Continuum

ARTIFICAL INTELLIGENCE AND THE INSURANCE GAP

RISK INSIGHT SERIES

MARCH 2026

Contents

Introduction

The insurance industry is undergoing a structural shift. Artificial intelligence is no longer a peripheral technology being tested at the edges of the business. It is being embedded across underwriting, claims, fraud detection, and customer operations at a pace that has few precedents in the industry's history. With 90% of insurers now in some stage of generative AI evaluation and 55% in early or full adoption, generative AI has shown the strongest uptick among all AI categories tracked across the sector.[1]

The operational case is clear. Claims processing times have been reduced by between 55% and 75% through AI automation, with routine claims moving from a seven to ten day cycle to resolution within 24 to 48 hours.[2] Fraud detection, risk evaluation, and pricing models are being rebuilt around machine learning outputs rather than actuarial judgment alone. The global insurance sector saw an 87% year-on-year increase in AI deployments, with generative AI and agentic technologies together accounting for 68% of all new rollouts in the final quarter of 2025, led by claims management, underwriting, and customer engagement.[3]

But the speed of adoption has outpaced the development of the frameworks needed to govern it. As AI becomes embedded in consequential decisions about which risks to accept, how to price them, and whether a claim is valid, questions of accountability, liability, and coverage respondability are becoming urgent. The same models driving operational efficiency are creating exposures that existing policy language was not designed to address.

This is particularly acute in Technology Professional Indemnity and Cyber insurance, where AI is both reshaping the risk being insured and creating new uncertainties about whether coverage responds when something goes wrong. Silent AI exposure, where policies neither explicitly include nor exclude AI-related losses, is widespread. Exclusion language is narrowing in ways that many buyers have not fully assessed. And as AI governance failures attract regulatory attention, the liability landscape for firms deploying AI is becoming more complex and less predictable.

This report examines those dynamics. Drawing on market observations and jurisdictional analysis, it considers how AI is reshaping the risk profile of Tech PI and Cyber policies, where the coverage gaps are emerging, and what firms and underwriters need to understand as the gap between AI adoption and insurance alignment continues to widen.

Background

How AI Entered the Insurance Coverage Debate

The insurance industry's relationship with AI risk did not begin with a deliberate policy decision. It began with silence. As businesses integrated AI tools into their operations over the past several years, existing policies simply did not address the question of whether AI-related losses were covered or excluded. The risk was absorbed by default rather than by design.

This phenomenon, known as silent AI, refers to AI-driven risks that are neither explicitly included nor excluded in insurance policies, leaving room for potential coverage gaps that can lead to significant financial losses for both insurers and policyholders.[4] The parallel with an earlier episode in the market's history is instructive. The situation closely resembles the evolution of cyber insurance, where cyber losses were initially paid under property or liability policies unintentionally before insurers added explicit exclusions and created dedicated cyber policies once the risk became significant and better understood.[5]

That inflection point has now arrived for AI. Artificial intelligence risk has captured the attention of the business world, with some 72% of S&P 500 companies now discussing AI and its related risks in their annual securities filings.[7] Insurers have responded not by developing comprehensive new frameworks but by moving to limit their exposure under existing ones. As insurers file for broad generative AI exclusions or move to exclude AI-related errors and omissions claims altogether, relying on silent AI coverage, where coverage is assumed because it is not explicitly excluded, is no longer a viable strategy for businesses.[8]

The exclusion trend has accelerated markedly. Some insurers have introduced what is described as an absolute AI exclusion, eliminating coverage for any claim arising out of or attributable to the actual or alleged use, deployment, or development of artificial intelligence, including AI-generated content, failure to detect AI-produced materials, inadequate AI governance, chatbot communications, and regulatory actions related to AI oversight.[9]

At the same time the market is narrowing, the complexity of the exposures it is leaving behind is growing. No single policy currently covers all AI perils. Companies often rely on a patchwork of policies to address AI risks, with different policies covering different aspects and each carrying its own gaps where an AI-related loss may not fit neatly.[10] This creates a structural problem. The risks are becoming more defined and more foreseeable, while the coverage available to respond to them is simultaneously becoming more fragmented and more conditional.

