→ cyberresilience.com

How To Do Rapid Risk Interviews

Since the publication of How To Measure Anything In Cybersecurity Risk, author and Resilience Chief Risk Officer Richard Seiersen has had the opportunity to consult with dozens of CISOs and their security teams. One thing he hears frequently is, **"how do I get started with the methods found in your book?"** This document written by Seiersen addresses that concern.

Note from the author: Consider this document the first rung on a cyber risk quantification ladder. If the rapid risk interview is the first rung then the second rung is the 1-or-1 substitution introduced in Chapter 3, and the spreadsheet tool found here. The third rung would be a full Cyber Risk Quantification (CRQ) engagement.

resilience

First Step On The Cyber Risk Quantification Journey

RAPID INTRO

The table on the following page is a “rapid risk assessment summary.” You can think of it as a type of risk register. Its goal is to get a quick feel of both expected and unexpected financial losses. In turn, those values can be used to quickly inform spending on security programs and insurance. Perhaps most importantly, the rapid risk assessment can pinpoint areas of high risk requiring immediate focus.

Example Rapid Risk Assessment Summary

Peril	Loss Ranges			Impact		Assessed Orgs						
	10% Low	Median	90% High	Mean Event Loss	Yearly Expected Losses	CIO	CTO	CFO	CMO	CRO	CISO	OTHER
Ransomware Breach	$2.9M	$20.3M	$29M	$17.7M	$443K	✓					✓	
Ransomware Disruption	$1.7M	$4.9M	$13.4M	$6.5M	$165K	✓					✓	
Ransomware Extortion	$300K	$1.5M	$5M	$2.2M	$55K	✓					✓	
BEC Fraud	$200K	$2M	$5M	$2.4M	$60K	✓		✓			✓	
Cloud Data Breach	$1.74M	$11.6M	$17.4M	$10M	$250K	✓	✓				✓	
Cloud Disruption	$370K	$1.34M	$2.7M	$1.5M	$38K	✓	✓				✓	
SaaS Data Breach	$290K	$1.2M	$5.8M	$2.3M	$57K	✓	✓	✓	✓	✓	✓	
SaaS Disruption	$1M	$4M	$6.5M	$3.9M	$100K	✓	✓	✓	✓	✓	✓	
			Mean Event Total:	$46.5M								
			Total Expected Yearly Losses:		$1.2M							

$1.2M Is Roughly What You Might Expect To Pay Per Year To Cover $46.5M In Aggregate Losses Over 40 Years.

The key to understanding your rapid risk assessment summary is found in the column headers. Each main header is detailed below, with the remainder of the document explaining how the summary is built.

How To Do Rapid Interviews

Peril	Insurance uses the term “peril.” A peril consists of a threat and loss. Threats include things like ransomware, business email compromise, cloud service compromise (simplified as cloud), saas compromise (simplified as saas), etc. Losses include things like data breach, business disruption, extortion, and fraud.
Loss Ranges	Initial values are in the form of records lost, time lost or direct dollars lost. Some of this comes from interviews – but not all. Simple arithmetic is then used to create the ranges you see below.
Impact	If a material event happens, then you could see something approximating the mean loss. The yearly expected loss is simply the cost of doing business as it relates to security and these particular perils.
Assessed Orgs	A list of business unit senior leaders is used for simple tracking purposes. Not all risks require interviews with every org. Making transparent who was, and was not considered, helps ground rationale for values.
Loss Severity	The loss ranges are multiplied by 30%, 40%, and 30% respectively. The products of the multiplication are then summed to get the Mean value. The 30-40-30 weights, when applied to a three point range, provide a standardized way to estimate the mean value of an uncertainty for an 80th percentile prediction interval.
Yearly Expected Losses	We are using a flat event likelihood of 2.5% throughout. We multiply 2.5% times the mean event to get the “expected value.”
Mean Event Total	This is the sum of all the mean events. If you can, you should transfer as much of this away from the business as reasonably possible.
Total Expected Yearly Loss	This is the sum of all the yearly expected losses.

