ONAM SECURITY

TECHNICAL WHITEPAPER · 02

Cloud risk in dollars

How Onam prices a verified attack path with FAIR — using named, external inputs a board can defend.

FAIR / ALE | IBM 2024 costs | Sensitivity × | Regulatory × | Exposure roll-up

THE ARGUMENT, IN ORDER

00 Executive summary

Severity tells you what is broken. Dollars tell you what to fix first.

Security teams are drowning in "criticals". Boards are not asking how many — they are asking how much. This paper explains how Onam turns a verified attack path into a defensible financial figure using FAIR, the open standard for risk quantification, so the conversation moves from a spreadsheet of severities to a single number a board can act on.

FAIR — Factor Analysis of Information Risk — is a published, peer-reviewed methodology maintained as an open standard by The Open Group. It decomposes risk into estimable factors and produces an Annualised Loss Expectancy (ALE) in currency. Onam implements FAIR as the fourth layer of its pipeline: after a path is found, verified and mapped to MITRE ATT&CK, it is priced. The pricing is not a marketing gauge invented in-house; it is a recognised model fed with named, external inputs.

> **Why this is the wedge.** A security graph and attack paths are now common across the CNAPP market. A defensible dollar figure — derived from a published model and named external cost data, attached to a specific verified path — is not. It is the difference between "we have 200 criticals" and "this one over-privileged role carries the largest single slice of our quantified exposure."

01 The FAIR model, briefly

The model is not ours. FAIR is an open, published standard, which is exactly why the number survives scrutiny.

FAIR expresses risk as the product of how often a loss happens and how large it is when it does.

The strength of FAIR is that each factor is estimable and inspectable. When Onam prices a path, a reviewer can open the calculation and see which input drove the number — not a black-box score between 0 and 100.

02 The inputs — named, external, defensible

The magnitude inputs come from a published benchmark, not from us.

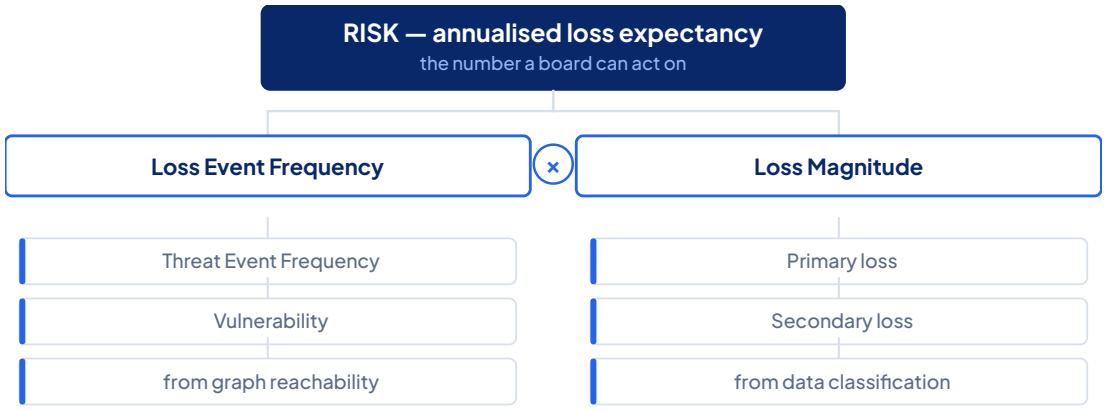

FIGURE 1 — FAIR DECOMPOSITION

Every leaf is computed, not estimated. That is the whole difference between a priced exposure and a workshop number.

A dollar figure is only as credible as the numbers behind it. Onam deliberately sources its magnitude inputs from a recognised external benchmark and applies transparent multipliers.

Input	Source / basis	Effect
Per-record breach cost	IBM Cost of a Data Breach 2024, by industry	Healthcare ≈ $10.93/record; default ≈ $4.45/record
Data-sensitivity multiplier	Onam data classification (DSPM)	Restricted ×3.0 → Confidential → Internal → Public
Regulatory multiplier	Applicable regime (GDPR, HIPAA, PCI)	Raises magnitude where statutory penalty applies
Frequency & vulnerability	Graph-derived reachability and control state	Replaces a workshop estimate with evidence

Every multiplier is visible and adjustable. A customer who disagrees with an input can change it and see the figure move — which is the point of an inspectable model.

> **The honest limit.** A quantified figure is an estimate, not an invoice. FAIR produces a defensible range built on stated assumptions; it does not predict what a specific breach will cost. Anyone presenting it as a guarantee is misusing the model, including us.

03 From path to price

Pricing happens after verification, never before. An unverified path is not priced, because a number attached to a route that does not exist is worse than no number.

1. **Find.** Graph traversal from an entry point to a crown jewel across 7 clouds.
2. **Verify.** Five-domain verification before a path is called confirmed.
3. **Map.** MITRE ATT&CK technique mapping for the steps on the path.
4. **Price.** FAIR applied to the verified path, using the inputs above.

04 Rolling up — from one path to a portfolio

A remediation queue you can argue with is worth more than one you have to trust.

Individual path prices aggregate into an exposure roll-up: total quantified exposure, the largest contributors, and the choke points where a single fix removes the most dollars. That is the view a board can act on, and the view that makes a remediation queue arguable rather than arbitrary.

Because the same graph and the same 9,853 CSPM posture rules feed both the compliance and the risk lenses, a control gap and a priced path are two readings of one underlying fact — not two systems to reconcile.

05 Summary

Every number in this paper can be taken apart — that is the point of it.

- FAIR is an open, published standard — the model is not ours, only the implementation is.
- Magnitude inputs are named and external (IBM Cost of a Data Breach 2024), not invented.
- Every factor is inspectable; a reviewer can see which input drove the number.
- Pricing follows verification. An unverified path is never priced.
- The output is an estimate with stated assumptions, not a prediction of a specific loss.