This is the environment in which Technology Professional Indemnity and Cyber policies now operate. Understanding how AI has changed the risk profile of those lines, and where the gaps in existing coverage are most acute, requires examining each in turn.

How AI Changes the Risk Profile

From Technical Risk to Operational and Legal Exposure

The introduction of AI into business operations does not simply add a new category of risk to existing ones. It restructures how familiar risks materialise, who is responsible when they do, and whether existing coverage responds. This restructuring is most acute in the two lines of insurance most relevant to technology-dependent businesses: Technology Professional Indemnity and Cyber.

Understanding why requires examining how AI changes the underlying risk profile across three dimensions.

The Accountability Problem

Traditional Tech PI and Cyber policies were written with a relatively clear model of human agency in mind. A professional makes a decision or provides advice, a system fails or is compromised, and liability flows from an identifiable act or omission. AI disrupts this model in a fundamental way. When AI generates outputs that cause harm, determining accountability is complex. If a business is using a service provider for AI, both parties should allocate liabilities in their contractual agreements. AI models typically add liability waivers, meaning responsibility sits with the business directly providing the service to the customer rather than with the AI provider.[11]

This creates a structural gap. The business deploying AI bears the liability but may not have designed, trained, or fully understood the model producing the output. Professional indemnity policies written around human professional judgment may not respond cleanly to losses caused by autonomous or semi-autonomous AI decisions. Policies that are silent on whether AI services constitute professional services may lead to future coverage disputes when an AI-related claim arises. Even where technology errors and omissions coverage exists, policies may restrict coverage only to failures of software developed or created by the insured, which could limit coverage where a third party's AI malfunctions and triggers litigation against the user.[12]

The Hallucination and Output Quality Problem

Generative AI introduces a category of loss that has no real precedent in traditional Tech PI: the confident, plausible, and wrong output. Gen AI outputs may deviate from intended purposes through hallucination, where the system produces factually incorrect or entirely detached outputs presented with high confidence. Unlike traditional AI, which enables traceable decision pathways, Gen AI outputs cannot be mapped to specific parameters and the creative generation process is inherently difficult to explain, complicating error attribution.[11]

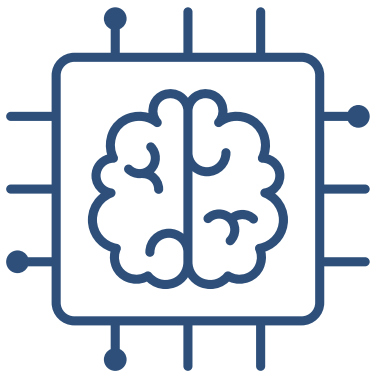

For insurers, this creates a challenge that goes beyond typical software error claims. A software bug produces a predictable failure. A hallucinating AI produces a failure that may be indistinguishable from a correct output at the point of delivery. Claims arising from AI-generated legal advice, medical information , financial guidance, or technical documentation therefore involve a degree of epistemic uncertainty that traditional PI underwriting was not designed to address.

The Cyber Risk Amplification Problem

AI also materially changes the cyber threat landscape that Cyber policies are written to address. AI is driving how cyber insurance is priced, assessed, and delivered. Newer risks like model manipulation, IP infringement, hallucinated outputs, and data poisoning are rarely addressed explicitly in current policy language. Companies are also struggling to collect cyber insurance payouts when a breach originates from a vendor in their supply chain, especially when the AI involved is a black box with no transparency, traceability, or audit trail to support the claim. [14]

At the same time AI is expanding the attack surface for businesses, it is also being used by threat actors to increase the scale and sophistication of attacks. Gen AI has lowered the technical barrier to entry for cyber criminals, enabling the proliferation of AI-driven malware, phishing, and distributed denial-of-service attacks. Cybercrime-as-a-service groups now leverage AI tools to specialise in particular phases of the attack chain, making attacks more scalable, believable, and effective. [15]

The combined effect is that the cyber risk being transferred under a Cyber policy today is materially different from the risk that was underwritten when most current policy language was developed. Policies calibrated to a pre-AI threat environment may carry significant unmodelled exposure.