The analysis will show you where your biggest potential losses are. In this case, ransomware breach has the largest losses. Furthermore, the analysis will guide you on transferring risk away from the company. As a rule of thumb, the **total expected yearly losses** represent the premium you could pay for a limit that covers the **mean event totals.**

Keep in mind that the total expected yearly loss is a type of average loss. You are going to want to invest in controls that reduce your tail risk, which can be much larger than the total expected yearly loss. That level of analysis requires a complete cyber risk quantification (CRQ) analysis. Until then, the rapid risk assessment and associated summary will get you started.

Summary

Imagine it's your first day on the job as the new CISO of a mid-sized organization with revenue around $250M. They have grown large enough to warrant establishing a security organization. You're the first hire. You have no staff. You have no budget. To build and fund a plan, you need a rapid method of assessing your enterprise's cyber risk.

This is what the rapid risk interview was made for. Its goal is to quickly assess high-impact and "readily knowable" risks. These risks are juxtaposed to "knowable-unknowns" that require a thorough, and oftentimes technical, risk assessment.

The rapid risk interview should be the first step for the new CISO, with deeper technical and quantitative assessments to follow.

APPROACH

Each rapid risk interview takes maximally a couple of hours to complete. A complete summary table, as seen above, can be assembled in one week. The goal is a basic forecast in terms of unexpected and expected losses. Basic forecasts are juxtaposed to more rigorous cyber risk quantification (CRQ) approaches.

During a rapid interview it was discovered that a company hosted many terabytes of personally identifiable information (pii) in the cloud. (Arguably one of the largest regulated data stores in the cloud at that time.) Key access controls were missing and security assessments were all but non-existent. Adding fuel to the fire, cyber insurance was significantly underfunded relative to the risk. This was discovered in a couple of days.

The results? A doubling of insurance limits. The rapid risk interview revealed material risks within days. Risk was reduced quickly in terms of transfer to insurance, while allowing for broader technical assessments and follow-up investments in mitigation.

DEFINE YOUR INTERVIEW SCOPE

The first step is defining which perils raise the greatest concern and who you are going to consult for details about them. You can see a starting list of threats and their losses (perils) in the summary to determine how perils are aligned to relevant organizations. The green checkboxes guide you in your interview targets.

The next step is to schedule rapid risk interviews with each major organizational unit owner and/or key members of their team.

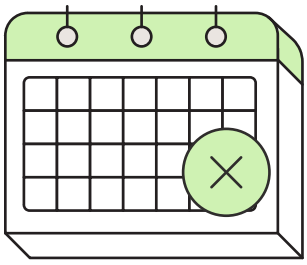

ASSESSING RISK

The bulk of the interview will focus on assessing how much of your organization's value is exposed to threats. Exposed value can be compromised, which can lead to impactful losses. Losses of value fall into at least three macro categories:

- **Data Loss:** Regulated data, ip, and other non-public information that is compromised
- **Business Disruption:** Service disruption that has a material impact on revenue
- **Direct Dollars Lost:** Extortion, fraud, brand (sales, stock, etc) and errors leading to direct losses of cash

Understanding and assessing loss ranges can be a challenging concept. Resilience experts spend a considerable amount of time on this subject in our training events. To ease into the topic, we have created a quick conceptual tutorial below. If this type of analysis piques your interest, see Resilience's Cybersecurity Risk Management Leader, Robert Brown's great book: Business Case Analysis with R.

Before we discuss how to assess a range, remember that ranges can be thought of as distributions of uncertainty about something we want to measure, like costs, but are unsure what the actual cost will be.

There are many ways distributions can represent uncertainty, and not all distributions have to look like a centered bell curve like the central orange distribution in the chart to the right. They can be skewed to the right (e.g., the dark green distribution on the left) or the left (e.g., the lime green distribution on the right). The amount they skew reflects how much information we have about the range, above and below the peak of the distribution.

