CSURFACE

WHITEPAPER · CYBER RISK QUANTIFICATION

Converting technical risk into financial decisions

Cyber Risk Quantification: how the FAIR methodology translates vulnerabilities, assets, and threats into annual loss expectancy (ALE) and Value at Risk — in dollars, in the language of the audit committee.

CAPABILITY	AUDIENCE	EDITION
Cyber Risk Quantification	CISO · Risk · CFO · Audit	2026

READING TIME	CLASSIFICATION	SOURCE
10 pages	Public	CSURFACE Platform

EXECUTIVE SUMMARY

Risk that has no price has no decision

The thesis of this whitepaper, in one page — what the board needs to know before the methodological detail.

AVERAGE COST OF A BREACH
$10.22M
United States average in 2025 (IBM Cost of a Data Breach)

STATUTORY BREACH DAMAGES
$100–750
per consumer, per incident (CCPA private right of action)

TIME TO IDENTIFY AND CONTAIN
241 days
global average in 2025 — the lowest in nine years (IBM)

Mature security programs produce data in abundance: number of open CVEs, mean time to remediate, scan coverage, control posture. These are useful indicators for the operational management of the SOC. They are, however, practically **useless for the audit committee, the board, and the CFO** — who must decide on budget, capital reserves, and risk appetite.

The root cause is a failure of language. Security operates in CVE, CVSS, and MTTR; the board operates in exposure, EBITDA, and return on capital. Without a common denominator, security investment is approved without objective criteria, the decision to accept a risk carries no explicit trade-off, and cyber insurance renewal rests on the broker's worst-case assumptions rather than on the organization's own.

Cyber Risk Quantification (CRQ) closes this gap. Applying the FAIR methodology — Factor Analysis of Information Risk — CRQ converts technical telemetry into **measurable financial exposure**: annual loss expectancy (ALE) and Value at Risk, expressed in dollars. This is the metric that the board approves, that external audit documents, and that the regulator expects to see.

The thesis. "We have 47 critical vulnerabilities" is a data point with no unit of decision — it does not support budget approval, reserve sizing, or the definition of risk appetite. "Our cyber exposure is $11 million per year, with VaR P90 of $20.5 million" is a financial quantity that the board already knows how to assess and compare against other exposures in the portfolio.

What this whitepaper covers

- Why technical indicators do not support capital decisions.
- Where qualitative heat maps and external posture scores stop.
- The FAIR methodology applied to the US regulatory context.
- CSURFACE's quantitative model, with a fully worked example.

THE LANDSCAPE · PART 01

Two languages, one decision

Why security and the board do not understand each other — and what that gap costs the organization.

THE LANGUAGE OF SECURITY

Technical indicators

Open CVEs · CVSS · MTTR · Scan coverage · Control posture

→

THE LANGUAGE OF THE BOARD · CRQ TRANSLATES

Financial decision quantities

Annual loss expectancy (ALE) · Value at Risk · P90 / P99 · Exposure in dollars · Risk appetite · Capital reserve · Remediation cost-benefit

The technical indicator does not reach the decision table

The security team operates with technical indicators — number of open CVEs, mean time to remediate, number of scans, coverage percentage. They are essential for day-to-day management. But they do not answer the questions that define the program:

- The **audit committee** needs to assess whether the security program is adequate to the business risk.
- The **board** needs to approve a security budget proportional to the real exposure.
- The **CFO** needs to size capital reserves for cyber risk.
- The **cyber insurance broker** needs to quantify the exposure in order to price the policy.
- The **external audit** assesses the program under frameworks such as COSO ERM.

None of these questions is answered with "we have 47 critical ones." All of them are answered with a financial value.

What the absence of quantification produces

1. **Chronic underinvestment.** The board prioritizes what it understands. Without a financial value attached to cyber risk, the security budget competes at a disadvantage with initiatives that already arrive quantified.
2. **Arbitrary risk decisions.** Without quantification, accepting or treating a risk results from undocumented judgment — without the explicit trade-off between the cost of remediation and the expected cost of the incident.
3. **Mispriced cyber insurance.** Without its own quantification, the company accepts the broker's worst-case assumption — and pays a premium for coverage misaligned with its real exposure.

The gap has a direct cost. The average cost of a data breach in the United States reached **$10.22M** in 2025 (IBM Cost of a Data Breach 2025), and in the United States breach exposure arises from a 50-state patchwork of notification laws, FTC Act §5 enforcement, and state statutes such as the CCPA (statutory damages of $100–$750 per consumer per incident). Without quantifying this exposure, the organization has no criterion to assess whether its level of security investment is adequate, insufficient, or excessive — and each of these deviations has distinct and measurable financial consequences.

