



Rupesh Vashist

Associate Director
Technology Risk: Cybersecurity

Tel: +27 66 101 6590

Email: rupesh.vashist@kpmg.co.za

Insuring the algorithms

How cyber and artificial intelligence (AI) risks are converging and what insurers need to do about this

The emerging prominence of AI alongside the increased incidence of cyber threats presents a shifting risk profile. The critical question for insurers to consider is no longer about whether to underwrite AI, but whether their current cyber, liability and professional indemnity portfolios are already exposed to AI risks and how to incorporate the appropriate level of underwriting discipline in an environment that is rapidly changing.

AI is now embedded in the value chain

Underwriting, pricing, claims, fraud detection and customer engagement have largely started to incorporate machine learning, generative AI and automation techniques. KPMG's 2025 Global Insurance Industry CEO Outlook Survey identified that over the past year insurance organisations have started to increase budget spend in respect of AI innovation, with the majority of insurers also increasing spend specifically to address the cyber risk that AI adoption introduces¹. Three out of four insurance CEOs expect to hire dedicated cybersecurity and risk management talent within the next twelve months, signalling that insurers are viewing AI adoption and cyber resilience as interconnected risks that need to be addressed together.

Another recent global KPMG publication, Advancing AI across insurance, describes generative AI's capabilities as evolving at an unprecedented pace, with accuracy and robustness improving even as deployments scale². Zurich Australia described its own trajectory, moving from bespoke AI pilots toward scalable, largely automated underwriting workflows². With AI adoption now being the norm, the risk exposure created deserves as much attention as the efficiency it delivers.

The cyber-AI convergence: a faster, sharper attack surface

The most immediate risk is unlikely to stem from a sophisticated AI model failure, but rather from AI enabling traditional cybercrime to become faster, more scalable and increasingly persuasive.

Munich Re's 2026 Cyber Outlook identifies ransomware, data breaches, business email compromise and denial-of-service as the dominant loss drivers, and flags agentic AI as a force that adds speed and scale to all four, estimating that if cybercrime were a national economy, it would already be the world's third largest, heading toward USD14 trillion in global cost by 2028³.

Travelers' Q1 2026 cyber threat data demonstrates that ransomware claims are up 80% since 2022, with losses from social-engineering fraud and business email compromise up more than 30% since 2023 and now accounting for roughly half of all cyber claims⁴. Interestingly, claims data from cyber insurer Resilience, covering the first half of 2026, found that AI is being used as a force multiplier for well-established attack methods, such as phishing and social engineering, rather than enabling entirely new categories of autonomous attacks at scale. Notably, the insurer reported no incurred losses from AI-specific threat vectors such as prompt injection or the misuse of agentic AI⁵.

The primary risk insurers now face is not the emergence of advanced AI-driven threats, but the way in which AI is already exacerbating traditional phishing and ransomware attacks, making them faster, more cost effective, and more difficult to prevent.

¹ KPMG LLP, <https://assets.kpmg.com/content/dam/kpmgsites/xx/pdf/2026/01/ceo-outlook-survey-insurance-report.pdf>

² KPMG, <https://assets.kpmg.com/content/dam/kpmgsites/xx/pdf/2024/09/advancing-ai-across-insurance-final-pdf.pdf.coredownload.inline.pdf> (KPMG UK / KPMG Australia, 2025)

³ Munich Re, "Cyber Insurance Outlook 2026" report, as summarised in Cyber Insurance News <https://cyberinsurancenews.org/cyber-insurance-risks-trends-2026-munich-re/> (26 March 2026)

⁴ Travelers, "Q1 2026 Cyber Threat Report", as reported by Risk & Insurance, <https://riskandinsurance.com/ransomware-hits-near-record-highs-as-ai-governance-gaps-widen-inside-organizations/> (8 June 2026)

⁵ Resilience, H1 2026 cyber claims data, PR Newswire, <https://www.prnewswire.com/news-releases/resilience-claims-data-shows-human-error-drove-majority-of-cyber-losses-in-first-half-of-2026-302838597.html> (30 July 2026)

This risk has important implications for underwriting discipline. According to security-ratings firm Bitsight, AI is reshaping cyber risk through increased concentration risk. The AI ecosystem relies on a limited number of model providers, shared orchestration platforms and common agentic architectures. As a result, the failure of a single dependency could transform an isolated incident into a correlated, portfolio-wide loss event⁶. This brings us to exploring the next less familiar category of risk.

