlier about event frequency and loss magnitude, these findings become more intuitive. Most nonprofit organizations will not suffer a security incident and those that do probably won’t experience a worst-case scenario. But some will, and that’s why it’s a useful exercise to ask “what if?”

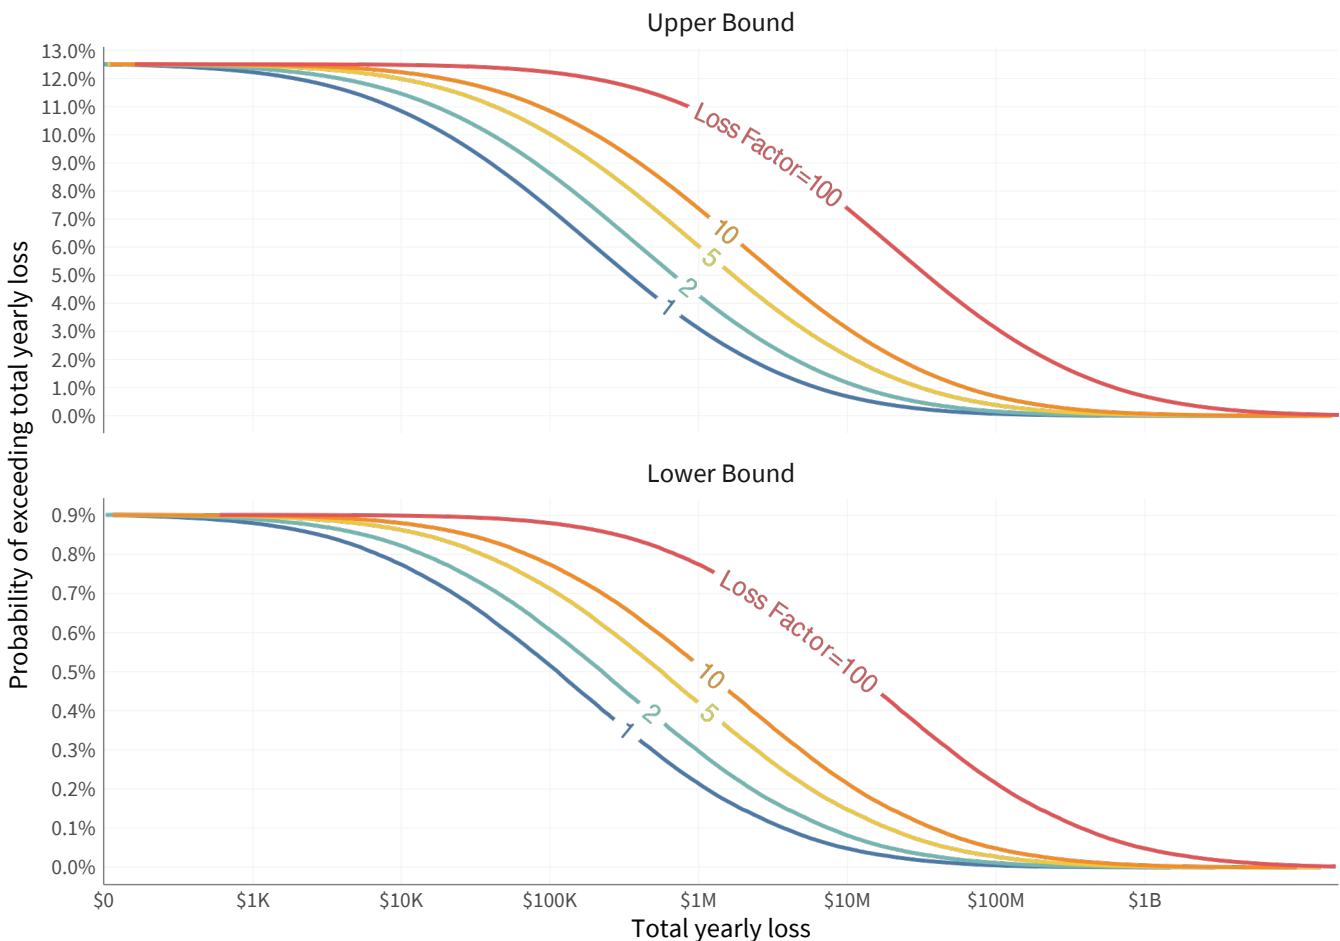
Table 6: Loss exceedance estimates for a typical nonprofit organization