The Exclusion Landscape

From Technical Risk to Operational and Legal Exposure

Regulatory Approach

The insurance market's response to AI has followed a familiar pattern. Faced with a new and poorly understood risk category, insurers have moved first to limit exposure under existing lines rather than to develop comprehensive new frameworks. The result is an exclusion landscape that is rapidly evolving, inconsistently applied, and in many cases poorly understood by the buyers most affected by it.

The current situation closely parallels the evolution of cyber insurance. Initially, cyber losses were paid under property or liability policies unintentionally before insurers added explicit exclusions and created dedicated cyber policies once the risk became significant and better understood. For AI, we are at that same inflection point. Insurers are starting to tighten terms, adding wording to clarify whether losses from autonomous decisions or algorithmic errors are covered. [16]

The Silent AI Problem

The starting point for most businesses today is silent AI coverage — exposure that is neither explicitly included nor excluded under existing policies. AI risks are neither explicitly mentioned, limited, nor excluded in much current policy language, and different exposures may be covered by different policies already in existence.[17] This is not a stable position. As AI becomes more deeply embedded in business operations and loss events accumulate, the implicit coverage that silent AI provides is being actively withdrawn.

As insurers file for broad generative AI exclusions or move to exclude AI-related errors and omissions claims altogether, relying on silent AI coverage, where coverage is assumed because it is not explicitly excluded, is no longer a viable strategy for businesses. [18]

The Move Toward Explicit Exclusions

The exclusion trend has accelerated through 2025 and into 2026. Some insurers have introduced what is described as an absolute AI exclusion, eliminating coverage for any claim arising out of or attributable to the actual or alleged use, deployment, or development of artificial intelligence. This includes AI-generated content, failure to detect AI-produced materials, inadequate AI governance, chatbot communications, and regulatory actions related to AI oversight.[16]

Other insurers have taken a narrower but still significant approach, focusing exclusions specifically on generative AI. Hamilton Insurance Group's Generative Artificial Intelligence Exclusion for professional liability policies removes coverage for any claim, wrongful act, damages, or defence costs based upon, arising out of, or in any way involving any actual or alleged use of generative artificial intelligence by the insured.[13]

The practical consequence is that two firms with superficially similar policies can have materially different AI coverage depending on which carrier they are with, when their policy was last renewed, and how their broker has navigated exclusion negotiations.

The Coverage Gap In Tech PI

For Technology Professional Indemnity specifically, the AI exclusion problem intersects with a pre-existing structural ambiguity. Companies that provide professional services need to perform extra due diligence when reviewing the scope of covered services, as many professional liability policies may limit coverage solely to those services being provided by natural persons, not artificial systems. Policies that are silent on whether AI services constitute professional services may lead to future coverage disputes when an AI-related claim arises. [8]

This creates a scenario where a firm using AI to deliver professional services — whether legal, financial, technical, or advisory — may find that its PI policy responds to the human advisor but not to the AI-assisted output, even where those outputs are indistinguishable to the client.

