EKSELENS

RISK INTELLIGENCE PRACTICE

From Heat Maps to Loss Curves

Reinventing Enterprise Risk Management for an Age of AI, Cyber and Geopolitical Fragmentation

An EKSELENS White Paper · November 2026 · ekselens.com

Executive summary

For two decades, enterprise risk management has been built on a comforting set of assumptions: that risk is periodic and can be reviewed on a calendar; that it can be sorted into tidy categories owned by separate functions; that a colour on a heat map is an adequate summary of exposure; and that the world, broadly, is stable enough for last year's risk register to describe next year's threats. Every one of those assumptions has now failed.

The risk environment has not merely intensified — it has changed character. Three forces are responsible. **Geopolitical fragmentation** has turned the global economy from a source of efficiency into a source of exposure, weaponising trade routes, supply chains and standards. **Cyber insecurity** has industrialised, compressing the time between intrusion and impact from months to hours. And **artificial intelligence** has become both the most powerful accelerant of those threats and, paradoxically, the only instrument capable of managing risk at the speed they now move.

These forces do not act independently. A cyber-attack is increasingly a geopolitical act; a supply-chain shock is increasingly a financial one; an AI failure is increasingly all three at once. Risk has become a **system**, not a list — and the discipline built to manage lists is no longer fit for purpose.

This paper argues that organisations need a new operating model for risk, defined by five principles: it must be **continuous** rather than periodic; **quantified** in financial terms rather than expressed in colours; **interconnected** rather than siloed; **resilience-led** rather than prevention-obsessed; and **explainable** enough that machine intelligence can be trusted inside the decision. The boards and risk leaders who make this shift will not simply be better protected. They will be able to take risk deliberately, price it, and convert resilience into competitive advantage. Those who do not will continue to manage this year's crisis with last year's map.

---

1. A step-change, not another turn of the cycle

Risk leaders are conditioned to think in cycles. Recessions come and go; regulations tighten and loosen; threats peak and recede. The instinct, when the environment worsens, is to add controls, commission another assessment, and wait for normality to return.

That instinct is now dangerous, because what we are living through is not a cyclical downturn in the risk weather but a structural change in the climate. The World Economic Forum's *Global Risks Report 2026* captures the shift plainly: it identifies **geoeconomic confrontation** as the risk most likely to trigger a material global crisis in the coming year, with **state-based armed conflict**, **misinformation and disinformation**, and **cyber insecurity** all clustered in the top tier of near-term concerns.[1] The report describes a "contested multipolar landscape" in which confrontation is replacing co-operation — precisely the conditions under which the efficiency-maximising business models of the last thirty years become liabilities.

The defining feature of this new climate is not any single threat but their **interconnection and velocity**. Risks now propagate across domains and borders faster than governance processes can convene. A quarterly risk committee, an annually refreshed register, and a heat map updated for the board pack are artefacts of a slower world. They were never designed to track a threat that can move from a third-party supplier in one jurisdiction to a frozen production line in another within a single trading day.

The remainder of this paper examines the three forces driving the change, explains why the inherited operating model is structurally unable to cope, and sets out what a fit-for-purpose model looks like.

---

2. Three forces remaking the risk landscape

2.1 Geopolitical fragmentation: efficiency becomes exposure

The globalised supply chain was engineered for cost, not resilience. Single-source suppliers, just-in-time inventory and offshore concentration delivered decades of margin — and a brittleness that is now being tested daily. In 2025, **56% of companies reported supply-chain disruption arising from geopolitical issues, and 67% from trade regulations and tariffs**; a trade survey found **72% of professionals identified tariff volatility as the most impactful regulatory change** of the year.[2][3] The financial consequences are not marginal: industry estimates put the average cost of supply-chain disruption at roughly **$184 million per year**, with a single major disruption capable of erasing **30–50% of a year's EBITDA**.[2]

Crucially, geopolitics has moved from being a *macro* concern — something for the strategy team's scenario deck — to an *operational* one. Trade routes have become chokepoints; standards have become trade barriers; and supply-chain dependencies have become instruments of statecraft. A 2025 US congressional analysis warned explicitly of the **weaponisation of supply-chain chokepoints**, noting that control over even a fraction of critical imports could impose "multiple years of economic pain."[4]