	Chance of equal or greater loss			
	10.0%	5.0%	1.0%	0.1%
Upper Bound				
\$1B to \$10B	\$9,674	\$274,024	\$5,481,871	\$68,171,067
\$100M to \$1B	\$23,560	\$366,364	\$6,373,507	\$76,106,446
\$10M to \$100M	\$19,707	\$344,512	\$6,147,390	\$73,857,063
Lower Bound				
\$1B to \$10B	\$0	\$12,592	\$1,954,016	\$34,692,131
\$100M to \$1B	\$0	\$0	\$174,717	\$10,785,102
\$10M to \$100M	\$0	\$0	\$0	\$2,934,209

Figure 10 allows for the possibility that reported financial losses in our dataset may not reflect the totality of impact that cybersecurity incidents can inflict on organizations. The standard LEC for the nonprofit industry has been modified to account for losses that are substantially higher than what we observe in the historical data. We also show both the upper and lower bound curves for those who really go nuts with analytical possibilities.

To us, the main takeaway from Figure 10 is that even cranking up loss magnitude by a factor of 100X still doesn't produce an absurdly high annualized risk exposure (less than 8% chance of exceeding \$10M). Even so, it's those extreme exposures that tend to cause the most concern. That's why we give them special attention in the next section.

Figure 10: Loss Exceedance Curves with amplified loss factors for a nonprofit organization



Tail Risk Analysis

As indicated by the previous section, cyber risk has a long tail of rare but highly impactful events. Most executives and risk managers worry far more about that tail than the bulk of more predictable loss scenarios. We studied 100 of the most impactful cyber incidents in the IRIS Xtreme and sought to understand the actors, techniques, and contributing factors behind them.

We won't reproduce those details here. But we do want to understand the nonprofit industry's propensity for experiencing extreme cyber loss events compared to what is typically seen in other sectors. As it turns out, there's only one example in our historical data of a nonprofit cyber event meeting the criteria for an extreme event per the IRIS 20/20 Xtreme. So not much we can infer from the historical data alone.


To dig a little deeper into what extreme events for nonprofit organizations might look like, we turn to a technique employed frequently by actuaries called Tail Value at Risk (TVaR).⁷ TVaR is a simple enough concept: Given the top X% of losses an organization is likely to experience in a year, what is their average value? TVaR is useful because it helps illustrate how scary the heavy tail we see in losses can be.

Table 7: Tail Value at Risk (TVaR) analysis for nonprofit loss events

	90%	95%	99%
Upper Bound			
\$1B to \$10B	\$176M	\$351M	\$2B
\$100M to \$1B	\$199M	\$397M	\$2B
\$10M to \$100M	\$193M	\$385M	\$2B
Lower Bound			
\$1B to \$10B	\$80M	\$159M	\$775M
\$100M to \$1B	\$25M	\$50M	\$251M
\$10M to \$100M	\$9M	\$17M	\$87M

TVaR is a simple enough concept: Given the top X% of losses an organization is likely to experience in a year, what is their average value?

Table 7 highlights upper and lower bound TVaR estimates for nonprofits. While the 95th percentile for the largest loss events might be a mere \$366k (see Table 6), the 95% TVaR is upwards of \$397M. This means that the minimum loss magnitude (marked by the 95th percentile) of a once in 20 year kind of event might not seem so bad, but averaging the full length of the long tail becomes very pricey indeed.