Why this matters. Risk that has no price does not enter the capital decision process. It stays outside the budget, outside the corporate risk map, and off the board's agenda — which means the organization cannot manage it with the same discipline applied to the other risks of the business.

CURRENT LIMITATIONS · PART 02

The heat map is opinion dressed as data

Why qualitative matrices, external scores, and CVSS do not replace financial quantification.

Qualitative 5×5 matrices do not support decisions

Red-yellow-green risk matrices are popular because they are simple. That simplicity, however, conceals four structural limitations:

WARNING **Ambiguity**
Does "high probability" mean 30% or 70%? The definition varies among participants, and the matrix has no mechanism to resolve that divergence.

WARNING **They do not allow comparison**
Two "high" risks are not equivalent — but the matrix treats them as equal.

CRITICAL **They do not aggregate**
Two "highs" do not add up coherently to "very high." There is no arithmetic behind the colors.

CRITICAL **They do not connect to risk appetite**
The board can say it will not tolerate a loss of $100 million. It cannot say it will not tolerate "red."

The point. A heat map communicates relative priority among already-known risks. It does not communicate magnitude, does not allow capital trade-offs, and does not speak to the board's risk appetite. For the decision table, it is a starting point — never the destination.

Other proxies — and where each one stops

Two other approaches are common. Both have value, and both have a clear limit.

1. **External posture scores.** Ratings from 0 to 1000 based on visible signals — certificates, configurations, exposures — are useful for benchmarking. But they do not translate into dollars, they ignore asset criticality and regulatory context, and they do not answer "how much to invest?".
2. **CVSS as a measure of risk.** CVSS quantifies theoretical technical severity. A score of 9.8 on a test machine with no data is not the same risk as 9.8 on a production server with card data — and CVSS, on its own, assigns the same number to both.

None of these instruments is dispensable — all of them feed the quantification model. The error is to treat any one of them as the financial answer the board expects.

THE CSURFACE APPROACH · PART 03

FAIR: decomposing risk to be able to calculate it

The principle of FAIR and the data the platform uses to feed each variable of the model.

THE DECOMPOSITION OF FAIR · FROM OPAQUE RISK TO CALCULABLE VARIABLE

1 **Risk**
Loss Event Frequency × Loss Magnitude — the whole, decomposed into two factors.

→

2 **Loss Event Frequency**
Threat frequency × Vulnerability — the annual chance of a material event.

→

3 **Loss Magnitude**
Primary loss + Secondary loss — the expected cost per event.

→

4 **Distribution**
Each variable becomes a probabilistic range for the Monte Carlo simulation.

The principle: separate and quantify

FAIR — Factor Analysis of Information Risk — is an open standard that decomposes risk into observable variables, instead of treating it as an opaque whole. The structure is that of the diagram above:

$$\text{Risk} = \text{Loss Event Frequency} \times \text{Loss Magnitude}$$
$$\text{Loss Event Frequency} = \text{Threat frequency} \times \text{Vulnerability}$$
$$\text{Loss Magnitude} = \text{Primary loss} + \text{Secondary loss}$$

Each variable is estimated as a **probability distribution** — which acknowledges the inherent uncertainty and enables the Monte Carlo simulation. The primary loss is the direct cost of the incident; the secondary loss includes fines, loss of customers, and brand damage.

FAIR does not replace technical judgment — it organizes it. Each estimate becomes explicit, auditable, and revisable.

What feeds the model

CSURFACE's CRQ does not require the organization to estimate parameters without an empirical basis. The model is fed by the data that the platform itself already collects continuously:

- **Technical telemetry.** Number of assets, vulnerabilities, and criticality — in real time, coming from the platform's continuous discovery.
- **Threat telemetry.** CSURFACE's *threat sensor* goes beyond static CVSS scores and exploitation-probability indices and active-exploitation catalogs: it observes threat activity, exploitation attempts, exploit intelligence, and campaign indicators — real-time signals that calibrate the frequency of events in the model.
- **Business context.** Sector, revenue, and number of employees, used to calibrate probability and magnitude against market references.
- **Incident history.** The organization's own events, when they exist, combined with public sector references.
- **Control posture.** Log coverage, MFA, patch cycle, and backups — factors that modulate probability.

The accuracy of the model grows in proportion to the breadth and currency of the telemetry collected. The calibration is continuous, tracking the telemetry as it evolves.

