

Cyber Risk in the Age of AI: From Threats to Financial Impact

An actuarial perspective

Boris Kouambo Ndam · Founder & CEO, Üsta Digital

• THE QUESTION YOU'VE ALL HAD

"What's our exposure to a major cyber event — in Rand?"

Every one of you has been asked a version of this by a board, an audit committee, or a regulator. And these are the answers our tools produce.

Amber

HEAT MAP

Level 3

MATURITY

62 / 100

COMPOSITE SCORE

98.2%

PATCH COMPLIANCE

≠

R?

WHAT WAS ASKED

None of these are denominated in Rand. The question was.

AI is the reason quantification stopped being optional.

Not because it added a threat to your register. Because it moved five things the old approach assumed would hold still. Two of them I want to spend real time on.

01

Frequency

You know it went up. Nothing you own can tell you by how much



02

Severity

Your MTTD has not got worse. What it is worth has



03

Correlation

Your register has rows. It has no edges



04

Exposure

The one you did to yourself, and the only one you can stop doing



05

Stationarity


And underneath all four: your report was accurate the day it was signed. This is the assumption that made point-in-time assessment survivable, and it is the one AI broke.



Your register assesses risks one row at a time. Nothing fails one row at a time.


A register is a list, and a list has no edges. Every risk scored on its own. Every vendor rated on its own. No field on the form for "these eleven rows go down together." You can complete every row correctly and still not see it.

ASSESSED SEPARATELY



- 01 Concentration.** You and your vendors are converging on the same few model providers, endpoints and tooling. You can see your own. You cannot see theirs.
- 02 Monoculture.** The same models now write code and configuration across the whole market, so one flaw class travels further than it used to.
- 03 No throttle.** Attacker capacity was quietly one of your better protections. A crew could only run so many campaigns at once. That constraint is gone.

EACH ROW SCORED ON ITS OWN · 3 SHARED DEPENDENCIES UNDERNEATH ALL OF THEM




Body barely moves. Tail thickens.

You already know which dependency keeps you up at night. What you do not have is a way to show it is worth more than one row.

Correlation only appears when the losses are modelled jointly. What comes out is a ranking the register cannot produce: not which dependency is most critical on its own, but which one takes the most of you down with it.

Point-in-time assessment assumed the world holds still between assessments.

• What your assessment says — What is actually true — Risk you are carrying unknowingly — Improvement you get no credit for



A stale estimate that always errs the same way is at least usable. One where you cannot tell the sign of the error is not.

What the problem requires is a measurement that updates when the evidence does: external attack surface, internal control state and business context feeding the same model continuously, with a method underneath it that does not change when the number does.

A colour, a maturity level and a security rating are all the same kind of thing. None of them is a quantity.

Amber

HEAT MAP · RISK
REGISTER

Level 3

MATURITY MODEL · NIST
CSF

A · 740

SECURITY RATING ·
OUTSIDE-IN

Two firms. Same rating. Not the same risk.

FIRM ONE

50,000 records
No client funds held

BOTH RATED
A · 740

FIRM TWO

1.8m records
R400m of client money

Thirty-six times the records, and four hundred million rand of other people's money. The rating cannot see any of it.

- 01 No unit.** "High" is not a quantity, and neither is 740. You cannot add them, compare them across business lines, or test them against an appetite statement.
- 02 No exposure term.** A rating scores how you *look*. It knows nothing about what you hold. Loss is how often multiplied by how much, and every one of these instruments is missing the second half.
- 03 No tail.** A rank cannot separate a bad year from a company-ending one. The correlation effect is invisible to all of them by construction, and so is everything we just covered.
- 04 It cannot price a decision.** The board is asking whether to spend R4m. A grade cannot answer, so the decision falls back on advocacy.

And to be fair to the ratings: they are real observed evidence, they update continuously, and a heat map never does. **We consume exactly that class of evidence ourselves.** The point is not that the signal is bad. It is that a signal without an exposure term is not a loss. A security rating is an **input** to quantification, not a substitute for it.

IF YOU DON'T AND THE RULES CHANGED UNDERNEATH YOU

Four instruments in two years. All of them put cyber on the governing body, not in IT.



Joint Standards 1 & 2 — the core obligation

Banks, insurers, market infrastructures, CIS managers, retirement funds, discretionary and administrative FSPs. The **governing body is ultimately responsible**. Together they require a cyber strategy and framework, identification and assessment of cyber risks, controls that keep risk **within acceptable limits**, classification of processes and assets by criticality, and **metrics to track and manage cyber risk, reported in both a technical and a business context**.

Banks Act · Insurance Act — where a number is already compulsory

ICAAP for banks and **ORSA** for insurers both require the institution to assess its own material risks and hold capital against them. Capital is denominated in Rand by definition. If cyber is material — and for a digital-first institution it is — there is no version of that assessment that ends in a colour.

SARB Directive 01 of 2024 — the payment system

Payment institutions, operators and payment-system FMIs. Requires a cyber-resilience framework and a stated **cyber-risk tolerance**, identification of **interconnected risk across the payment system**, third-party assessment at onboarding and annually, and controls **proportionate to the threat landscape, risk tolerance and systemic role**.