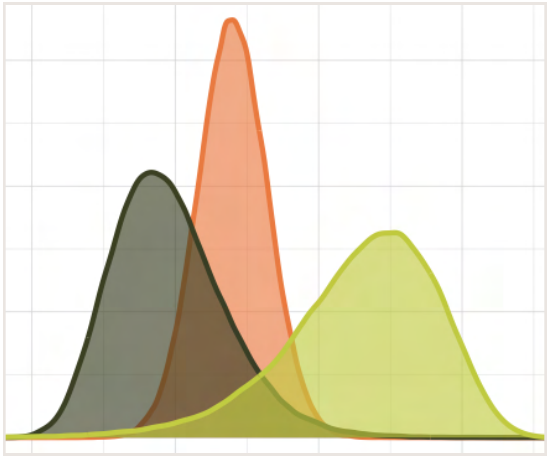

By comparison, the orange distribution is narrower and symmetric around its peak. This means that we are more sure about the actual value of the orange distribution than the value of the dark green or lime green.

The dark green distribution is shifted to the left of orange, but its tail is skewed to the right. For the measurement that the dark green distribution represents, we are saying that it demonstrates less confidence about its range than the orange (because it's wider), and less confidence about its range to the right of its own peak than to the left of its peak (because its tail is longer to the right).

Similarly, the lime green distribution is shifted to the right of orange, but it is skewed to the left. For the measurement that the lime green distribution represents, it demonstrates less confidence about its range than the orange or dark green (because it's wider than both), and less confidence about its range to the left of its peak than to the right of its peak.

In all three examples above, the actual value will more likely be located near the peak than out in the tails. Although the tails reflect less certainty than the peaks, we haven't totally disregarded them, and should consider their possibility to avoid surprise. But some tails are more likely than others; therefore, it is a good practice to get used to thinking that centered distributions (like the orange one) are the exception, rather than the rule. Simply put, we usually possess more information about one side of a range's central value than the other side.

Let's walk through a simple example to see how to put this into practice.

Assessing Value Ranges Using Elephant Weights

You're likely not a veterinarian, zoologist, or a big game hunter. Yet I bet you know something about the weight of an elephant – without having to use Google. You know that an elephant weighs more than 10 lbs. You may also assume it weighs less than 100,000 lbs. In fact, you could likely come up with a much more reasonable range.

We will call our reasonable range your **80% prediction interval**. Start by forecasting the highest number on the range. It should be a "surprisingly high" value – but not shocking. For example, you might be surprised if an elephant weighed more than 20,000 lbs. (I would be shocked if it was more than 50,000 lbs). The 20,000 lbs value is your **90% value** in the

prediction interval. You believe there is a 90% chance the true value is less than 20,000 lbs. We show this 90% value with the black dotted line on the right side of the distribution.

Next we select the lower bound value. You may be surprised if the value was less than 5,000 lbs and shocked if it was below 3,000 lbs. You think there is a 10% chance the weight is less than 5,000 lbs. We show this 10% value with the black dotted line on the left side of the distribution. (See how the 10% to 90% range equals an 80% interval)

You may feel the true weight tends more toward one end of the range than the other. Let's say you feel there is an equal chance the true weight is above 15,000 lbs as it is below. Formally, we would say you believe the **median** of the 80% prediction interval is 15,000 lbs. We place this median value where the black dotted line is.

Is the median off to the left, right, or dead center? In this case, it's off to the right. We would say that our assessment of the weight of an elephant is skewed. This is why detecting skewness matters. This same process is applied in the case study below.

USE CASE: RANSOMWARE BASED DATA LOSS

Most, but not all, data has monetary value. That's why the goal of many threats is exfiltrating data – particularly in ransomware attacks. According to the 2021 Verizon DBIR report, 10% of all data breaches are Ransomware based. Stealing data is becoming part of an emerging "double extortion" ransomware strategy. From 2020 to 2021 ransomware-based double extortion increased over 900%. While this example focuses on ransomware, it can be applied to countless threat scenarios.