The Coverage Gap in Cyber

Cyber policies present a different but equally significant gap. Cyber policies typically do not cover bodily injury or property damage resulting from AI failures. Professional Indemnity policies may exclude AI-assisted professional services. New specialty AI policies often have significant limitations, high premiums, and extensive governance requirements. This creates what insurance lawyers describe as the Swiss cheese effect — multiple policies, each with its own holes, theoretically combining to provide complete coverage but leaving gaps where holes align.[11]

Companies are struggling to collect cyber insurance payouts when a breach originates from a vendor in their supply chain, especially when the AI involved is a black box with no transparency, traceability, or audit trail to support the claim. Policy exclusions for state-sponsored attacks, unclear attribution, and invisible AI risk compound the problem. [10]

The Emerging Affirmative Market

Against this backdrop of narrowing coverage, a nascent affirmative AI insurance market is beginning to take shape. In April 2025, Armilla Insurance Services introduced an AI liability insurance product underwritten by Lloyd's of London syndicates, explicitly addressing AI-specific perils such as hallucinations, degrading model performance, and mechanical or algorithmic failures. Google also announced a partnership with Beazley, Chubb, and Munich Re, introducing a tailored cyber insurance solution specifically designed to provide affirmative AI coverage for Google Cloud customers. [7]

These products represent an important development but remain nascent. No single policy currently covers all AI perils. Companies often rely on a patchwork of policies to cover AI risks, with different policies covering different aspects and each carrying its own gaps where an AI-related loss may not fit neatly.[6] The affirmative market has not yet developed to the point where it replaces the coverage that silent AI once provided.

Where the Gaps Are

The Coverage Gaps Firms Are Most Exposed To

Understanding that AI exclusions exist is not the same as understanding where specific exposures sit uninsured. The gaps that matter most are not always the obvious ones. They tend to cluster around four areas where policy language, regulatory expectations, and AI-driven risk intersect in ways that buyers have not fully mapped.

Gap 1: AI-assisted professional advice

The most immediate and widespread gap affects any firm using AI to support the delivery of professional services. This includes legal research tools, financial modelling platforms, underwriting support systems, compliance monitoring software, and any workflow where AI output informs a professional recommendation delivered to a client.

The gap arises because most PI policies were written to respond to human professional error. Where AI assists or generates the output, the policy may not respond, either because the insurer argues the service was not provided by a natural person, because the policy excludes AI-generated content, or because the chain of causation between the AI output and the loss is contested. Policies that are silent on whether AI services constitute professional services may lead to future coverage disputes when an AI-related claim arises. Organizations that have secured technology errors and omissions coverage may assume protections are already in place for any tech or software failures involving AI, but E&O policies may restrict coverage only to failures of software developed or created by the insured, limiting coverage where a third party's AI malfunctions. [8]

For firms operating in regulated industries where AI is being deployed at scale in client-facing functions, this is a material and largely unaddressed exposure.

Gap 2: Third-party AI failures flowing through to the insured

A significant and growing exposure involves AI tools developed and operated by third parties but embedded in the insured's products or services. When a third-party AI model fails, hallucinates, or produces a harmful output, the liability often sits with the business that deployed it rather than the developer, but the coverage may not follow.

Companies are struggling to collect cyber insurance payouts when a breach originates from a vendor in their supply chain, especially when the AI involved is a black box with no transparency, traceability, or audit trail to support the claim. Policy exclusions for state-sponsored attacks, unclear attribution, and invisible AI risk compound the problem, and many organisations silently absorb losses rather than pursue claims they believe will be disputed. [9]

This gap is compounded by the contractual reality that most AI service providers include broad liability waivers in their terms of service, effectively pushing liability downstream to the deploying business while leaving that business without clear policy coverage for the resulting loss.

> “The gaps that matter most are not always the obvious ones. “

Gap 3: Regulatory and Governance Failures Related to AI

As AI governance becomes a formal regulatory expectation, the liability exposure associated with governance failures is growing. Directors and senior management face increasing personal exposure where AI-related failures can be attributed to inadequate oversight, poor governance frameworks, or failure to implement appropriate controls.

Berkley's absolute AI exclusion eliminates coverage for any claim based upon, arising out of, or attributable to inadequate or deficient policies, practices, procedures, or training relating to artificial intelligence, as well as any actual or alleged failure to develop or implement such policies. This extends to regulatory actions related to AI oversight. [10]

Where D&O policies contain similar exclusions, the individuals most exposed to regulatory scrutiny over AI governance may find themselves personally uninsured at exactly the moment enforcement attention is highest.