The exposure is most acute in complex, concentrated value chains. **Semiconductors** require more than a thousand production steps across a globally distributed network with recognised single points of failure — a vulnerability underlined in 2025 when disruption at a single North Carolina quartz mine rippled through global chip production.[4][5] **Pharmaceutical manufacturing** is structurally similar: the first half of 2025 saw **214 active medication shortages in the United States and 135 in the United Kingdom**, driven heavily by single-source active-ingredient dependence, concentrated production sites and geopolitical disruption to inputs.[6] In both sectors, the cost-optimal configuration of normal times is precisely the configuration that fails catastrophically under stress — and the risk lives not in any one factory, but in the *flow* of materials, components and data between them.

2.2 Cyber insecurity at machine speed

Cybersecurity has completed its transition from an IT housekeeping function to a top-tier enterprise and systemic risk. The numbers describe an industrialised threat. Researchers recorded **7,419 ransomware**

attacks worldwide in 2025, a 32% increase on the prior year, with **manufacturing the hardest-hit sector** — attacks up 56% and average ransom demands more than doubling to **$1.16 million**.[7] Global ransomware damage for the year has been projected at around **$57 billion**.[8]

Two shifts matter more than the headline volume. The first is **the collapse of dwell time**: where attackers once lurked undetected for months, 2025 reporting describes intrusion-to-impact windows compressed to **days or even hours**, while the typical organisation still needs roughly two weeks to detect and contain.[9] The defender's response cadence and the attacker's operational tempo are now mismatched by orders of magnitude. The second is the **dominance of third-party and supply-chain vectors**: Verizon's 2025 breach analysis found that **30% of breaches involved a third party — double the previous year's share** — and supply-chain compromise ranks among the costliest initial vectors, averaging close to **$4.91 million** per breach.[10] Your security posture is now only as strong as that of suppliers you may never have assessed.

The result is that cyber risk can no longer be ring-fenced. A breach at a logistics provider is a supply-chain risk; a ransomware event at a hospital network is a patient-safety and regulatory risk; an outage at a cloud dependency is an operational-resilience risk. The category boundaries that organisational charts impose on cyber risk are precisely the boundaries that attackers exploit.

2.3 Artificial intelligence: accelerant and amplifier

If geopolitics supplies the motive and cyber the method, AI supplies the acceleration. It is reshaping the risk landscape from three directions simultaneously.

As a **weapon**, AI has industrialised social engineering. Reporting through 2025 documented a surge of over **1,200% in phishing linked to generative AI**, with the overwhelming majority of analysed phishing emails containing AI-generated content, alongside a sharp rise in **deepfake voice cloning and fabricated video calls** used for fraud.[11] The marginal cost of a convincing, personalised attack has fallen to near zero.

As a **new attack surface and source of internal risk**, AI introduces failure modes most governance frameworks have never contemplated. IBM's 2025 breach research found that **AI-driven attacks featured in 16% of breaches and "shadow AI" — unsanctioned employee use of AI tools — in 20%**, adding an average **$670,000** to breach costs. Most strikingly, **97% of AI-related breaches involved systems lacking proper access controls, and 63% of organisations had no AI governance policy at all**.[12] AI is being adopted faster than it is being governed — the textbook precondition for systemic risk.

As an **amplifier of complexity**, AI accelerates decision-making across the enterprise, which means errors, biases and model failures also propagate faster and at greater scale. This is why regulators have moved quickly: the **EU AI Act** began applying prohibitions and AI-literacy obligations in February 2025, expanded in August 2025, and reaches its main application phase on **2 August 2026**.[13] AI is now simultaneously a threat to manage, an asset to govern, and — as we argue below — the only practical instrument for managing risk at the velocity the other forces demand.

2.4 The convergence: from a list of risks to a system of risk

The decisive point is not that these three forces are each individually severe. It is that they have **fused**. The WEF's framing of an interconnected risk landscape reflects a lived operational reality: a geopolitical decision triggers an export control, which idles a single-source supplier, which an opportunistic ransomware crew exploits while defences are distracted, while AI-generated disinformation amplifies the reputational fallout — all within days.