ASSESSING RECORDS AT RISK

The first question to ask yourself is, "Does my company process, store or otherwise handle data that will have financial impact if compromised?" If the answer is yes, then your next step is to assess a rough range of data and records that could potentially be exposed.

We build the first example one row at a time. Portions of the table will inform the first row in the summary above.

Example Rapid Risk Assessment Summary

80% Prediction Intervals For Ransomware Data Breach			
Range	Low (10%)	Median (10%)	High (90%)
Records	5M	35M	50M
Loss	$2.9M	$20.3M	$29M
Probability Weighted Loss	0.3 * $2.9M = **$870K**	0.4 * $20.3M = **$8.7M**	0.3 * $29M = **$8.7M**
Loss Severity	$870K + $8.1M + $8.7M = **$17.7M**		
Expected Loss	0.025 * $17.7M = **$443,000**		

RECORD RANGES

We asked our example client how many records they store on systems that they directly manage. They were uncertain, but would be surprised if the record count was over 50M. On the lower end, they assumed as low as 5M records. They also assumed the count could be just as likely to be above or below 35M records.

This line of thinking follows the same logic in the elephant weight tutorial above. With this data, progressively add rows to the table below. Each row builds out the losses. The first row expresses the basic unit of loss like records. Later we consider time as well as direct dollars lost.

80% Prediction Intervals For Ransomware Data Breach			
Range	Low (10%)	Median (50/50)	High (90%)
Records	5M	35M	50M

COST PER RECORD

Next, convert the record counts to dollar values. To do that for this example, we are using the average value provided by the Verizon DBIR team of $0.58 per record.* Multiply the record counts by that value to get losses.

*We are using $0.58/record as a very simple and conservative example for our discussion; however, we recommend that you use a log-log regression analysis to capture a more robust description of how the cost per records lost varies with the number records lost.

NOTE

Verizon's basic analysis and rationale lines up with our own models. But keep in mind, this is just a shortcut to keep things agile. Full CRQ analysis goes well beyond static multipliers.

Steer clear of the forecasts that put the per record averages up in the multiple hundred dollars per record. Such analysis isn't accounting for how the cost per records scale downwards based on the amount of records at risk. You can see that in the graph at the end of this document extracted from Verizon. Note, the loss values show up in the first row in the summary above.

80% Prediction Intervals For Ransomware Data Breach			
Range	Low (10%)	Median (50/50)	High (90%)
Records	5M	35M	50M
Loss	$2.9M	$20.3M	$29M

A SIMPLE PROBABILITY WEIGHTED LOSS

Our next step uses a low-math method for approximating a risk-adjustment of the range of losses. We do that by multiplying the loss values by a set of probabilities – which "probability weight" our losses. There are two types of probability weightings we use here. The selection of each weighting is based on whether the records are **skewed** or **centered**. We will stick with the skewed solution throughout this exercise:

- **Skewed Data:** If the middle value in your data leans toward the 90% or the 10% range then your data is skewed. With skewed data you multiply losses by **0.3, 0.4, 0.3**
- **Centered Data:** If the values are not skewed we multiply by **0.25, 0.5, 0.25** respectively

NOTE

I must emphasize that this approach should be considered a short-hand method for approximating the distribution of values. It should be followed with rigorous cyber risk quantification.

80% Prediction Intervals For Ransomware Data Breach			
Range	Low (10%)	Median (50/50)	High (90%)
Records	5M	35M	50M
Loss	$2.9M	$20.3M	$29M
Probability Weighted Loss	$870K	$8.1M	$8.7M

AVERAGE LOSS SEVERITIES

Next, we will add the weighted losses together. This will give us what the losses might be on average in the event of a loss. These are unplanned as the event probability is quite low and often not included in a financial budget plan. This specific Loss Severity is represented in the last row of the next summary table below. You can think of this value as one that is representative of the full Loss range in a single summary number.