Gap 4: Model training data and intellectual property disputes

A less widely understood but increasingly litigated gap involves claims arising from the data used to train AI models. Copyright infringement, unauthorised use of personal data, and IP disputes linked to AI training datasets are generating a new category of third-party claim that sits awkwardly across existing policy lines.

Professional Indemnity policies may exclude AI-assisted professional services while Cyber policies typically do not cover bodily injury or property damage resulting from AI failures. New specialty AI policies often have significant limitations, high premiums, and extensive governance requirements. No single policy currently covers all AI perils. [11]

For firms that have built proprietary AI tools using third-party data, or that use commercial AI platforms where training data provenance is unclear, this is an exposure that sits in the gap between existing lines with no obvious home in current coverage structures.

What Firms Should Be Doing

The practical implication of these gaps is that most firms currently operating with AI-assisted workflows are carrying uninsured exposure they may not have fully identified. Firms should start by identifying their own unique AI risk profile, then carefully review all liability insurance at renewal, paying close attention to the application of AI exclusions and requesting their removal or replacing coverage with a carrier providing more appropriate terms. For companies with specific AI exposures, consideration should be given to new affirmative AI products and whether they can fill gaps that may remain in legacy policies. [12]

Regulatory Dimensions

AI Governance Failures and the Liability They Create

The regulatory environment surrounding AI is evolving quickly across the major insurance markets, but it is doing so unevenly. Across Asia and beyond, the absence of a single unified AI regulatory framework does not mean an absence of regulatory risk. It means that liability exposure is diffuse, accumulating across multiple regulatory channels simultaneously, and increasingly difficult for firms to anticipate from policy language alone.

The EU AI Act as a Global Reference Point

While not directly applicable to most Asian markets, the EU AI Act has established a reference framework that is shaping regulatory thinking globally. The EU AI Act imposes comprehensive regulation on AI development and usage, requiring businesses to carefully balance the management of AI risks against the promotion of technological innovation. For insurers and their clients, this matters because it formalises the concept of AI risk tiering, assigns liability to deployers of high-risk AI systems, and introduces audit and documentation obligations that are beginning to influence how underwriters assess governance quality.[32]

The significance for insurance buyers is not primarily about direct compliance. It is that the Act establishes a template for how liability attaches to AI governance failures that regulators in other jurisdictions are actively considering. Firms that have not mapped their AI systems against an emerging governance standard face the prospect of retroactive exposure as local frameworks catch up.

Regulatory Scrutiny Across Asian Markets

Across the jurisdictions most relevant to Continuum's audience, AI-specific regulation remains at an early stage but supervisory expectations are rising.

In Singapore, the Monetary Authority of Singapore has issued guidance on responsible AI use in financial services, emphasising fairness, ethics, accountability, and transparency. In Hong Kong, the Securities and Futures Commission has begun examining how AI is being used in trading, investment advice, and compliance functions, with particular attention to model governance and audit trails. In Japan, the Financial Services Agency has indicated that AI-related risks will be treated as part of existing operational risk frameworks, placing the compliance burden on licensed entities.

The common thread across these jurisdictions is that AI governance failures are increasingly treated not as technology incidents but as operational risk events with regulatory consequences. Where a firm cannot demonstrate that its AI systems are operating within defined parameters, that outputs are being monitored, and that governance structures exist to identify and remediate failures, it faces enforcement exposure that existing policy structures may not adequately address.

