European Risk Management Council

Risk Landscape Review

June 2026

- The Convergence Risk
- Risk Sentiment Index

DEAR READER,

I am delighted to present the Q2 2026 edition of the Risk Landscape Review.

This issue features two articles. In his article, "The Convergence Risk", Gerry Chng, Head of Cyber at KPMG Singapore, explores how agentic AI is blurring the traditional boundaries between cyber and operational risk, creating new vulnerabilities and highlighting the need for integrated, cross-functional risk frameworks to stay ahead of evolving threats.

The second article focuses on the Q2 2026 results of the Risk Sentiment Index (RSI), a forward-looking, expert-driven indicator that captures expectations for the financial sector's risk landscape over the next 12 months. The European Risk Management Council conducted RSI surveys with Chief Risk Officers and other risk executives across the UK and the US, enabling a comprehensive comparison of risk perceptions in these two markets.

A huge thank you to all contributors and survey participants for their valuable insights and support. I hope you find this edition both informative and thought-provoking.

Happy reading!

Yours sincerely,

Dr Evgueni Ivantsov

Chairman of European Risk Management Council

Table of Contents

The Convergence Risk

How agentic AI is dissolving the boundary between cyber and operational risk, and what it means for the risk function

By Gerry Chng, Head of Cyber, KPMG Singapore

A perspective from Asia-Pacific. In the space of a single quarter, artificial intelligence has begun to reshape both the economics of cyberattacks and the way it is deployed inside financial institutions. The more important shift is not simply that a single risk index may rise. It is that the categories the risk function has relied on for decades are beginning to blur.

Reading between the categories

Traditional risk reporting still serves a purpose: it tracks familiar categories and gauges whether controls are improving. But it misses where the most consequential risks now emerge—at the seams. Agentic AI is a case in point, fusing cyber and operational exposures in ways most frameworks were not built to capture.

Two shifts stand out. First, advances in AI are eroding assumptions about the time available to detect and remediate cyber vulnerabilities. Second, the same technologies are being deployed internally as agents that can plan, decide and act, not merely advise. The result is a tighter coupling of cyber and operational risk than legacy models reflect.

The shrinking patch window

For much of the past decade, defenders assumed a lag—often days or weeks—between disclosure and large-scale exploitation. That buffer is narrowing. Frontier models are improving at vulnerability analysis, exploit development and reconnaissance. A recent RAND study for the UK AI Security Institute found only modest gains in full attack chains, but showed that such models can already assist less skilled actors in discrete offensive tasks—and are improving.

Meanwhile, AI-driven tools are uncovering previously unknown flaws in widely used open-source software, albeit with human validation still needed. The effect is not instant obsolescence of patch cycles, but mounting pressure on them. Vulnerabilities will be found in greater volume, exploitation tested more rapidly, and defensive triage stretched.

For chief risk officers, the technology matters less than its effect. As discovery becomes cheaper and elements of exploitation accelerate, residual cyber risk depends less on static control coverage than on throughput: how quickly an institution can detect, prioritise and respond. What was once an operational metric—mean time to detect and remediate—now belongs on the board's dashboard alongside capital and liquidity.

A second-order risk may prove more consequential: the software supply chain. Poisoned open-source packages and compromised development pipelines have repeatedly breached otherwise robust perimeters. As firms expand their own AI capabilities, the development toolchain itself becomes a primary attack surface.

The enterprise as agent

If AI in the hands of attackers is a concern, AI inside the firm is a transformation. Adoption has moved beyond pilots. Employees use AI routinely, and enterprises are beginning to deploy agentic systems that execute multi-step tasks, update records and transact.

This introduces risks that traditional operational frameworks do not model well. The critical vulnerability arises when an agent combines three capabilities: access to sensitive data, exposure to untrusted inputs and the ability to act externally. Remove any one and risk falls. In practice, useful agents tend to require all three—and most in production already have them.