Sponsored by


VISIBLERISK

IRIS 20/20

Information Risk Insights Study

Xtreme

Analyzing the 100 largest cyber loss events of the last five years



IRIS
20/20

Extreme loss events like those discussed in this section are scary. But knowledge overcomes fear. That's why the IRIS 20/20 Xtreme focuses on the 100 largest cyber incidents of the last five years, totaling \$18 billion in reported losses and 10 billion compromised records. We once again started with Advisen's Cyber Loss Data and then collected hundreds of additional data points on each of these extreme cyber loss events. Our goal was to breakdown the costs, categorize incident types, identify the actors behind these events and the actions they employed, and better understand how these events impacted the organizations involved. Our primary goal remains the same as the IRIS 20/20—to clear the fog of fear, uncertainty and doubt (FUD) surrounding cyber risk and help managers see their way to better data-driven decisions.

[Download the full report](#)

⁷ Sometimes called Conditional Tail Expectation (CTE).

Common Incident Patterns

So far we've treated all cyber loss events the same. That's fine at the macro level when our main goal is to quantify overall risk to the organization. But when the goal turns to mitigating that risk, it helps to know something about the types of incidents that are most common and costly.

To help with that, we categorize incidents according to common patterns of threat actors, techniques, vectors, and technical impacts as defined in the list below. These patterns are intended to represent the high-level scenarios we often see on risk registers for assessment and reporting purposes.

IRIS Risk Retina Incident Patterns

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

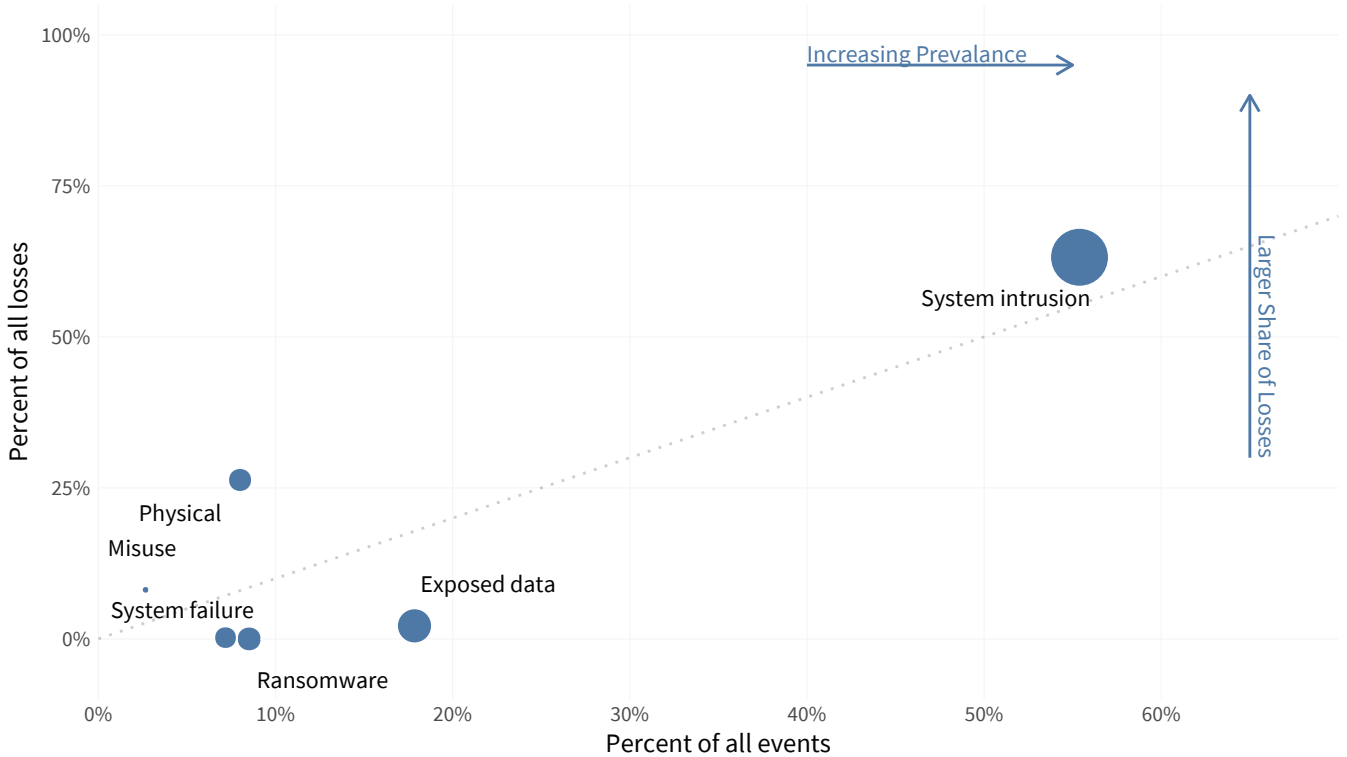
1. DDoS: Any attack intended to render online systems, applications, or networks unavailable, typically by consuming processing or bandwidth resources.
2. Exposed data: Data stores that are inadvertently left accessible to unauthorized parties, typically through misconfigurations on the part of the data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

