Sentinelcore-Ai

WHITEPAPER · CYBER RISK QUANTIFICATION



Cyber Risk in Euros: A Practical Guide to FAIR-Style Quantification

Replace heat maps with loss distributions the board can act on, compare controls by return on investment, and do it without a data-science team.

By Harsukh Singh

Eight years in security engineering, information security governance and identity

September 2026 · sentinelcore-ai.com

Free to share with attribution

EXECUTIVE SUMMARY

'High' is not a number

Heat maps are easy to produce and hard to act on. Two 'high' risks can differ by orders of magnitude; multiplying ordinal scores (4 × 3 = 12) produces numbers that look precise but can't be added, compared with a budget, or checked against reality.

Quantitative risk analysis expresses risk as a range of possible financial losses per year. The most widely used model is FAIR (Factor Analysis of Information Risk), published as an Open Group standard. This paper explains the model in plain terms, shows how Monte Carlo simulation turns expert estimates into loss distributions, and how to use the results for budget and control decisions.

€	10,000	P90	ROSI
risk in the board's currency	simulated years per risk	a 1-in-10-year loss	return on security spend

FAIR on one page

Risk	Annualised loss exposure			
Frequency	Loss event frequency	= threat event frequency × vulnerability		
Magnitude	Primary loss	Secondary loss × probability		
Primary	Response	Downtime	Replacement	Productivity
Secondary	Fines & judgements	Notification	Reputation	Competitive

- **Loss event frequency**: how many times per year a loss is expected to happen. 0.1 means once in ten years.
- **Primary loss**: what the organisation bears directly per event: response, downtime, lost revenue, recovery.
- **Secondary loss**: what follows from reactions of others (regulators, customers, courts), with the probability it happens at all.
- **Annualised loss exposure**: the resulting distribution of losses per year, not a single number.

Estimating with ranges, not points

Nobody knows the exact cost of a future ransomware incident, but most teams can give a calibrated range. For each factor, ask for a **minimum**, **most likely** and **maximum** value, where the minimum and maximum are the values the estimator would be surprised to see exceeded. Sources include your own incident history, insurance claims and policy limits, finance data on downtime cost, regulatory fine guidance, and industry reports.

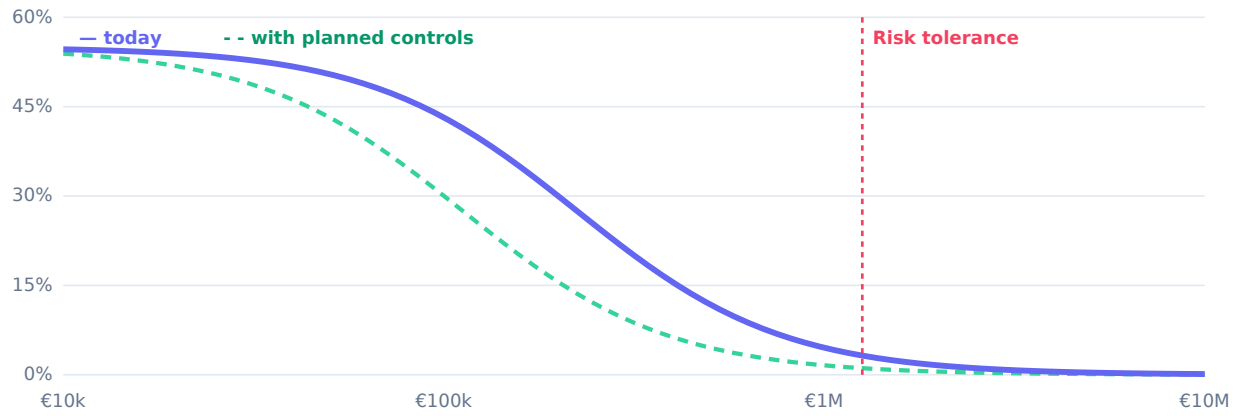
> **Calibration beats precision**
> A wide range that contains the truth is more useful than a precise number that doesn't. Record where each estimate came from; that rationale is what makes the analysis defensible to auditors and the board.