None of these says “calculate an expected annual loss in Rand.” All of them ask for something you cannot produce without one. A limit needs a unit, and “we accept up to amber” is a preference, not a limit.

One question, asked properly: what does a year of this cost us, and how bad can a bad year be?

01 • EVIDENCE

Your control posture, from the assessments you already run. Your external attack surface, which is observable. And business context: **records held, revenue at risk, funds held, what an hour of downtime costs**. Almost none of this is new work.

02 • HOW OFTEN × HOW MUCH

Per scenario, a **frequency** — how often this lands on a firm like yours, conditioned on your evidence rather than a benchmark average. And a **severity** — built up from the cost components that actually apply: response, notification, business interruption, legal, regulatory.

03 • AGGREGATE

Simulate the year tens of thousands of times and keep **every** outcome, not just the average one. Most years are quiet. A few are not. You need to see both.

04 • READ IT OFF

An **expected annual loss**, and the shape around it: a typical year, a bad year, a 1-in-100 year. All in Rand, all comparable to every other number your business runs on.

Why the number is worth trusting

Every assumption is written down. You can disagree with any one of them and see exactly what it changes. That is the difference between a model and an opinion, and it is the only reason anyone signs off a forward-looking number.

It is conditioned on your evidence, not a benchmark. Change a control and the number moves for a reason you can trace back to a specific input.

It does not need perfect data. Where the data runs out you put judgement, you write the judgement down, and you test how far the answer moves when you change it. That is standard practice in every catastrophe line ever priced — nobody has enough earthquakes either.

None of this is new. It is the **loss distribution approach**: written into Basel for operational risk, and used by every insurer that prices catastrophe cover. Banks have published forward-looking credit loss estimates inside audited financial statements for years, and **I built those models**. Same machinery, pointed at a different risk.

This is the moment the gap becomes visible.

Illustrative composite — a mid-market South African digital bank. Not a client, not a real company.

TODAY'S ANSWER

Amber

Composite score 62/100. Third-party risk: Medium.



SAME EVIDENCE, QUANTIFIED

R3.88m

 expected annual loss

Breach and fraud only. A floor, not a total.

01 · EVIDENCE

Control posture

your own assessments

External attack surface

observed from outside

Business context

1.8m records · R400m deposits

↓ CONDITIONS THE MODEL

02 · HOW OFTEN × HOW MUCH

Data breach

1 event every 4–5 years × ~R12m when it lands

R2.75m
a year

Cyber fraud

More often × far smaller each time


R1.13m
a year

EXPECTED ANNUAL LOSS

R3.88m

↓ 03 · SIMULATE TENS OF THOUSANDS OF YEARS

04 · EVERY YEAR THE BUSINESS COULD HAVE, RANKED BY HOW BAD IT WAS



Simulated losses, producing financial metrics you can read straight off: **a typical year, a bad year, and the 1-in-100.**

Roughly five years in six cost less than the average. The average is not the typical year, and it is not the bad year either. A colour cannot tell you that either of those exist.

"How much cyber risk can we afford to take?" Risk appetite stops being an adjective.

R1.0m

TYPICAL YEAR

R3.88m

EXPECTED ANNUAL LOSS

R5.8m

1 YEAR IN 10

R9.5m

1 YEAR IN 20

R35.9m


1 YEAR IN 100

- 01 The expected loss is a P&L question.** This is what a normal year costs you, so this is what should be sitting in a budget line rather than arriving as a surprise.
- 02 The 1-in-20 is the bad-year question.** Can the balance sheet absorb R9.5m without a conversation nobody wants to have? That is a question your CFO can answer in about four seconds, and has never been asked.
- 03 The 1-in-100 is the survival question.** It is not a budgeting number. It is what cover and capital exist for, and it is the one that decides whether a bad year is survivable or terminal.

Now the appetite statement can be written. "We accept up to **R4m** expected, and up to **R35m** at a one-in-a-hundred year" is a limit you can test against and be held to. "We have a low appetite for cyber risk" is a sentence. Joint Standard 2 requires controls that keep risk **within acceptable limits**; the SARB Directive requires a stated **cyber-risk tolerance**. A limit needs a unit, and this is the unit.

"Which security investments are actually worth making?" Two of these score identically. One is worth six times the other.

WHAT YOUR SCORES SAY



01 Same score. Six times apart. Vulnerability management and mobile controls come out of this assessment scored identically. One is worth R0.83m a year, the other R0.13m. Three more share a different identical score and run from R1.36m down to R0.49m. Whatever is driving that ranking, it is *not* how much room each one has left to improve.

02 And you already have the other half. This tells you what closing each gap is *worth*. You are the person who already knows what closing it costs. Put the two together and you have a ranked, costed remediation roadmap, which is a thing no maturity model has ever produced.

03 Which is what lets you say no. "We are not doing that this year, and here is the Rand value of what we are doing instead." That sentence survives a CFO. "This improves our score by four points" does not.