Figure 11 compares the relative frequency (% of events) and impact (% of total losses) for each incident pattern. The proportion of compromised data records is also reflected by the size of the dot. The intent is to enable managers to focus on the most risky types of loss events. Anything above the dotted line has a higher percentage of financial losses relative to its frequency of occurrence. Patterns below the line are relatively less costly.

Positioned in the extreme upper right, system intrusions are far and away the most prominent incident pattern afflicting nonprofit organizations. They accounted for over half of all events and two-thirds of total losses across the industry over the last decade. While not part of this analysis, our in-depth examination of extreme and massive multi-party events points to the exploitation of [valid user accounts](#) as the primary initial access technique for system intrusions.

Exposed data stores appear as a distant second on the frequency and dataloss scales. Whether your nonprofit collects donor information, develops cutting-edge IP, or handles other forms of data, take care how it’s protected. Cloud storage is often cheap and convenient—a very attractive combo for nonprofits short of IT staff and cash—but it’s also prone to misconfiguration that leaves your data open to any stranger on the internet.

Figure 11: Relative frequency and losses associated of common incident patterns



Depending on the question being asked, it may be useful to have exact values for the proportion of events, losses, and records shown in Figure 11. Table 8 captures all this, along with a ranking for each measure and comparison against the overall population. Here we see both similarities and differences highlighted for the nonprofit industry that we hope supports more informed risk management decisions.

Table 8: Ranking of common incident patterns with relative frequency, financial loss, and data loss statistics

	Frequency			Financial impact			Records affected		
	Sector	Sector rank	Overall rank	Sector	Sector rank	Overall rank	Sector	Sector rank	Overall rank
System intrusion	55.38%	1	1	66.08%	1	1	9.31%	2	1
Exposed data	17.85%	2	2	2.45%	3	6	89.20%	1	2
Ransomware	8.51%	3	6	0.00%	6	2	1.05%	3	6
Physical	8.00%	4	3	30.56%	2	5	0.22%	4	4
System failure	7.18%	5	5	0.36%	5	3	0.00%	6	5
Misuse	2.67%	6	4	0.54%	4	7	0.21%	5	3
DDoS	-	-	8	-	-	8	-	-	8

We could dive into each of these incident patterns individually but that goes beyond the scope of this baseline report for the nonprofit industry. But such deeper dives into the pattern pool are available as an add-on to your organization’s IRIS Risk Retina package.

Viewing Other Dimensions

Our IRIS Risk Retina service begins with a baseline analysis of risk parameters for a particular sector—everything you’ve seen thus far for nonprofits. Beyond that, you can choose selections from our library of reports that target risk dimensions relevant to your organization.

A thorough analysis of exteme events (including TVaR calculations previewed in an earlier section) plus details on who’s behind them, what techniques were used, and how firms responded (similar to the [IRIS 20/20 Xtreme](#)) is one such dimension. We’d like to use this final section to preview two other Risk Retina dimensions we’re developing around the Vocabulary for Event Recording and Incident Sharing (VERIS) and multi-party events. If interested, ask your friendly Cyentia rep about adding them to your Risk Retina package.