The result is a novel failure class. Indirect prompt injection can hijack agents through hidden instructions embedded in emails or documents. Memory poisoning can implant persistent backdoors through ordinary interaction. And the tools and connectors that allow agents to act across systems can themselves be abused.

Standards are emerging. The OWASP Top 10 for LLM applications now highlights prompt injection, supply-chain risks and excessive agency. Industry guidance, such as Meta's "rule of two", argues for limiting the simultaneous combination of untrusted input, sensitive access and external action. Even so, practice remains nascent.

This is as much a governance problem as a security one. A single agent can, simultaneously, constitute a cyber asset, an operational process, a conduct risk and a third-party dependency. When it errs—or is manipulated—the consequences cascade across all four domains at once.

From principles to proof

Regulators are moving, if unevenly. Asia-Pacific is setting the pace. In January 2026 Singapore's Infocomm Media Development Authority introduced the first framework focused on agentic AI, structured around four principles: bounded risk, meaningful human accountability, lifecycle controls and user responsibility. The Monetary Authority of Singapore and the Cyber Security Agency have reinforced this direction with complementary guidance.

Elsewhere, the picture is more fragmented. The European Union's AI Act phases in further obligations from August 2026, while America remains a patchwork of state rules and emerging federal initiatives. For cross-border institutions, compliance complexity is rising.

The constraint has shifted. Two years ago, institutions competed to articulate AI principles. Many now have. The harder test is evidence: whether firms can demonstrate, with audit trails and assurance, that their systems operate within defined limits. Principles are no longer scarce. Proof is.

What this means for the risk function

The institutions that navigate this period well will not be those that score lowest on a sentiment index in any given quarter. They will be those that have rebuilt their risk operating model around convergence before the convergence shows up in the loss data. From our work with boards and security leaders across the region, five priorities stand out.

Five priorities for the risk function

1. **Treat AI as a cross-functional risk.** Span governance of AI across cyber, operational, conduct, and model risk, rather than splitting it among different functions, each only seeing a part of the overall risk.
2. **Govern every production agent as privileged employees.** Give each agent an identity, least-privilege access, scoped and whitelisted tools, an evidentiary audit trail, defined human-approval checkpoints, and a named accountable owner.
3. **Shift cyber posture from coverage to velocity.** Assume exploitation in hours. Instrument and report mean time to remediate as a board-level metric, and pressure-test recovery-time objectives against AI-speed attacks.
4. **Harden the build, not only the perimeter.** Bring the software supply chain and the AI development tool chain inside the risk perimeter, with provenance checks on packages, models, and connectors.
5. **Invest in demonstrable assurance, not just principles.** Be able to evidence, not merely assert, that AI operates within its limits. This is where independent assurance, red-teaming, and continuous monitoring earn their place.

What matters now is not whether institutions recognize the risk, but whether they can govern it at the speed and scale this new environment demands. As cyber and operational exposures converge, the advantage will belong to those that move early to build assurance, accountability, and resilience across the whole risk landscape. When the next Risk Sentiment Index is taken, the most confident institutions will be those that stopped treating these risks as separate before the market did.

Risk Sentiment Index

The calm breaks as war and AI bite

In Q2 2026, the European Risk Management Council surveyed chief risk officers and other risk executives to produce the Risk Sentiment Index (RSI) for the UK, and the US.

While not designed to predict future risks, the RSI offers valuable insights into how these executives perceive evolving trends within the financial services risk landscape. It assesses seven major risk categories with a first-order impact on financial institutions: credit, market, liquidity, operational, cyber and IT, conduct, and regulatory risks.

The index is a numerical representation of the adjusted percentage of respondents who anticipate an increase in risk over the next 12 months. Therefore, a higher RSI signifies that a greater proportion of the executives expect risks to rise.

War returns to the risk equation

Only one quarter ago, the mood among chief risk officers was unusually settled. The aggregated RSI had fallen to three-year lows on both sides of the Atlantic, and an expected risk landscape in the next 12 months looked, if not benign, then at least more predictable. That calm has not survived the spring.