When the algorithm becomes the exposure

Insurers traditionally underwrite risks arising from human error, operational failure or identifiable external threats. AI introduces a new risk landscape: loss exposures with no malicious actor or clear point of failure. For example, a claims-triage model may reject legitimate claims at scale; a pricing engine may unknowingly embed bias; a generative AI tool may leak confidential data; an autonomous agent may act within its technical permissions but outside its intended purpose.

Attribution is what is difficult to identify. When a system output causes loss, responsibility may sit somewhere across the deploying organisation, software vendor, foundation-model provider, cloud host, systems integrator, data supplier or humans responsible for oversight. Since most legal frameworks do not yet recognise AI as a party capable of bearing liability, insurers who pay a claim may find subrogation against a responsible third party far harder than in a conventional loss scenario. That does not just create legal ambiguity, it shifts the economics of risk transfer as insurers may be left absorbing losses without being able to rely on traditional recovery mechanisms.



... AI is already exacerbating traditional phishing and ransomware attacks, making them faster, more cost effective, and more difficult to prevent.

⁶ Bitsight, <https://www.bitsight.com/blog/how-frontier-ai-is-reshaping-cyber-insurance-risk> (7 July 2026)

⁷ <https://www.cisco.com/site/us/en/learn/topics/artificial-intelligence/what-is-a-frontier-model.html#tabs-35d568e0ff-item-4bd7dc8124-tab>

⁸ <https://www.underwriting-agents.com/> (2026)

⁹ The Insurer, <https://getcaic.org/playbooks/ai-insurance-policy-and-exclusions> (23 July 2026); Verisk/ISO exclusion endorsement CG 40 47 (effective 1 January 2026)

The frontier-model problem: concentration, autonomy and tail risk

Current AI deployments are powered by frontier AI models as the foundational layer. These large models provide the advanced capabilities needed to perform complex tasks and carry out multi-step reasoning tasks at massive scale⁷.

For insurers using frontier models, accumulation risk is emerging as a significant area of concern. A recent study on AI-liability underwriting estimates that more than 80% of AI deployments now depend on just a handful of foundation-model providers, meaning that a single upstream failure, vulnerability or safety incident could generate correlated losses across a large share of an insurer's book simultaneously⁸. The same paper notes that agent capability is currently improving faster than agent reliability and the ceiling on how badly an autonomous system can fail is rising. In addition, fewer than one in five businesses report a mature governance model for the autonomous agents they are already running⁹.

Frontier AI is transforming from a technology-adoption narrative into a familiar catastrophe-modelling challenge, defined by correlated triggers, limited loss experience, and a small number of critical shared dependencies that underpin an expanding population of insureds.

The coverage reckoning: from “silent AI” to explicit exclusion

The industry's response to managing frontier model risk has been reasonably fast. Since 1 January 2026, Verisk, an American-based entity providing data analytics and technology services, exclusion endorsements have enabled American insurance carriers to remove wording related to generative-AI exposures from standard commercial general liability contracts, with more than 60 property and casualty insurance groups having filed to adopt some form of AI exclusion⁹. This trend echoes how the market responded more than a decade ago to exposures to “silent cyber”, where insurers discovered that they had unintentionally paid for cyber losses under traditional property and liability covers that had not explicitly defined cyber-related risks. The industry is now moving deliberately to close the equivalent “silent AI” gap before accumulating unpriced exposures⁹.