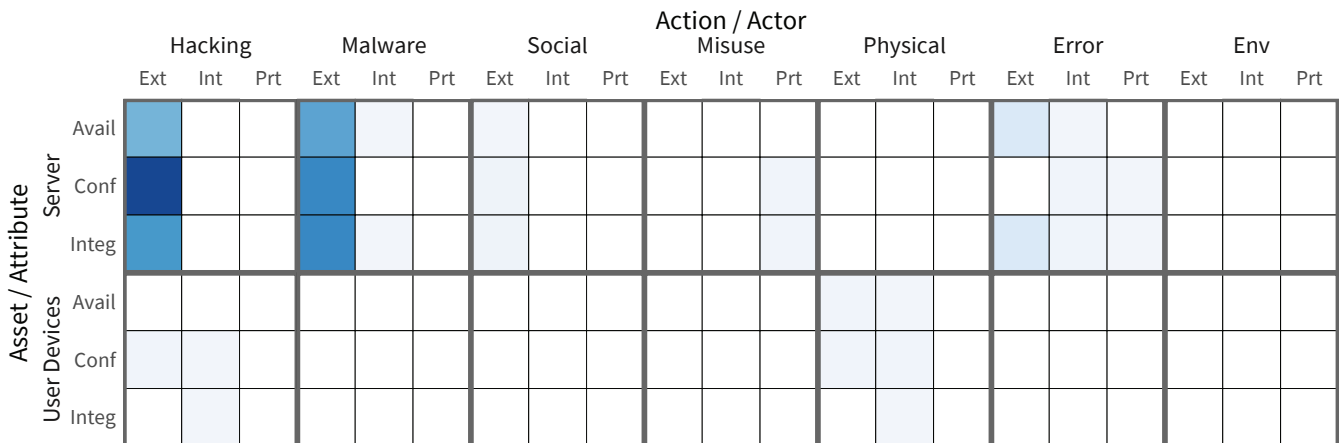
VERIS Threat Events

VERIS is a framework for describing security incidents in a structured and repeatable manner. To accomplish this, VERIS employs the A4 Threat Model developed originally for Verizon’s Data Breach Investigations Report (DBIR). Describing an incident essentially means identifying all the actors, actions, assets, and attributes involved (the 4 A’s).

- **Actors:** Whose actions affected the asset?
- **Actions:** What actions affected the asset?
- **Assets:** Which assets were affected?
- **Attributes:** How the asset was affected?

The four A’s are intended to represent the minimum information necessary to adequately describe any threat or incident scenario (a chain of threat events). Figure 12 presents a consolidated view of threat events recorded across the historical incidents in the nonprofit industry through the VERIS A4 Grid. The Grid is designed to give a big picture view of all possible threat events that form an incident scenario. Each intersection is a unique combination of the 4 A’s. For example, the top left intersection describes a threat event involving an external actor using a hacking technique to compromise the availability of a server (probably an application-layer DDoS attack) somewhere in the incident scenario.

Figure 12: VERIS A4 Grid depicting the relative frequency of high-level threat events contributing to nonprofit incidents



We’ve truncated the grid to show the two most commonly affected asset categories (servers and user devices) for this Retina dimensional preview. It normally shows 315 intersections. It’s easy to see that Hacking and Malware dominate, being involved in over 80% of events! If you like this concept, there’s an optional Risk Retina dimension that categorizes all incidents according to the 4 A’s, constructs a complete version of the Grid you see here, and more!

Multi-Party Incidents

Modern organizations operate in an increasingly interconnected environment, which enables cyber events and their effects to propagate via business-to-business relationships. Such events spread outward from the initial victim organization to impact multiple organizations, which is why we call them ripple events. From a risk perspective, ripple events are difficult to manage because your organization can be impacted through the actions (or inactions) of another.

How common are multi-party incidents in the nonprofit industry? Table 9 contains the answer and the good news is that it's not that often. The industry ranks in the bottom half (#11 out of 18) of the top-level sectors designated by NAICS. Nonprofits are slightly more likely to be on the receiving end of ripple events, however (#8 of 18). As such, nonprofits should be mindful of the security posture of key 3rd party partners and service providers.