This is what commentators have labelled the **polycrisis**: a state in which risks are correlated and compounding rather than independent and additive. It has a profound implication for measurement. If risks were independent, you could assess them one at a time and sum the results — which is exactly what a risk register does. But correlated risks produce *non-linear* losses; the whole is far worse than the sum of the parts. An enterprise risk model that scores risks in isolation will systematically understate exposure precisely when it matters most. Risk has become a system. The tools built to manage a list cannot see it.

---

3. Why the inherited operating model is failing

The dominant enterprise-risk operating model — the one embodied in most GRC deployments, risk committees and board reports today — fails along four structural fault lines.

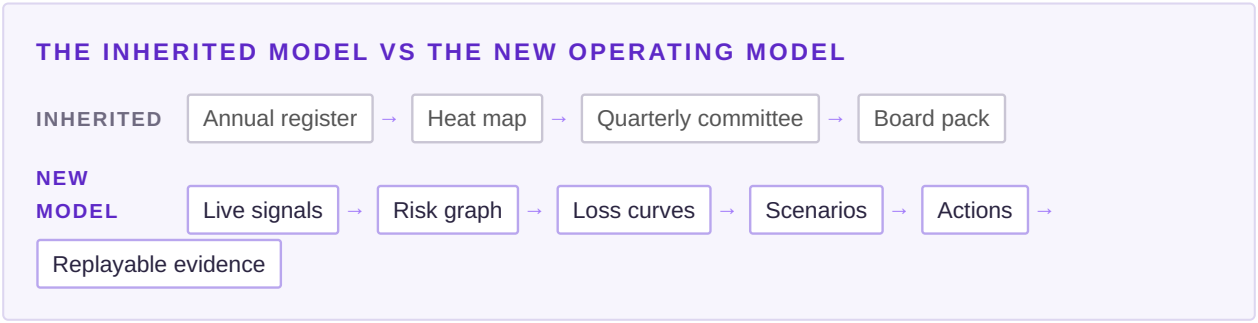
It is periodic in a continuous world. Annual registers, quarterly committees and point-in-time assessments impose a calendar cadence on threats that move at machine speed. By the time a risk is reviewed, escalated and acted upon, the dwell-time data tells us the damage is often already done. The mismatch between governance cadence and threat tempo is the single most important failure of the inherited model.

It is siloed in an interconnected world. Cyber sits with the CISO, supply chain with operations, market risk with the CFO, compliance with legal. Each function optimises its own domain and reports up separately. But the most consequential risks are precisely those that cross these boundaries — and no one owns the connections. The organisational chart becomes the attack surface.

It is qualitative where it needs to be quantified. The red/amber/green heat map is comforting, ubiquitous and analytically hollow. It cannot be aggregated, cannot express probability, cannot be compared against a budget, and cannot answer the only question a board ultimately cares about: *what is this likely to cost us, and how likely is it?* Two "red" risks may differ in expected loss by three orders of magnitude; a heat map renders them identical.

It is location-based where risk has become flow-based. Traditional registers and asset inventories describe where things *are* — which suppliers, which systems, which jurisdictions. But modern exposure lives in the *flows* between them: the movement of components through chokepoints, of data across borders, of dependencies through a supply chain you can only partially see. A static map of locations cannot capture a dynamic of flows.

The common thread is that the inherited model was designed to **report** risk — to assure a board, satisfy an auditor, and document a position at a moment in time. The environment now demands something different: a model designed to **sense, quantify and act** on risk continuously.



4. The new operating model: five principles

A fit-for-purpose risk operating model rests on five principles. None is exotic; the difficulty is that they must be adopted together, because each reinforces the others.

1. Continuous, not periodic. Risk posture must be sensed continuously from live internal and external signals — threat intelligence, supplier health, geopolitical developments, control performance — and updated in near-real-time, with the calendar reserved for governance, not detection. The annual register becomes a continuously maintained, living model.

2. Quantified, not coloured. The lingua franca of risk must become **financial**. Probabilistic loss quantification — expressing exposure as a distribution of potential losses in currency, with explicit likelihoods — replaces the heat map. This is the approach long championed in disciplines such as Factor Analysis of Information Risk and now spreading across enterprise risk. A loss-exceedance curve tells a board what a colour never can: the range of plausible outcomes, their probabilities, and the marginal value of each pound spent reducing them.