The D&O Dimension

As regulatory scrutiny intensifies, personal liability exposure for directors and senior management is growing in parallel. As regulatory regimes formalise accountability, exposure extends to directors and senior management. D&O insurance responds to regulatory investigations and enforcement actions, allegations of governance failure, claims arising from inadequate supervision or internal control, and shareholder or investor actions following material loss events. In regulated AI markets, enforcement focus increasingly examines governance oversight. [16]

The risk is particularly acute where boards have approved AI adoption without establishing commensurate governance frameworks. A board that can demonstrate it considered AI risks, implemented oversight structures, and monitored outcomes is in a materially different position from one that approved AI deployment and left governance to operational management. The proliferation of AI exclusions in D&O policies means that the individuals most exposed to AI governance scrutiny may find themselves personally uninsured at exactly the moment enforcement attention is highest. [17]

The Compliance and Coverage Alignment Problem

There is a structural misalignment emerging between what regulators expect of firms deploying AI and what insurers are willing to cover. Regulators are expanding the scope of what constitutes a governance or compliance failure in AI contexts. Insurers are simultaneously narrowing the coverage available for AI-related losses. The result is a widening zone of regulatory exposure that sits outside both the firm's risk management framework and its insurance programme.

Insurers have started adjusting underwriting practices for AI exposures, using analogies to known risks and running scenario analyses. Data scarcity means there is not yet sufficient historical loss data on AI incidents to support more precise underwriting. In the meantime, the gap between regulatory expectation and insurance coverage is likely to widen before it narrows. [10]

For firms operating in regulated financial services environments, the practical implication is that regulatory alignment and insurance alignment need to be treated as connected problems rather than separate workstreams. A governance framework that satisfies a regulator but does not satisfy an insurer's underwriting requirements may leave a firm exposed on both fronts.

Implications for Firms and Underwriters

What Both Sides of the Market Need to Do Differently

The coverage gaps described in the preceding sections are not static. They are widening as AI adoption accelerates, as exclusion language tightens, and as regulatory expectations around AI governance become more explicit. For both insurance buyers and underwriters, the gap between current practice and what is needed is becoming a material risk management problem.

For Firms Buying Insurance

The starting point for any firm using AI in its operations is an honest assessment of where AI sits in its risk profile and how existing coverage responds to AI-related losses. Most firms have not done this work systematically. Firms should start by identifying their own unique AI risk profile, then carefully review all liability insurance at each renewal, paying close attention to the application of any AI exclusions and requesting their removal or replacing coverage with a carrier providing more appropriate terms. All policies should also be re-reviewed should the organisation decide to expand into AI offerings or develop its own technology. [7]

Several practical steps follow from this. First, firms should map their AI use cases against their existing policy schedule to identify where AI-assisted activities are taking place and whether those activities are explicitly covered, silently covered, or excluded. Second, they should engage their broker in a specific conversation about AI exclusions rather than treating renewal as a routine process.

The introduction of broad AI exclusions by major carriers means that a policy renewing on similar terms to last year may carry materially different coverage. Third, firms should understand the contractual allocation of liability in their AI vendor agreements. Where a third-party AI provider has excluded liability for model failures, the firm is absorbing that exposure and needs to ensure its own policy responds.

For firms in regulated industries, there is an additional dimension. Where AI governance frameworks are becoming a regulatory expectation, firms that can demonstrate strong governance, explainable models, and documented oversight processes are better positioned both with regulators and with underwriters. Transparent AI could move from best practice to competitive advantage, with insurers offering better coverage terms to firms that can show how their models work and prove they are fair and secure. [4]

For Underwriters

The underwriting challenge posed by AI is not simply a question of whether to include or exclude AI risk. It is a question of how to assess and price a risk category where loss data is limited, causation is complex, and the underlying technology is changing faster than policy language can keep pace with.