80% Prediction Intervals For Ransomware Data Breach			
Range	Low (10%)	Median (50/50)	High (90%)
Records	5M	35M	50M
Loss	$2.9M	$20.3M	$29M
Probability Weighted Loss	$870K	$8.1M	$8.7M
Loss Severity	$17.7M		

EXPECTED LOSS

Our next step is to multiply the average Loss Severity by the ransomware event probability. All of the events in this exercise are based on the assumption of “materiality.” We define materiality as simply “reportable.” Therefore, we aren’t asking about the likelihood of having a ransomware event; rather, we are asking about those events in which the financial impact is large enough to report to your insurance carrier.

To the CFO, reporting this loss to the insurance carrier is an unexpected or unbudgeted expense. Your reported loss is an empirical, mathematically unambiguous, and contractually binding example of your company's risk tolerance. Your CFO is paying a premium to be reimbursed should that loss be large enough to recover damages. Note that if you exceed your coverage limit you may find it hard to renew your coverage. This will affect the tradeoff decisions between getting enough coverage without committing moral hazard and implementing controls to avoid excessive (i.e., capitally inefficient) coverage.

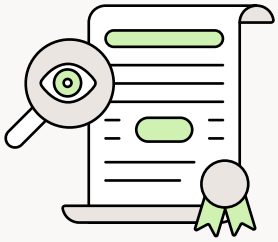

The probability of experiencing a reportable loss event for this, and all use cases in this exercise, is **2.5%**. Strictly speaking, this is a naive forecast based on quantitative surveys with CISOs and meta-analysis conducted for the book and ransomware claims experience. Again, a full cyber risk quantification engagement would create more robust values that take into consideration the local value of controls or lack thereof.

Below, $17.7M is multiplied by 2.5%. This creates the expected loss, which also shows up in the summary analysis above.

80% Prediction Intervals For Ransomware Data Breach			
Range	Low (10%)	Median (50/50)	High (90%)
Records	5M	35M	50M
Loss	$2.9M	$20.3M	$29M
Probability Weighted Loss	$870K	$8.1M	$8.7M
Loss Severity	$17.7M		
Expected Loss	$443,000		

WHAT DOES THE ANALYSIS MEAN

According to our model you should be willing to pay someone $443k to take on all the loss of ~$18M when it occurs.

If and when that will happen is a function of time. With a yearly rate of 2.5% you can expect that in 5 years there's a 12% chance of having one or more losses. In 10 years there is roughly a 22% chance of having one or more losses. [However, when an event occurs, you will face the full range of value from $2.9M - $29M, with an average centered around $17.7M.]

RANSOMWARE EXTORTION

For extortion we can use ranges influenced by publicly available data. In this case I referenced the Palo Alto Networks 2022 Unit 42 Ransomware Threat Report. The average payment for all of their cases was ~$530K. Demands ranged upwards of $50M, but such requests were met with relatively tiny payments. We will use this data to inform our ranges – in part.

Research also shows extortion requests ranging between 0.07% and 5% of revenue. In reality, the payments made are on average 50% of what is demanded. With a stated revenue of $250M and an average extortion request of 2.9%, we get $7.25M. We can then discount that by 50% to get $3.6M. The loss row below is informed by the average paid extortion value of $530K and the revenue-derived value of $3.6M. Here is the procedure to get the remaining values following the Ransomware example:

- Each loss is multiplied by the skewed distribution weights of 30%, 40%, and 30% respectively
- The distribution of losses is summarized to get the average Loss Severity
- Expected losses are derived by multiplying the average Loss Severity by 2.5%

80% Prediction Intervals For Ransomware Extortion			
Range	Low (10%)	Median (50/50)	High (90%)
Loss	$300K	$1.5M	$5M
Probability Weighted Loss	$90K	$600K	$1.5M
Loss Severity	$2.2M		
Expected Loss	$55K		

RANSOMWARE BUSINESS DISRUPTION

What about business disruption? A small handful of threats have disruption based impacts, whereas ransomware disrupts business operations almost every time. Disruption is typically measured in minutes of downtime. For small companies, that can start at $100 a minute. For larger companies the average value is closer to $5600 a minute, although I have seen questionable research that puts it closer to an average of $9000 a minute. These are just rough estimates with some stating that the cost of disruption is 50 times that of payment. This is why people pay; the ROI on maintaining availability far outweighs the current average cost of extortion.