The model is fed by observed evidence. The assets, vulnerabilities, and criticality that enter the risk calculation come from the continuous discovery of the CSURFACE platform itself — the same one that, across 68 organizations, mapped **122 thousand external assets** and, in 16 detailed analyses, **2,361 vulnerabilities**. The quantification is as solid as the inventory that sustains it.

THE CSURFACE APPROACH · PART 03

The model, parameter by parameter

How probability and magnitude are calculated, with the multipliers calibrated by the 2025 benchmark.

Sector base probability

The annual chance of a material incident starts from a per-sector reference, adjusted by a maturity multiplier (Basic 1.4× · Intermediate 1.0× · Advanced 0.65×).

SECTOR	P(INCIDENT / YEAR)	SECTOR	P(INCIDENT / YEAR)
Health	34%	Telecom	26%
Financial	32%	Education	24%
Government	30%	Industry	22%
Retail	28%	Energy	20%

Magnitude — the SLE formula. The expected loss per event (Single Loss Expectancy) starts from the national average cost and is adjusted by sector, size, and surface size:

$$SLE = \$10.22M \times \text{sector mult.} \times (\text{revenue} \div \$200M)^{0.35} \times (\text{no. of assets} \div 500)^{0.25}$$

The fractional exponents avoid naive linearity: doubling revenue does not double the expected loss.

Cost multipliers by sector

Derived from the ratio between the sector's average cost and the national average from the IBM Cost of a Data Breach 2025.

SECTOR	MULT.	SECTOR	MULT.
Health	1.59×	Telecom	0.92×
Financial	1.24×	Education	0.78×
Services	1.18×	Retail	0.72×
Energy	1.00×	Government	0.68×
Industry	0.95×	Others	1.00×

CSURFACE estimates derived from IBM 2025 sector proportions: Health $11.43M; Financial $8.92M; Services $8.51M. National average (IBM 2025): $10.22M.

From loss per event to the board's number

- **ALE** = adjusted probability × SLE — the annual loss expectancy.
- **VaR P90** — the 90th percentile of the loss distribution, which exceeds ALE and guides capital decisions.
- **VaR P99** — the 99% tail, which materially exceeds ALE and serves for reserves and insurance coverage.

APPLICATION · PART 04

From technical inventory to a number in dollars

A complete calculation, step by step: a retailer with $1.2 billion in revenue.

ANNUAL LOSS EXPECTANCY (ALE)

$11.1M

adjusted probability × SLE

VALUE AT RISK · P90

$20.5M

the number for capital decisions

REGULATORY & LITIGATION EXPOSURE

$17.5M

CCPA + state-AG + class action, weighted by probability

The scenario and the calculation

An omnichannel retailer, revenue of **$1.2 billion**, 4,500 employees, **8,700 assets** discovered by the platform, and intermediate security maturity. The model runs through each parameter:

- **Probability.** Retail base (28%), adjusted by maturity and surface size — adjusted probability of approximately **56%**.
- **SLE.** $10.22M × 0.72 (retail) × $(1{,}200 \div 200)^{0.35}$ × $(8{,}700 \div 500)^{0.25}$ ≈ **$19.8 million**.
- **ALE.** 56% × $19.8M ≈ **$11.1 million** per year.
- **VaR P90 / P99.** As a parametric estimate of the tail, 1.85 × ALE ≈ **$20.5M**; 3.6 × ALE ≈ **$39.9 million**.
- **Regulatory & litigation exposure.** CCPA statutory damages ($100–$750 per affected consumer) plus state-AG settlement and class-action exposure × retail sensitivity (1.3) × probability ≈ **$17.5 million**.

The result, side by side

METRIC	VALUE	WHAT IT IS FOR
SLE	$19.8M	Loss per event
ALE	$11.1M	Annual budget
VaR P90	$20.5M	Capital decision
VaR P99	$39.9M	Reserve / insurance
Regulatory & litigation	$17.5M	Regulatory exposure

Values calculated with CSURFACE's recalibrated model, based on IBM Cost of a Data Breach 2025.

The language of the board. "Our annual cyber exposure is $11.1 million, with VaR P90 of $20.5 million and regulatory exposure of $17.5 million. Reducing ALE by 20% requires an incremental investment whose cost-benefit ratio is favorable." This is the formulation that allows the board to deliberate with the same rigor applied to other risks in the portfolio.

THE CSURFACE APPROACH · PART 03

Exposure in dollars is a living metric

How ALE tracks every change in the inventory and in the threat — and how CRQ supports regulatory documentation.

TRADITIONAL RISK ASSESSMENT

1x

per year — a static snapshot that ages the day after delivery.