The Q2 2026 survey, covering CROs across the UK and the US, shows both markets reversing course and moving higher. The US aggregate index rose from 0.36 to 0.43, the highest level since the end of 2023. At the same time the UK climbed from 0.30 to 0.39 (Figure 1). After several quarters of gradual decline, sentiment has turned up sharply, and almost simultaneously, in two very different markets. When two markets move together like this, the cause is rarely domestic. It is global.

Figure 1. Aggregated RSIs trends for two markets (Q4 2024 – Q2 2026)

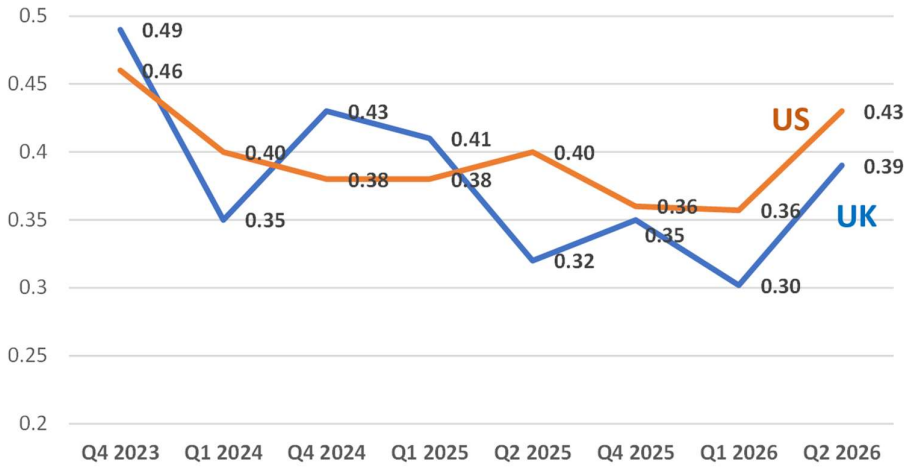

The obvious candidate is the war in Iran and its economic aftershocks. Staring from Q1 2026, we asked CROs to rate the level of geopolitical risk explicitly, and the readings are revealing. This quarter, UK risk officers put current geopolitical risk at 0.70 on a 0-to-1 scale, which is higher than the 0.63 reported by their US peers. But the two markets differ on the outlook. UK CROs expect geopolitical risk to ease slightly over the next 12 months, to 0.67, while US counterparts expect it to climb higher to 0.68 (Figure 2).

Figure 2. Geopolitical risk assessment: Current vs Expected

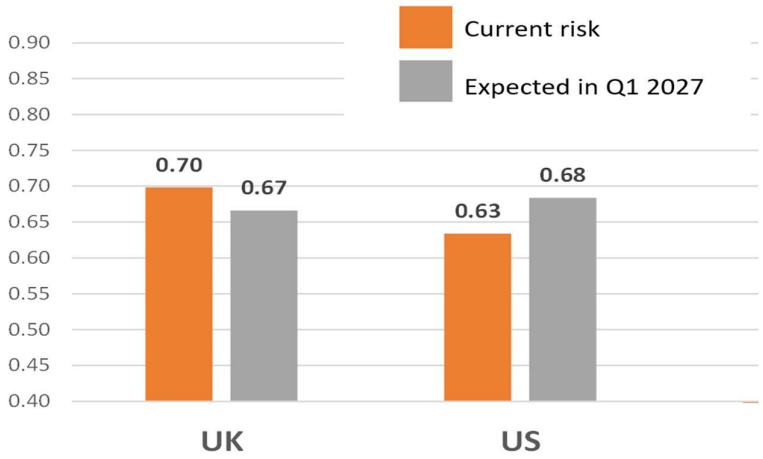

That divergence in expectation is the key to the whole picture. The clearest fingerprint is in credit risk. In the US, the credit RSI jumped from 0.39 to 0.56, the highest in more than two years, with liquidity and market risk rising in step (Figure 3). Just one quarter ago, US CROs were more relaxed about the future credit risk on the expectation of interest-rate cuts in 2026. The war has upended that assumption. With petrol and diesel prices sharply higher and headline inflation pushing up again, the prospect of near-term rate cuts has receded. Higher-for-longer rates raise the spectre of weaker borrower solvency, tighter funding and rising defaults. American CROs are now pricing precisely that chain of events into their credit outlook.