Insurers are currently using analogies to known risks, comparing AI scenarios to similar past claims in cyber and technology, and running scenario analyses. The slow appearance of AI activity exclusions means that silent AI is perceived as sufficient for many seeking coverage, but this client perspective differs from that of insurance companies actively developing AI products. These differing incentives lead to varying approaches regarding the future of AI coverage. [17]

The most important shift underwriters need to make is from treating AI as a binary inclusion or exclusion decision to developing the underwriting tools needed to assess AI governance quality as a risk factor. As regulatory frameworks mature, insurers are increasingly assessing firms based on governance quality, custody architecture, and regulatory alignment rather than technical security measures alone. The same principle applies to AI. A firm with documented AI governance, regular model audits, human oversight protocols, and clear escalation procedures presents a materially different risk profile from one deploying AI without governance infrastructure. [18]

This requires investment in underwriting capability. Most underwriters currently lack the technical fluency to assess AI governance quality in meaningful depth. Building that capability, whether through specialist underwriting teams, third-party assessment tools, or partnerships with technology assessment providers, is a precondition for developing sustainable AI coverage rather than simply cycling through phases of silent coverage and broad exclusion.

The Market Direction

The trajectory of the AI insurance market over the next two to three years is likely to follow the pattern established by cyber insurance in the early 2010s. A period of broad exclusion and market uncertainty will give way to more defined coverage structures as loss data accumulates, underwriting tools mature, and regulatory frameworks establish clearer standards. Whether existing insurance products or new standalone AI solutions will come to dominate the market remains to be seen, but the direction of travel is toward explicit, affirmative coverage with governance requirements embedded in policy terms as a condition of cover. [19]

For firms, the implication is that governance investment made today translates into insurance access tomorrow. For underwriters, the implication is that those who develop AI underwriting capability ahead of the market will be better positioned to write profitable business when the market for affirmative AI coverage matures.

Summary

Artificial intelligence is reshaping the insurance industry at a pace that existing coverage structures were not designed to accommodate.

The efficiency gains are real and the competitive pressure to adopt is intensifying. But the risk profile of businesses deploying AI is changing faster than the policy language governing what is and is not covered.

The central problem is not that AI is uninsurable. It is that the insurance market is in the middle of a transition from silent AI coverage, where AI-related losses were absorbed by default under existing lines, to a more explicit framework where coverage must be affirmatively obtained, governance requirements must be met, and exclusion language must be carefully navigated. That transition is incomplete, inconsistently applied, and poorly understood by many of the buyers most affected by it.

Across Technology Professional Indemnity and Cyber insurance, the gaps are structural rather than incidental. AI-assisted professional advice may fall outside PI coverage where policy language was written around human professional judgment. Third-party AI failures may not be covered where the insured did not develop the model. Regulatory and governance failures related to AI may be explicitly excluded under the very D&O policies that directors are relying on for personal protection. And model training data disputes are generating IP and copyright claims that sit in the gap between existing lines with no clear coverage home.

The exclusion trend is accelerating. Some carriers have introduced absolute AI exclusions that eliminate coverage for any claim arising from the use, deployment, or development of artificial intelligence. Others have narrowed coverage more selectively. The affirmative market that would replace this coverage is nascent, expensive, and governance-intensive. For most businesses, the result is an insurance programme that has not kept pace with the AI risk it is being asked to cover.

The regulatory dimension adds a further layer of complexity. Across Asian markets, AI governance is transitioning from voluntary best practice to regulatory expectation. Firms that cannot demonstrate documented governance, model oversight, and escalation procedures face enforcement exposure that existing policy structures may not adequately address. The misalignment between what regulators expect and what insurers will cover is likely to widen before it narrows.

For firms, the most important action is also the most straightforward: map AI use cases against existing coverage, engage brokers specifically on AI exclusions at renewal, and treat governance investment as both a regulatory and an insurance requirement. For underwriters, the path forward lies in developing the tools to assess AI governance quality as a risk factor rather than defaulting to broad exclusion as the primary response to uncertainty.

The gap between AI adoption and insurance alignment is not permanent. But closing it requires action from both sides of the market, and the window for doing that proactively is narrowing.