Table 9: Relative prevalence of multi-party incidents in the nonprofit industry

Relative prevalence of ripple events	
	Sector rank
Ripple related	11 out of 18
Ripple generators	11 out of 18
Ripple receivers	8 out of 18

Some ripple events grow so large that they're better described as tsunamis. These massive multi-party incidents were the focus of our IRIS Tsunami report, which examined the 50 largest ripple events of the last decade. Understanding what differentiates these cyber tsunamis from their smaller scale ripple events can be very useful for managing risk. Fortunately, our records show the seas are clear of historical tsunami events in the nonprofit industry!

riskrecon

mastercard

interos

Cyber

GRX

Information
Risk
Insights
Study

IRIS TSUNAMI

Following the wake of damage from major multi-party cyber incidents

Cy

centia

In almost every way imaginable, we live in a hyperconnected world. This connectivity has brought many benefits to modern business models, but it has also introduced myriad challenges and risks. If you take the time to deconstruct even the simplest of business transactions, you'll find in the mix a surprising number of parties from technical components supporting the transaction to the completed delivery of products to the customer. But what happens to all these parties when something goes wrong?

That is ultimately the question the new IRIS Tsunami seeks to explore. We identified 50 of the largest multi-party cyber incidents over the past several years in an effort to understand their causes and consequences from beginning to end. If you are familiar with our other research in the Information Risk Insights Study (IRIS) series, Tsunami draws from the same rigorous methodology. We started with a huge dataset of cyber loss events, identified those that involved multiple organizations, and then researched each event to understand who was behind it, what happened, how the after effects propagated through the supply chain, and the financial losses for all parties involved.

[Download the full report](#)

Appendix β: BetaPERT

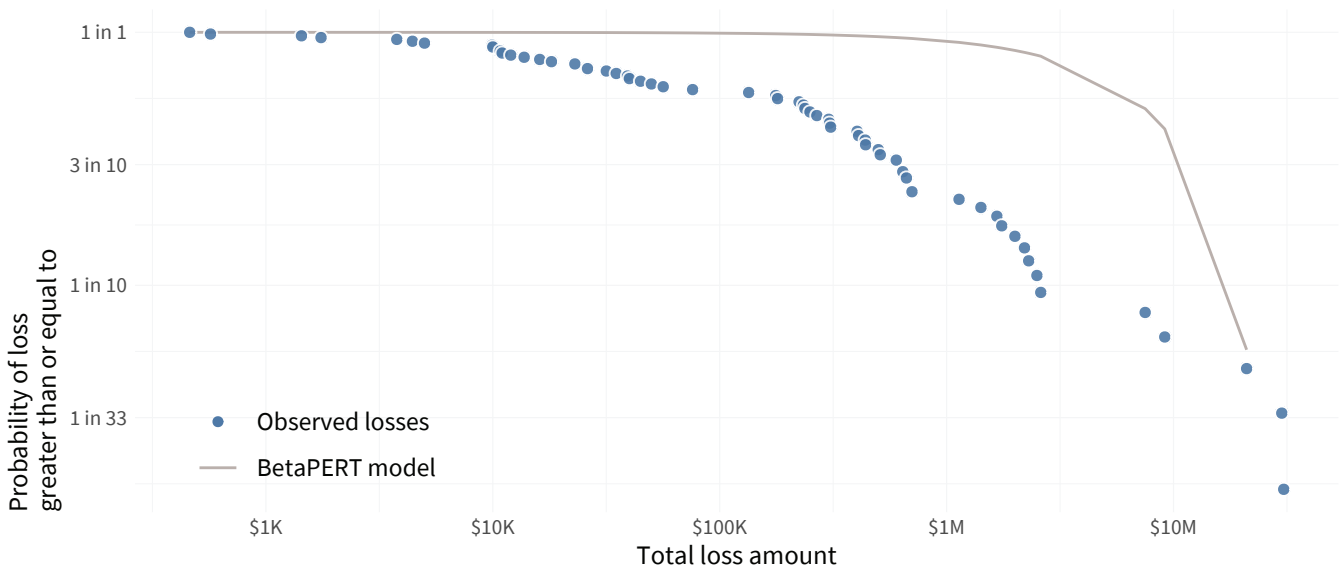
Below you'll find risk parameters to feed into BetaPERT distributions. For estimating the probability of experiencing one or more security incidents, the beta model fits the observed data quite well. Just keep in mind this is only appropriate for probabilities (values between 0 and 1), but not for modeling the frequency or number of events.