3. Interconnected, not siloed. The model must represent risk as a connected graph — suppliers, services, controls, threats and financial impacts linked so that a shock can be traced as it cascades from a tier-3 supplier through to a critical business service and onto the income statement. This is the only way to capture the non-linear, correlated losses the polycrisis produces, and the only way to give a single executive a defensible enterprise-wide view.

4. Resilience-led, not prevention-obsessed. The regulatory and strategic centre of gravity has shifted decisively from *prevention* to *resilience* — from "stop it happening" to "assume disruption, and prove you can absorb it." This is now codified. The EU's **Digital Operational Resilience Act (DORA)** is in force; the UK's operational-resilience regime passed its transition deadline in March 2025; and both turn on the concept of **impact tolerance** — the maximum acceptable level of disruption to an important business service, which firms must map, test against severe-but-plausible scenarios, and evidence.[14][15] Resilience, not prevention, is the new compliance bar — and the better strategic posture.

5. Explainable, not opaque. As machine intelligence enters the risk decision, a new requirement emerges: every material risk figure must be **reconstructable and defensible**. A regulator, auditor or board member must be able to ask "why is this number what it is?" and receive a precise, replayable answer. This is the principle that allows AI to be trusted in the loop — and, as the next section argues, it is where most organisations are least prepared.

---

5. AI in the risk function: instrument and liability

There is an apparent paradox at the heart of this paper. We have described AI as a primary driver of the new risk landscape — yet the operating model we propose is only achievable *with* AI. Continuous sensing across thousands of signals, real-time financial quantification, cascade modelling across a supply graph: none of this is humanly tractable at enterprise scale. AI is not optional to the new model; it is its engine.

This is the resolution of the paradox: **AI is simultaneously the risk and the instrument for managing it.** The same capability that lets an adversary generate a thousand tailored attacks lets a defender continuously re-price exposure across a thousand dependencies. The organisations that will thrive are not those that resist AI in risk, nor those that adopt it naively, but those that deploy it *under governance*.

And governance is exactly where the gap lies. Recall that **63% of organisations have no AI governance policy**, and **97% of AI-related breaches involved inadequate access controls**.[12] An AI that generates risk insight no one can explain, audit or reproduce does not reduce risk — it relocates it, from the threat itself to the decision about the threat. A board that acts on a black-box risk score has simply substituted one opaque exposure for another.

This is why **explainability and determinism** are not technical niceties but governance necessities. A risk figure derived from live, time-varying data feeds, probabilistic simulation and language-model reasoning is, by default, irreproducible: re-run it tomorrow and it changes, with no way to prove the original was sound. To be trustworthy — and to satisfy the auditability that DORA, the AI Act and operational-resilience regimes increasingly demand — such a figure must be **replayable**: the exact inputs, model versions, assumptions and computations must be captured so that any number can be reconstructed and independently verified after the fact. The future of AI in risk is not the most autonomous system; it is the most *accountable* one. Determinism, versioning and an auditable computation trail are what convert a clever model into a defensible decision.

> **WHAT THIS LOOKS LIKE IN PRACTICE: A RISK DECISION-INTELLIGENCE LAYER**
>
> The principles above converge on a recognisable architecture: a layer that sits *above* the systems an organisation already runs — GRC, ERP, procurement, cyber tooling, finance — rather than replacing them. It fuses their internal signals with live external intelligence, expresses every material exposure as a calibrated financial range (never a single point), simulates how shocks cascade across suppliers and services, and preserves the data, assumptions and model versions behind every figure so the decision can be replayed and defended later. The layer does not make decisions; it makes decisions *quantified, comparable and accountable*.

6. What it means for the board and the CRO

The implications for leadership are concrete.

The board's risk conversation must change currency. Directors should stop accepting heat maps and start demanding loss distributions. The right questions are no longer "what colour is this risk?" but "what is our aggregate expected annual loss, what does the tail look like, and what is the marginal return on the next pound of mitigation?" A board that governs risk in financial terms can make genuine trade-offs; one that governs in colours cannot.