References

1. Conning, "AI in Insurance: The C-Suite Verdict," Third Annual Survey on AI and Insurance Technology, June 2025. https://www.conning.com/about-us/news/ir-pr---ai-survey-2025

2. Datagrid, "42 Insurance AI Agent Statistics (Adoption and Impact)," December 2025. https://datagrid.com/blog/ai-agent-for-insurance-statistics

3. Evident, "Insurance AI Deployments Jump 87% as GenAI and Agentic Systems Expand," Reinsurance News, March 2026. https://www.reinsurancene.ws/insurance-ai-deployments-jump-87-as-genai-and-agentic-systems-expand-says-evident/

4. Kennedys Law, "Silent AI Cover: The Unforeseen Risks for Insurers," May 2025. https://www.kennedyslaw.com/en/thought-leadership/article/2025/silent-ai-cover-the-unforeseen-risks-for-insurers/

5. WTW, "Insuring the AI Age," Dr Anat Lior and Sonal Madhok, December 2025. https://www.wtwco.com/en-us/insights/2025/12/insuring-the-ai-age

6. Hunton Andrews Kurth, "How Insurance Policies Are Adapting to AI Risk," 2025. https://www.hunton.com/insights/publications/how-insurance-policies-are-adapting-to-ai-risk

7. Metropolitan Risk Advisory, "Major Insurers Are Pulling Back from AI Liability," November 2025. https://www.metropolitanrisk.com/major-insurers-are-pulling-back-from-ai-liability/

8. Zelle Law, "AI Update: The Growing Trend of AI-Related Insurance Policy Exclusions," October 2025. https://www.zellelaw.com/AI_Update_The_Growing_Trend_of_AI-Related_Insurance_Policy_Exclusions

9. TechLife Future, "Silent AI Insurance Crisis: SME Coverage Gaps in 2026," December 2025. https://www.techlifefuture.com/ai-insurance-exclusions-sme/

10. Geneva Association, "Gen AI Risks for Businesses: Exploring the Role for Insurance," October 2025. https://www.genevaassociation.org/sites/default/files/2025-10/gen_ai_report_0110.pdf

11. American Bar, "The Evolving Landscape of AI Insurance: Empirical Insights into Risks and Policy Gaps," Fall 2025. https://www.americanbar.org/groups/tort_trial_insurance_practice/resources/brief/2025-fall/evolving-landscape-ai-insurance-empirical-insights-risks-policy-gaps/

12. Hunton Andrews Kurth, "The Hidden C-Suite Risk of AI Failures," Harvard Law School Forum on Corporate Governance, September 2025. https://corpgov.law.harvard.edu/2025/09/22/the-hidden-c-suite-risk-of-ai-failures/

13. ISACA, "Cyber Insurance in Crisis with AI Blind Spots," 2025. https://www.isaca.org/resources/news-and-trends/isaca-now-blog/2025/cyber-insurance-in-crisis-with-ai-blind-spots

14. Metropolitan Risk Advisory, "Major Insurers Are Pulling Back from AI Liability," November 2025. https://www.metropolitanrisk.com/major-insurers-are-pulling-back-from-ai-liability/

15. Datagrid, "42 Insurance AI Agent Statistics (Adoption and Impact)," December 2025. https://datagrid.com/blog/ai-agent-for-insurance-statistics

16. Evident, "Insurance AI Deployments Jump 87% as GenAI and Agentic Systems Expand," Reinsurance News, March 2026. https://www.reinsurancene.ws/insurance-ai-deployments-jump-87-as-genai-and-agentic-systems-expand-says-evident/

17. Hunton Andrews Kurth, "The Continued Proliferation of AI Exclusions," 2025. https://www.hunton.com/hunton-insurance-recovery-blog/the-continued-proliferation-of-ai-exclusions

18. TechLife Future, "Silent AI Insurance Crisis: SME Coverage Gaps in 2026," December 2025 https://www.techlifefuture.com/ai-insurance-exclusions-sme/

19. IBM Institute for Business Value, "Insurance in the AI Era," September 2025. https://www.ibm.com/thought-leadership/institute-business-value/en-us/report/insurance-in-ai-era

Continuum

Risk. Insurance. Technology.

hello@continuuminsure.com

www.continuuminsure.com