Table 10: BetaPERT parameters for the probability of a nonprofit organization experiencing at least one cyber event

Probability of at least one event				
BetaPERT parameters for upper and lower bounds				
	Parameters			
	Minum	Mode	Maximum	Shape
Upper Bound				
\$1B to \$10B	0	0.098535464	1	17.72888
\$100M to \$1B	0	0.110874104	1	14.48968
\$10M to \$100M	0	0.102079592	1	13.95533
Lower Bound				
\$1B to \$10B	0	0.007486058	1	15.94390
\$100M to \$1B	0	0.000000000	1	26.12403
\$10M to \$100M	0	0.000000000	1	54.16875

We foreshadowed that BetaPERT distribution didn't fit the observed data for loss magnitude very well and Figure 13 makes that plain as day. The BetaPERT model dramatically overestimates the bulk of recorded losses between the extremes of the range (note the log scale here). That might not seem so bad if you're extremely risk averse, but the model completely ignores any potential loss above the max. There are far better ways to account for risk aversion than using a poor model.

Figure 13: BetaPERT distribution fit to historical cyber event losses in the nonprofit industry



Given the poor fit demonstrated in Figure 12, we debated whether to include beta parameters for loss magnitude. But we realize BetaPERT distributions might be the only option for some. Thus, we think it’s better to have this information—along with Figure 13—to support your analysis and decisions rather than making blind assumptions.

Table 11: BetaPERT loss magnitude model parameters for nonprofits

BetaPERT parameters for loss			
Minimum	Maximum	Mode	Shape
\$461	\$30,600,001	\$235,744	1.518063

Risk Pro Tip: Apologies for the repetition, but we’d like to stress once again that the BetaPERT distribution does not fit the data and that we do not recommend using it for modeling loss magnitude. We include it only to make that point and to encourage both cyber risk managers and product managers to use and support other distributions.

Appendix Γ: Lower and Upper Bounds

Why Gamma? Well, it’s the Greek letter that comes after Beta. It just seemed wrong to go with “C” after that precedent.

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 nonprofits 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: Uses all registered nonprofits in the United States according to Dun & Bradstreet (because we don’t have numbers for the whole world). That assumes incident frequency among U.S. firms is similar to those 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. nonprofit. We say this is a lower bound because it assumes that all registered nonprofits 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: Uses all nonprofits recorded in our dataset, which means they’ve experienced a known incident at some point in the past. While that’s clearly not the case for all registered nonprofits, this upper bound approach is based on the premise that not all nonprofits 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 nonprofits prone to incidents have had one already, 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 lower and upper bounds, so we’ve opted to share both to support your assessment. We keep things simple by presenting upper bound charts and including lower bound values in the tables. In general, the upper bound offers a more risk-averse view (higher values). Choose one or fuse both to suit your organization’s risk posture and tolerance.



IRIS Risk Retina

by  Cyentia

Ready for your own Risk Retina?

We hope you agree this analysis offers a strong foundation for cyber risk quantification. With Risk Retina, you benefit from the world's premier database of historical cyber loss events from Advisen. Cyber insurers use this data to build their risk models, so why not tap into that insight for your own needs? What's even better is that you don't have to wrangle with the data to obtain those insights – we've done that for you! You get the best data and the best analysis in one package that's tailored to fit your needs.

Another way you benefit from IRIS Risk Retina is immediately. There's no software to install, no calibration training to attend, no consulting engagements to contract. You get the data you need right away for the cyber risk assessments you're doing today.

[Learn more](#) or [reach out](#) to get yours!