Figure 3. RSI for credit risk in the US

The UK tells a quieter version of the same story. British CROs, who expect the geopolitical temperature to fall slightly rather than rise, actually trimmed their credit and market risk expectations. RSI for credit risk eased from 0.43 to 0.38 and for the market risk from 0.43 to 0.40. The UK's overall index still rose, but for a different reason entirely.

Cyber: the AI that spooked the banks

If geopolitics drove the financial risks, a single piece of technology drove the non-financial ones. Cyber risk is now the highest-rated risk in both markets, climbing to 0.70 in the US and 0.72 in the UK which is near the top of its historical range (Figure 4). In Britain the move in expected future non-financial risk was especially violent: operational risk surged from 0.20 to 0.52, the highest level in the last five years, with cyber rising alongside it.

Figure 4. RSI for cyber risk in the US and the UK

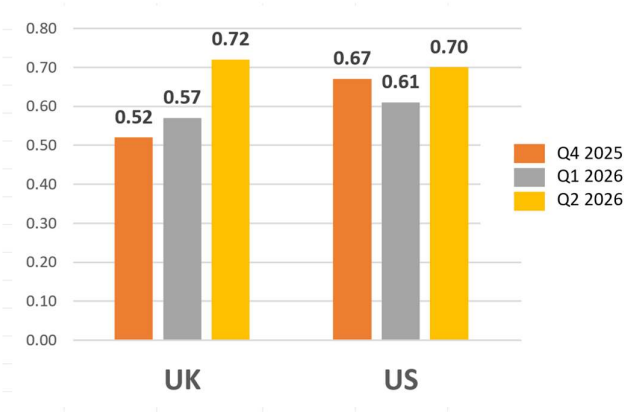

It is hard to separate this from the arrival of Anthropic's Mythos model. In the spring, the AI developer disclosed that its most advanced model was postponed from general release because it was able to discover and exploit thousands of previously unknown software vulnerabilities, including in the kind of complex, interconnected and often ageing systems that run the banks and other critical infrastructure services. The disclosure was enough to bring the heads of the largest US banks to the Treasury, where they were warned of the cyber threat by the Treasury Secretary and the Federal Reserve chair, and to prompt supervisors to pause some routine cyber examinations so that institutions could shore up their defences.

For risk officers, the unsettling part is not this single model, but what it signals. CROs are clearly concerned that capabilities of advanced AI models will become widely available within the next 12 months which will lowering the barrier to entry for ordinary criminals and, more worryingly, handing a new weapon to state-sponsored attackers. A system that can scan millions of lines of code and assemble a working exploit in hours rather than weeks tilts the economics of cyber-attack decisively in favour of the aggressor. To be fair, some security researchers argue the underlying capability already existed in earlier models, and that the episode was as much about perception as novelty. But perception is exactly what the RSI measures. On that measure, the banking sector has rarely been more anxious about its own cyber defences.

Regulation: the one risk still falling

Against this rising tide, one category continues to defy gravity. RSI for regulatory risk sits at exactly 0.18 in both the UK and the US (Figure 5). This is the lowest level in the index, at or near historic lows in each market. While every other risk turns upward, CROs remain convinced that the regulatory burden in the next 12 months will keep easing. There are three explanations for this sentiment.