We are going to go with the $5,600 average per minute rate, only considering a hard-down event time frame where the means of value creation are completely disrupted.

80% Prediction Intervals For Ransomware Business Disruption			
Range	Low (10%)	Median (50/50)	High (90%)
Minutes	300 Mins	870 Mins	2400 Mins
Loss	$1.7M	$4.9M	$13.4M
Probability Weighted Loss	$504K	$2M	$4M
Loss Severity	$6.5M		
Expected Loss	$163,000		

BUSINESS EMAIL COMPROMISE

Average losses to BEC are just under $200k. We have seen $43B in damages across 241,000 incidents internationally. Since this is a mid-market company with $250M in revenue, we raised the losses. The same process applies in terms of getting the distribution by multiplying losses by 30%, 40%, and 30%, respectively. That value is added together to get the Loss Severity, which in turn is multiplied by 2.5% to get the expected loss.

80% Prediction Intervals For Business Email Compromise			
Range	Low (10%)	Median (50/50)	High (90%)
Loss	$200K	$2M	$5M
Probability Weighted Loss	$60K	$800K	$1.5M
Loss Severity	$2.4M		
Expected Loss	$60K		

CLOUD DATA BREACH

For the company in question, they have regulated data distributed within a hybrid environment. Data that is in a third-party cloud is likely not subject to ransomware. It is, however, subject to compromise and exploitation. This summary table uses the exact same operation as the ransomware breach example in terms of using a $0.58 per record cost. All the other operations remain the same.

80% Prediction Intervals For Cloud Data Breach			
Range	Low (10%)	Median (50/50)	High (90%)
Records	3M	20M	30M
Loss	$1.74M	$11.6M	$17.4M
Probability Weighted Loss	$520K	$4.6M	$5.2M
Loss Severity	$10M		
Expected Loss	$250K		

CLOUD DISRUPTION

Third-party clouds typically have strong uptime SLAs. We use the same $5,600 multiplier, distributions and other operations to derive our values.

80% Prediction Intervals For Cloud Disruption			
Range	Low (10%)	Median (50/50)	High (90%)
Minutes	60	240	480
Loss	$370K	$1.34M	$2.7M
Probability Weighted Loss	$111K	$540K	$810K
Loss Severity	$1.5M		
Expected Loss	$38K		

SAAS DATA BREACH

Data owners are generally held accountable for third-party breaches. Based on interviews, we determined that a measurable amount of data is processed by SaaS vendors. The same operations apply here as they do for all breach examples seen above.

80% Prediction Intervals For SaaS Data Breach			
Range	Low (10%)	Median (50/50)	High (90%)
Records	500K	2M	130M
Loss	$290K	$1.2M	$5.8M
Probability Weighted Loss	$87K	$464K	$1.7M
Loss Severity	$2.3M		
Expected Loss	$57K		

SAAS BUSINESS DISRUPTION

There may be a handful of providers that are critical to your business. While an individual loss is rare, when you have several, the losses can add up. The same methods apply here in terms of multiplying by $5,600, etc.

80% Prediction Intervals For SaaS Disruption			
Range	Low (10%)	Median (50/50)	High (90%)
Minutes	180	720	1160
Loss	$1M	$4M	$6.5M
Probability Weighted Loss	$300K	$1.6M	$2M
Loss Severity	$3.9M		
Expected Loss	$100K		

ENTERPRISE MODELING

The likelihood per year of any single event is relatively low. Many threats have yearly probabilities that may be considerably lower than the 2.5%. Also, data breach is one of the more costly loss types so the expected losses may be lower for other examples.

The point is that you can consider the likelihood of having one or more events per year across 8 perils. With a likelihood of 2.5% (which again is a simplified value for this exercise), that becomes an 18% chance of having one or more material events a year. In 3 years you're looking at over a 45% probability with a 12% chance of having 2 or more events.