Their strongest domain, identity and access, lands fifth of six. Domains bearing only on scenarios outside this two-scenario view are excluded rather than shown as zero.

"Which cyber risks should we actually prioritise?"

It depends entirely on what you are protecting.

The average year WHAT YOU BUDGET FOR



The tail WHAT DECIDES WHETHER YOU STILL EXIST



POPIA civil liability Breach response & recovery Cyber fraud The administrative fine everyone argues about is under 1% of both.

- 01 Same three things. Completely different order.** Cyber fraud is nearly a third of your average year and effectively none of your risk of not existing. POPIA civil liability is the reverse, and it is the half nobody prices — a data subject sues directly, and fault is not an element: they show the breach and the harm, not that you were negligent.
- 02 Two lists, two jobs.** The average tells you what to *work on*. The tail tells you what to *insure, capitalise and escalate*. They are answers to different questions and they do not rank the same way.

Every instrument you currently own shows you a version of the top bar. None of them shows you the bottom one. Prioritise off the average alone and you will spend the year working hard on the thing that was never going to end you.

"What do we actually do about each one?" You already run all four. Now check you are running them on the right risks.

	THE QUESTION IT ANSWERS	ON THIS BOOK	WHAT IS ACTUALLY HAPPENING TODAY
MITIGATE	Which control area buys the most reduction in expected loss?	R1.36m top control domain	Ranked by maturity score, so the R0.47m domain gets funded ahead of the R1.36m one.
TRANSFER	Is the insurance cover we buy sized against anything other than last year's cover?	R50m = a 1-in-130 year	Nobody chose 1-in-130. It came out of a renewal, a broker call, or a peer benchmark.
ACCEPT	Can the balance sheet absorb a bad year without a conversation nobody wants?	R9.5m 1 year in 20	"Low appetite for cyber risk." Nobody has ever stated what we actually accept.
AVOID	Which exposure could we simply choose not to carry?	1.8m records drive 94% of the tail	Never asked, because nobody knew the record base was the tail.

This is your own vocabulary, not mine. You have been running all four of these for years. On this evidence, three of them are being decided by default.

"How do I show the board what my team is actually worth?" You are the only executive in that room without a number.

WHAT YOU CAN SAY TODAY

"We've moved from maturity level two to level three, closed 94% of critical findings, and patch compliance is at 98.2%."

All true. The board hears activity, and has no way to tell whether it was worth what it cost.

WHAT YOU CAN SAY WITH A NUMBER

"Our expected annual loss is R3.9m. The programme we're asking you to fund takes R1.4m off that, every year. And a bad year is R9.5m — that's the one we'd like your view on."

Same work. Same team. Now it is a business case, and it competes on the same terms as everything else on the list.

- 01 You are funded to produce an absence.** When the work succeeds, nothing happens. You cannot show a board the breach that did not occur, so every year you are asking them to pay for something invisible. A number is the only way to put a size on an absence.
- 02 You can price what you already prevented.** Run the same model against last year's posture and this year's. The difference is what your team delivered, in Rand. That is a performance figure rather than a plea, and most security functions have never had one.
- 03 It defends you in both directions.** It is also the only way to answer "what happens if we halve this?" — a question you will eventually be asked, and which currently has no answer except conviction.

You have been protecting an enormous amount of value. Almost none of it has ever been written down in Rand. That is not a failure of your work. It is a failure of the instrument you were handed to describe it.

You do not need a vendor to start this. You need one scenario and a spreadsheet.

- 01 Pick one scenario, not seven.** The one your board already worries about. For most of you that is a data breach or ransomware. You are not modelling your risk universe; you are pricing one thing properly.
- 02 Write down what you already know.** Records held, revenue at risk per day of downtime, recovery times, what last year's incidents actually cost. Most of it is already in your BIA and your post-mortems.
- 03 Put a range on how often, not a point.** Your own experience blended with what firms like yours see. Write the assumption down so someone can argue with it. That is the whole trick.
- 04 Report it as two numbers, not one.** An average year and a bad year, both in Rand. The moment those exist, every conversation we have talked about changes.

Walk into your next risk committee with three blanks

- 01** What is our expected annual cyber loss, in Rand? **R** _____
- 02** What is our 1-in-20-year exposure, and has it moved as we have adopted AI? **R** _____
- 03** Does our security spend move either number? **R** _____

Under **King V** the governing body now has to set an appetite for this. Those blanks are the appetite, and somebody is going to be asked to fill them in.

Nothing on this slide needs a budget, a procurement cycle, or my company. It needs one scenario and an afternoon.

• THANK YOU

Let's talk numbers, not scores.

You are already making this estimate. Every security budget you set, every insurance limit you buy, every risk you quietly accept is a quantitative decision. **The only question is whether the number behind it is visible and testable, or sitting in somebody's head.**

I am here for the rest of the morning. Come and tell me where this breaks in your environment. That conversation is more useful to me than the talk was.

Boris Kouambo Ndam — Founder & CEO, Üsta Digital
boris.kouambo@usta.co.za · usta.co.za



Slides & contact
usta.co.za/boris