→

CSURFACE'S CLOSED-LOOP CRQ

24/7

ALE tracks every change in the inventory, in the threat, and in remediation.

The closed loop in operation

CSURFACE's CRQ is a living calculation. It is integrated with the platform's other capabilities, and the financial exposure recalculates whenever reality changes:

1. **A critical vulnerability is remediated.** The adjusted probability drops and ALE is recalculated automatically on the same day.
2. **A critical asset is discovered.** Continuous discovery adds it to the scope, the surface size rises, and the SLE is re-estimated.
3. **The threat sensor detects active exploitation of a CVE in the inventory.** Confirmed exploitation raises the probability, and the exposure is repriced upward.
4. **A corporate credential leaks.** Credential monitoring incorporates the additional impact into the magnitude calculation.

Regulatory compatibility

Financial quantification is exactly what regulators and governance frameworks expect. CRQ directly supports:

- **FFIEC / GLBA.** Integrated risk management with financial quantification of cyber exposure.
- **SEC.** Board oversight (Caremark) and disclosure of material cyber risk (Reg S-K Item 106; Form 8-K Item 1.05).
- **FTC / state privacy law.** Reasonable-security documentation and data-protection impact assessments supported by quantitative data.
- **ISO 27005 and COSO ERM.** Risk management frameworks fully compatible with the FAIR approach.

From audit to decision. Because every estimate in the model is explicit and traceable, CRQ produces an auditable trail: the audit committee can verify the origin of each parameter, and the board receives a monthly digest with the updated exposure — instead of an annual report with data that no longer reflects operational reality.

APPLICATION AND RESULTS · PART 04

From the public calculator to the integrated module

From sector estimate to defensible report: the implementation path and the deliverables at each stage.

TIME-TO-VALUE · FROM SECTOR ESTIMATE TO DEFENSIBLE REPORT

1 **Day 1**
Immediate estimate via the public calculator, without internal data.

→

2 **Week 1**
Continuous discovery feeds the model with the real asset inventory.

→

3 **Week 2**
The threat sensor — observed exploitation and exploit intelligence — calibrates the probability.

→

4 **Month 1**
First executive CRQ report, ready for the audit committee.

→

5 **Month 3+**
Closed loop in steady state: ALE fluctuates daily and generates a monthly digest.

Two entry points

CRQ offers two modes of use. The first operates exclusively with public sector parameters:

- **The public risk calculator.** Available at csurface.io, it estimates the sector's ALE and VaR from a few parameters — revenue, sector, number of employees, and assets. It is the entry point: a first reading of exposure in seconds, without registration.
- **The integrated CRQ module.** The full version, inside the platform. It replaces the sector estimates with the organization's real inventory, discovered and classified continuously, and maintains the closed loop of recalculation.

The calculator gives the order of magnitude. The integrated module gives the number that is defensible before an audit.

What each stage delivers

Public calculator — order of magnitude
A first reading of exposure in seconds, without registration or internal data.

WARNING **Real inventory — replaces the estimate**
Continuous discovery swaps sector references for the discovered and classified inventory.

WARNING **Threat calibration — real probability**
The threat sensor adjusts the probability by observed exploitation and exploit intelligence, replacing static scores with operational evidence.

CRITICAL **CRQ report — defensible number**
Traceable ALE and VaR, ready for the committee's agenda and for external audit.

What changes in the boardroom. Cyber risk ceases to be a technical item presented once a year and becomes a quantified line in the corporate risk map — comparable, aggregable, and directly tied to the risk appetite approved by the board.

NEXT STEPS

The next risk meeting can begin with a number in dollars

Cyber Risk Quantification is one of the capabilities of the CSURFACE platform — a Continuous Exposure Management platform. In a single integrated platform, it brings together continuous discovery of the external surface with Machine Learning, analysis of the digital supply chain of vendors, active exploitability validation where test coverage exists, with passive monitoring across the rest of the surface, threat intelligence with dynamic prioritization, leaked-credential monitoring, and financial risk quantification. It operates in a fully external manner — without an agent and without installation — with optional cloud, WAF, and CIEM integrations for organizations that require in-depth visibility of those layers.

Use the risk calculator

Estimate your sector's ALE and VaR in seconds, without registration — the entry point to financial risk quantification.

Discover the CRQ module

In a demonstration, see the FAIR analysis applied to the organization's real inventory, with the closed loop of continuous ALE recalculation.

Read the Threat Intelligence whitepaper

Dynamic vulnerability prioritization — one of the critical inputs to CRQ's probability calculation.

See your exposure in dollars — start with the public risk calculator

Get a free preliminary analysis