Going back to the original table, recall that the expected loss for all eight perils is $1.2M. You should not think of this as the actual amount of money you will lose on an annual basis without controls. As we have seen, any one peril that materializes could cost many times more than $1.2M.

The $1.2M, then, refers to the rational maximum amount you should be willing to pay each year to transfer the full range of losses to someone else. Of course, no one will be willing to take on the full range of your and everyone else's potential losses; therefore, you would adjust down the amount you would pay based on the loss limit offered by an insurer. The insurer is going to require that you participate at some minimal level of due diligence and active control of risk to avoid moral hazard.

This sheds some light onto why cyber insurance and security controls must work together. By spending more on limit and security controls you can mutually lower the likelihood of having impactful material events.

SAAS BUSINESS DISRUPTION

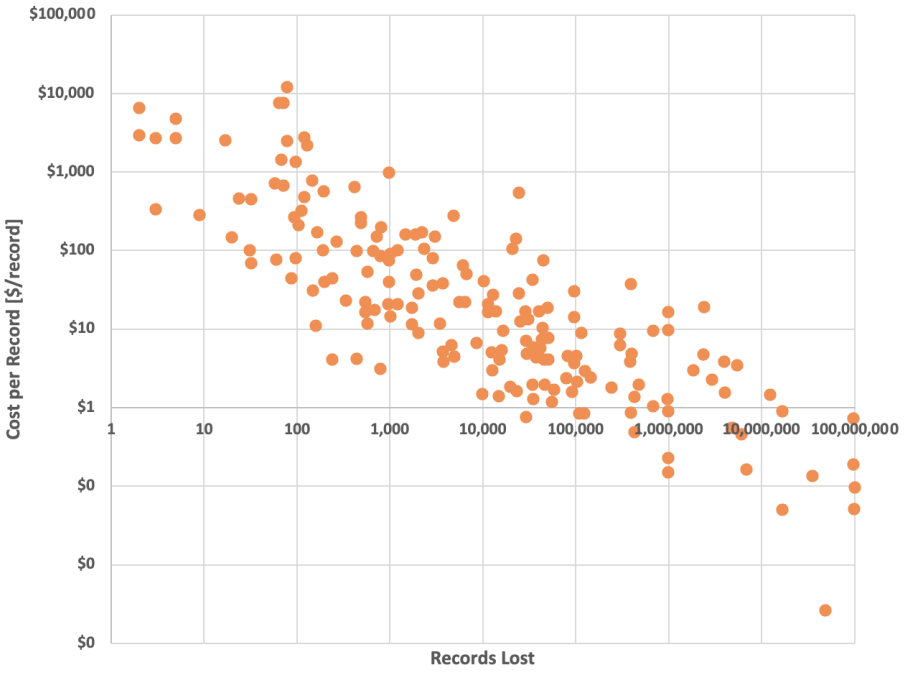

ADDITIONAL PERIL CONSIDERATIONS

Non-Exhaustive List Of Possible Perils	
Threat	Loss
Ransomware	Data Reach
Ransomware	Business Disruption
Ransomware	Extortion
Business Email Compromise	Fraud
Cloud Compromise	Data Breach
Cloud Compromise	Business Disruption
SaaS Compromise	Data Breach
SaaS Compromise	Business Disruption
Insider Threat	Data Breach
Insider Threat	Fraud
Insider Threat	Business Disruption
APT (State Sponsored)	Intellectual Property
APT (State Sponsored)	Data Breach
APT (State Sponsored)	Business Disruption
Software Supply Chain Compromise	Intellectual Property
Software Supply Chain Compromise	Data Breach
Software Supply Chain Compromise	Business Disruption

About The Author

Richard Seiersen, Chief Risk Officer for Resilience, is a 20+ year security veteran with 10 years as a Chief Information Security Officer (CISO). He has served as CISO at GE, Twilio and LendingClub. His books include "How to Measure Anything in Cybersecurity Risk" (2016) and "The Metrics Manifesto: Confronting Security with Data" (2022).