However, economies across the globe are adopting different approaches. In London, The Lloyd's Market Association has held back from a blanket exclusion pending clearer market instruction. Syndicates such as Chaucer and cover-holder Armilla have launched standalone AI-liability structures bundling significant AI and cyber limits together, with other syndicates having introduced AI sub-limits rather than outright exclusions. Interestingly, Lloyd's of London underwriters now identify professional indemnity risk as having the highest perceived AI exposure, with AI-related risk already representing a meaningful share of gross written premium^{10 11}.

The result is a bifurcated market: American insurance players are largely limiting exposure to "silent AI" risks through widespread policy exclusions, whilst in the United Kingdom, specialist capacity is stepping in to underwrite the same risk explicitly on its own terms and at its own price.

Building the discipline: governance, quantification and response

KPMG's 2025 Global Insurance Industry CEO Outlook Survey found that the vast majority of insurance CEOs are already prioritising stronger risk management and cyber resilience specifically to protect their organisations as AI scales¹. What is clear is that AI and cyber risks must be managed in tandem, supported by a consistent set of foundational principles that include robust data governance, rigorous model validation, clear lines of accountability and structured AI maturity assessments¹².

In practice, this requires underwriting frameworks that assess dependency on and concentration around foundation models rather than focusing solely on internal AI policies. It also calls for claims teams with the technical expertise to evaluate incidents involving cyber compromise, model failure and human-oversight simultaneously, as well as scenario analysis that treats a major frontier-model incident failure in the same way catastrophe teams approach hurricanes or earthquakes, as a correlated, portfolio-wide event rather than an isolated claim.

Conclusion

AI represents more than just a new category of risk for insurers. As AI becomes embedded in operational, product and decision-making aspects of the businesses being insured, it changes the frequency, severity, causation and correlation of potential losses. This makes established risk proxies and historical claims experience less reliable, requiring insurers to reassess how exposure is identified, quantified, priced and reflected in underwriting terms across cyber, liability and professional indemnity portfolios.

At the same time, AI-enabled cyber threats are increasing the speed, scale and persuasiveness of familiar attacks, while model failures, data leakage, biased outputs and autonomous actions create losses with uncertain attribution and more difficult recovery pathways.

Beyond these downstream exposures, there is a distinct accumulation risk. Many enterprise AI agents and applications depend on foundation models supplied by a small group of providers among the principal providers. A vulnerability, outage, safety failure or change in underlying model or orchestration layer could affect many insured businesses at the same time, turning apparently isolated claim events and exposures into a correlated, portfolio-wide event. For underwriting, understanding AI usage and dependency on the underlying business being insured now requires continuous visibility over model-provider dependencies, substitutability, safeguards, human oversight and business-continuity arrangements.

Insurers that will be better positioned to succeed will be those that treat AI and cyber risk as a single, converging discipline, supported by robust data governance, model validation, clear accountability, scenario analysis and technically capable claims teams. Their key advantage will not be in predicting every failure or relying solely on broad exclusions, but in continually adapting risk quantification and underwriting discipline as AI adoption evolves and prepare for when shared dependencies turn isolated incidents into systemic losses across the portfolio.

References:

Frontier Model Forum risk taxonomy, as referenced in "How frontier AI companies could implement an internal audit function", arXiv:2512.14902 (2026)

¹⁰ ResultSense, <https://www.resultsense.com/news/2026-07-30-cfc-affirmative-ai-cover-lloyds-market/> (30 July 2026)

¹¹ Landers, <https://landers.com.au/legal-insights-news/ai-risk-and-insurance-a-new-frontier> (23 April 2026)

¹² KPMG, <https://assets.kpmg.com/content/dam/kpmgsites/sg/pdf/2026/05/ai-in-insurance-sector-2026.pdf.coredownload.inline.pdf> and <https://assets.kpmg.com/content/dam/kpmgsites/xx/pdf/2026/06/intelligent-insurance-report.pdf> (KPMG International / KPMG Singapore, 2025–2026)