Figure 5. RSI for regulatory risk in the US and the UK (Q4 2023 – Q2 2026)

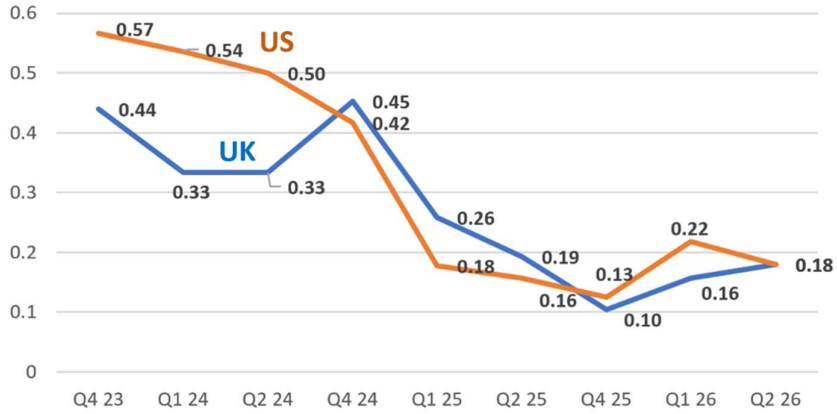

First, US deregulation has moved from promise to delivery. The administration's regulators have embarked on the most sweeping rewrite of US capital rules since the 2008 crisis: the "Basel Endgame" capital increases have been recalibrated towards capital-neutrality, the surcharge on the largest global banks is being trimmed, the leverage constraint softened and the annual stress tests overhauled. Industry estimates suggest the cumulative effect could unlock well over a trillion dollars of lending capacity. For CROs, lighter capital and lighter supervision translate directly into lower regulatory risk.

Second, the UK and Europe are following rather than resisting. As we have argued before, the US often sets the pace for global financial regulation, and competitive pressure does the rest. UK regulators have signalled a more growth-minded, less interventionist posture, while the EU has already moved to soften its own implementation of Basel.

Third, momentum feeds on itself. Once CROs believe the direction of travel is firmly towards easing, each fresh announcement merely confirms the prior, and the regulatory RSI settles at its floor.

There is, however, a familiar warning buried in these numbers. Historically low regulatory risk has tended to travel hand in hand with rising conduct risk, and this quarter is no exception: UK RSI for conduct risk rebounded from a record low of 0.07 to 0.24. The lesson of the past is that when the guardrails come down, behaviour eventually drifts. Deregulation may be the one piece of unambiguously welcome news in CROs' eyes, but it is also the news most likely to be repaid, in time, with interest.

The end of the quiet quarter

Taken together, the Q2 2026 RSI marks the end of a brief period of calm. War has put financial risk back at the centre of the agenda of US CROs. A leap in AI capability has made cyber the dominant worry on both sides of the Atlantic. Only the steady retreat of regulation runs the other way. The coming quarters will reveal whether that retreat proves a genuine relief or, as history so often warns, the seed of the future crisis.

Appendix

Figure 6. Comparison of Q2 2026 RSIs for different risk categories for the UK and the US