The CRO's mandate must expand from assurance to intelligence. The traditional risk function reports and assures. The emerging one *senses and advises* — operating a continuous, quantified, enterprise-wide risk model that informs strategic decisions in real time, not a register reviewed in arrears. This implies a convergence long resisted by organisational politics: cyber, operational resilience, third-party risk, financial risk and compliance must be managed through a **common model and a common language**, not as federated silos.

Resilience must be treated as strategy, not compliance. Mapping important business services, setting impact tolerances and testing severe scenarios are increasingly legal obligations — but firms that treat them as a compliance chore miss the point. The same exercise, done well, reveals the organisation's true dependencies and single points of failure, and turns resilience into a basis for confident risk-taking and competitive differentiation.

AI governance is now a board-level risk domain in its own right. With the EU AI Act's main provisions applying from August 2026, boards need an inventory of AI use, a classification of AI risk, and assurance that AI in the risk function itself is explainable and auditable. Ungoverned AI adoption is the clearest avoidable systemic risk on the current horizon.

---

7. A practical maturity model

The shift is significant but need not be done all at once. We observe organisations progressing through four stages.

Stage 1 — Reactive (the inherited model). Annual registers, qualitative heat maps, siloed functions, periodic board reporting. Risk is documented after the fact. Most organisations remain here.

Stage 2 — Quantified. The organisation adopts financial risk quantification for its most material exposures, replacing colours with loss estimates, and begins to express risk appetite and tolerance in currency. The board conversation starts to change.

Stage 3 — Connected. Risk data is integrated across domains into a single model: cyber, supplier, operational and financial exposure linked so that cascades can be traced and aggregated. Impact tolerances are mapped and tested. The risk function begins to sense continuously rather than review periodically.

Stage 4 — Intelligent and accountable. The model is continuous, AI-augmented and explainable: it senses across live signals, quantifies exposure in real time, simulates cascading scenarios, and produces decisions that are fully reconstructable and auditable. Risk becomes a real-time input to strategy — and AI operates inside the loop, under governance.

The practical path is to start where exposure is greatest and data is richest — typically third-party and operational resilience — quantify it financially, connect it across the relevant silos, and only then layer in continuous sensing and AI augmentation with explainability designed in from the outset, not bolted on.

---

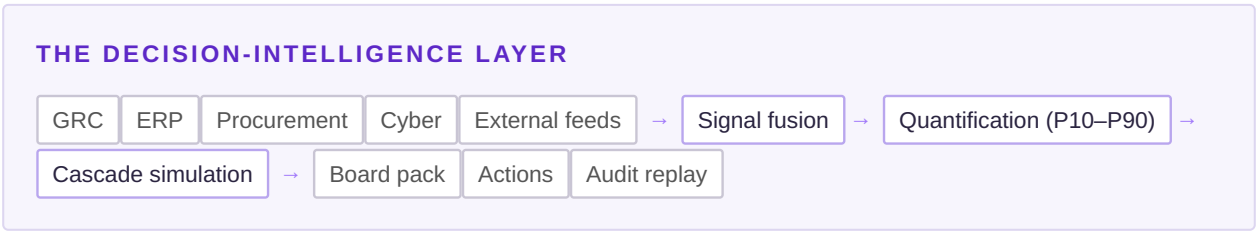
8. How forward-leaning firms are operationalising this

The principles above are no longer theoretical. A small but growing group of risk-mature organisations — concentrated, unsurprisingly, in financial services and in complex-supply-chain sectors such as pharmaceutical and discrete manufacturing — are building an integrated **risk decision-intelligence layer** that embodies them: a single model that ingests live internal and external signals, quantifies exposure financially across a connected graph of suppliers, services and controls, simulates how shocks cascade and what they cost, monitors resilience against defined impact tolerances, and uses AI to do continuously what was previously impossible to do even annually.

What distinguishes the serious efforts from the superficial ones is the fifth principle: **explainability built in**. The leading implementations treat every risk figure as an auditable artefact — capturing the exact data, model versions, simulation parameters and assumptions behind it, so the number can be replayed and defended to a regulator or board member months later. This is what makes it safe to put AI inside the decision rather than merely alongside it.