From estimates to a loss distribution

Monte Carlo simulation runs thousands of possible years. In each simulated year it samples an event rate from the frequency range, draws how many events happen (a Poisson process), and for each event samples a primary loss and, with the given probability, a secondary loss. Three-point estimates are commonly sampled from a PERT distribution, which weights the most likely value while respecting the minimum and maximum.

Output	Meaning	Typical use
Expected annual loss (mean)	Average loss across all simulated years	Budgeting; comparing risks; ROSI
P50 (median)	Half of years are better, half worse	Communicating a 'typical' year (often zero for rare events)
P90	A 1-in-10-year loss	Planning for a bad year; reserves
P99	A 1-in-100-year loss	Insurance limits; severe-but-plausible scenarios
Probability of loss	Chance of at least one event in a year	Explaining frequency to non-specialists

The loss exceedance curve



Illustrative loss exceedance curve: the probability that annual losses exceed each amount, today and with planned controls, against a stated risk tolerance.

The loss exceedance curve (LEC) shows, for every loss amount, the probability of exceeding it in a year. Adding a line for the organisation's risk tolerance turns an abstract discussion into a concrete question: is a 5% chance of losing more than €2M a year acceptable? Summing simulated years across risks gives a portfolio curve that captures bad years where several things go wrong.

Deciding with numbers: ROSI

Model a planned control as a reduction in frequency (for example, phishing-resistant MFA reducing credential-based compromise) or in magnitude (for example, tested backups reducing downtime). Re-run the simulation to get the residual expected loss. Then:

ROSI = (expected loss before − expected loss after − annual cost of control) ÷ annual cost of control

A positive ROSI means the control is expected to save more than it costs. It's a comparison tool, not a promise: use it to rank options, and pair it with the tail (P90/P99) for controls whose main value is limiting catastrophic years.

Common mistakes

- **Quantifying everything at once.** Start with the five to ten risks leadership actually discusses.
- **Scenarios that are too vague.** 'Cyber attack' can't be estimated; 'ransomware encrypting the ERP system' can.
- **Point estimates in disguise.** A range of €100k–€110k is a point estimate. Let ranges be wide.
- **False precision in reports.** Round to two significant figures and always show ranges alongside the mean.
- **Ignoring correlation.** Independent simulation is a reasonable start, but note where risks share a cause, such as one identity provider.
- **A one-off exercise.** Revisit estimates when controls, threats or the business change, and at least at each management review.

A first quantification in a week

- ☐ Day 1: pick five risks from the register and rewrite each as a specific loss scenario.
- ☐ Day 2: gather data: incident history, insurance, downtime cost per hour, regulatory exposure.
- ☐ Day 3: run a workshop to agree min / most likely / max ranges; record sources.
- ☐ Day 4: simulate, review the outputs for sense, and adjust obviously miscalibrated ranges.
- ☐ Day 5: present expected loss, P90 and the portfolio curve with one or two control options and their ROSI.

How Sentinelcore-Ai helps

Sentinelcore-Ai's Risk Quantification runs a FAIR-style Monte Carlo simulation directly on your risk register. It shows expected annual loss and 1-in-10 and 1-in-100-year losses in euros, a portfolio loss exceedance curve, and a what-if for planned controls with ROSI. Illustrative presets give a starting point for common scenarios, to be replaced with your own data.

References

The Open Group Risk Taxonomy (O-RT) and Risk Analysis (O-RA) standards (Open FAIR) · ISO/IEC 27005:2022 Information security risk management · ENISA publications on cybersecurity risk quantification · D. Vose, Risk Analysis: A Quantitative Guide (PERT and Monte Carlo methods).

About the author

Harsukh Singh has spent eight years in hands-on cybersecurity work across security engineering, information security governance, and identity and access management. That work has meant building and running information security management systems to ISO/IEC 27001:2022, auditing them against the standard, and designing identity and access controls on Microsoft Entra ID, from conditional access and privileged roles to access reviews. Harsukh founded Sentinelcore-Ai to give security and compliance teams the tooling that work called for: posture measured against vendor baselines, findings tied to risks and corrective actions, and evidence an auditor can follow.