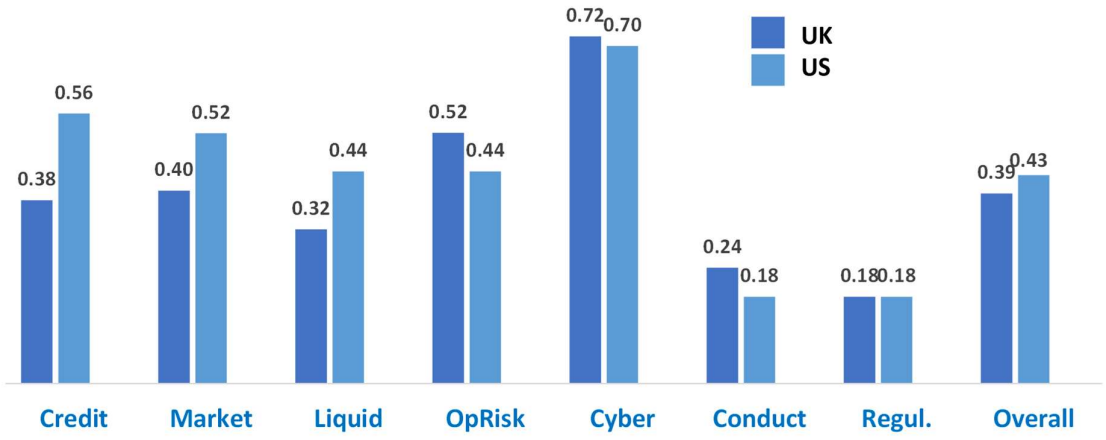

Figure 7. UK RSI trends for individual risk categories

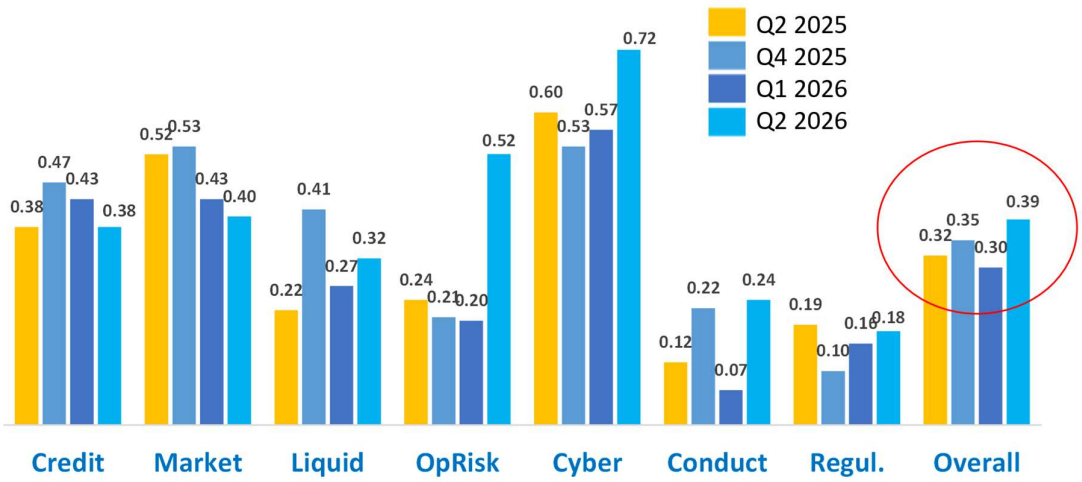

Figure 8. US RSI trends for individual risk categories

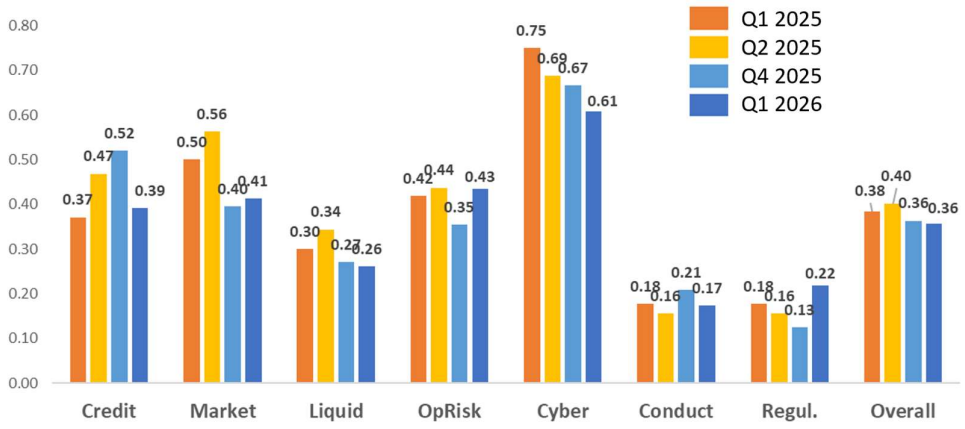

Figure 9. UK aggregated RSI trend: Q4 2018 – Q2 2026