EKSELENS has been developing a working prototype of exactly such an integrated, explainable risk decision-intelligence layer, as a way of pressure-testing these principles in practice rather than in theory. The methods underpinning its deterministic, replayable audit are now the subject of a filed UK patent application (GB2613313.2), and a guided demonstration is available to clients and partners who wish to see what a continuous, quantified, resilience-led and *accountable* risk model looks like in operation — including foresight capabilities (risk velocity, contagion simulation and counterfactual analysis) that express every

figure as a calibrated range rather than a single point. The purpose of sharing it is not to prescribe a product but to make the argument of this paper tangible: the new operating model is buildable today.



9. Conclusion: manage the climate, not the weather

Risk management has spent decades refining a model built for a stable, slow, compartmentalised world — and that world is gone. Geopolitical fragmentation, industrialised cyber threat and the rise of AI have not merely raised the temperature; they have changed the climate, making risk faster, more interconnected and more consequential than the inherited tools can perceive, let alone manage.

The organisations that endure will be those that stop managing this year's crisis with last year's map. They will move from periodic to continuous, from colours to currency, from silos to systems, from prevention to resilience, and from opaque judgement to explainable, accountable intelligence. They will treat AI not as a threat to be resisted nor a novelty to be adopted carelessly, but as a governed instrument that finally makes enterprise risk manageable at the speed it now moves.

This is not a counsel of fear. A organisation that can see its risk clearly, price it accurately and absorb shocks deliberately is not merely safer — it is freer to take the risks that create value. In an age of permanent volatility, that capability is not a cost of doing business. It is the business.

References

1. World Economic Forum, *Global Risks Report 2026* (Digest and full report), 2026 — weforum.org/publications/global-risks-report-2026
2. ElectroIQ / Procurement Tactics, *Supply Chain Disruption Statistics 2025*, 2025 — electroiq.com · procurementtactics.com
3. Thomson Reuters, *Confronting Complexity and Disruption in Global Trade (Trade Survey)*, 2025–2026 — thomsonreuters.com
4. U.S.–China Economic and Security Review Commission, *Beijing's Weaponization of Supply Chains*, 2025 — uscc.gov; Atlantic Council, *US–China Semiconductor Supply Chain Under Stress*, 2025 — atlanticcouncil.org
5. Z2Data, *Most Destructive Supply Chain Challenges 2025*, 2025 — z2data.com

6. PharmaBoardroom, *When Geopolitics Meets Pharmaceuticals*, 2025 — pharmaboardroom.com; ING, *EU Drug Shortage Analysis*, 2025 — think.ing.com
7. Comparitech / Industrial Cyber, *Global Ransomware Attacks Rose 32% in 2025; Manufacturing Top Target*, 2025 — comparitech.com · industrialcyber.co
8. Cybersecurity Ventures / Mimecast, *Ransomware Statistics*, 2025 — cybersecurityventures.com · mimecast.com
9. ExtraHop, *2025 Global Threat Landscape Report (dwell time)*, 2025 — extrahop.com
10. Verizon, *Data Breach Investigations Report (DBIR) 2025* — verizon.com/business/resources/reports/dbir; IBM, *Cost of a Data Breach 2025* (third-party vector) — ibm.com/reports/data-breach
11. StrongestLayer / BR Side, *AI-Generated Phishing and Deepfake Fraud Analysis*, 2025 — strongestlayer.com
12. IBM, *Cost of a Data Breach Report 2025* (AI-driven attacks, shadow AI, AI governance) — ibm.com/reports/data-breach
13. EU Artificial Intelligence Act, *Implementation Timeline*, 2025–2026 — artificialintelligenceact.eu/implementation-timeline
14. EU Digital Operational Resilience Act (DORA), in application 2025 — eur-lex.europa.eu/eli/reg/2022/2554/oj; ENISA / NIS2 implementation materials — enisa.europa.eu
15. UK Financial Conduct Authority, *Operational Resilience: Insights and Observations*, 2025 — fca.org.uk

---

© 2026 EKSELENS. This white paper is provided for general information and thought-leadership purposes and does not constitute legal, regulatory, financial or investment advice. Statistics are drawn from third-party reports cited above and are presented as reported estimates. For permission to reproduce, contact the EKSELENS Risk Intelligence Practice.