

June 15, 2026

The End of ‘Silent AI’? Emerging AI Exclusions, Coverage Fragmentation, and Practical Implications for Policyholders

By: [Heidi Lawson](#), [Sarah Hopkins](#)

Related Professionals

Heidi Lawson
Partner, Regulatory

Sarah Hopkins
Associate, Regulatory

Related Practices

→ Regulatory

→ Corporate Governance

→ Privacy & Cybersecurity

[+ Show More](#)

Related Industries

→ [AI & Machine Learning](#)

→ [Insurtech](#)

[Share](#)

→ [Subscribe](#)

What You Need to Know

- The insurance market is moving away from “silent AI” coverage, the practice of implicitly covering AI risks through existing cyber and technology errors and omissions (Tech E&O) policies without express reference to AI, as insurers introduce AI-specific exclusions and revised forms ahead of 2026 renewals.
- Coverage is fragmenting across cyber, Tech E&O, directors and officers (D&O), and employment practices liability insurance (EPLI) as insurers independently narrow AI protections within each line, creating “gap risk” where no single policy provides comprehensive coverage.
- Coverage erosion may be occurring quietly, through revised base forms, narrowing definitions, and restrictive carve-backs rather than a single conspicuous AI exclusion. Policyholders should

review insuring agreements, exclusions, and endorsements together to understand their exposure.

- Proactive program design is key. Policyholders should conduct comprehensive policy reviews; assess how potential claims map to specific coverage lines; and explore manuscript endorsements, captive structures, or difference-in-conditions layers to address gaps in commercial markets.

Artificial intelligence is being embedded across virtually all business operations, from product development to customer engagement and corporate governance. Until recently, the insurance market largely addressed AI-related exposures through existing policy forms, most often through cyber and Tech E&O policies, without expressly referencing AI (meaning that AI risk was implicitly covered because it wasn't specifically excluded). This coverage framework, often described as "silent AI," is rapidly reaching its end.

Insurers are increasingly reexamining AI risk and responding in divergent, and at times inconsistent, ways. Some carriers are clarifying coverage through endorsements, while others are narrowing or eliminating coverage through broad exclusions, revised coverage forms, or underwriting changes. When reviewing coverage for 2026 renewals, policyholders should expect more AI risk to be expressly allocated, or quietly removed, from policies that previously may have responded, though this coverage was "silent." The result is not clarity but fragmentation. Policyholders now face heightened risk of AI-related claims falling between traditional coverage lines or

being subject to competing exclusions across the insurance tower.

This article outlines key developments in the market, the implications for common lines of coverage (including cyber, Tech E&O, D&O, and EPLI), and practical considerations for structuring insurance programs in light of these changes.

From ‘Silent AI’ to Explicit Risk Allocation

Historically, [AI risks were addressed implicitly through traditional insurance products](#). For example, cyber policies could respond to data breaches involving AI systems, while Tech E&O policies could cover liability arising from software errors or failures. As AI has become more pervasive, however, this approach has created ambiguity regarding the scope of coverage.

The insurance market is now moving away from this ambiguity. Insurers are increasingly seeking to define the boundaries of AI coverage, either through affirmative grants of coverage or, more commonly, through [exclusions designed to limit exposure](#). For example, in January 2026, the Insurance Services Office (ISO), an industry organization that, among other things, develops standard policy forms for the industry, introduced a new [generative AI exclusion](#) in commercial general liability (CGL) policies. The endorsement [excludes coverage under the CGL policy](#) for bodily injury, property damage, and personal or advertising injury arising out of, or attributable to, generative AI.

This transition mirrors the evolution of cyber risk, which was initially treated as “silent cyber” exposure under traditional policies before insurers introduced exclusions, affirmative cyber endorsements, and standalone cyber products. AI

appears to be following a similar path, but on a more compressed timeline and with a more fragmented market response.

Emerging Trend: Retrenchment and ‘Quiet’ Coverage Erosion

A growing number of insurers have begun to [narrow coverage for AI-related risks](#), often without significant changes to headline policy language. For insureds reviewing their renewals for 2026, that narrowing may appear in revised base forms, AI-specific endorsements, new exclusions, definitions, application questions, underwriting file positions, or restrictive carve-backs rather than in a single conspicuous “AI exclusion.” Market participants report that carriers are, in some cases:

- Declining to underwrite AI-related exposures entirely or moving them outside the base coverage form
- Increasing premiums or imposing additional underwriting scrutiny
- Carving out liability associated with AI outputs, decision-making, system behavior, or use of third-party AI tools

These developments reflect [fundamental underwriting concerns](#). Insurers have cited the opacity of AI systems, the difficulty in modeling loss scenarios, and the potential for systemic or large-scale harm as key drivers of this retrenchment.

Notably, these changes are not always transparent to policyholders. Rather than introducing new exclusions in all cases, carriers may rely on existing policy limitations, narrowing interpretations, revised coverage forms, or

underwriting practices to manage exposure. The result is a form of “quiet” coverage erosion that may only become apparent at the claims stage, precisely the concern that made “silent cyber” such a challenging issue before the cyber market matured.

For 2026, the key question may be less whether an AI exclusion is present on the declarations page and more whether AI-related loss has been removed or narrowed through the architecture of the form itself. Policyholders should review insuring agreements, definitions, exclusions, carve-backs, application materials, and endorsements together, because the practical effect of those provisions may be to move AI risk outside the coverage grant even where the policy still appears AI-neutral.

Divergence Across Coverage Lines

While the market trend is broadly toward greater restriction of AI risk, the approach varies significantly by line of coverage.

A. Cyber Insurance

Cyber policies remain, for now, the most stable source of coverage for AI-related risks. Many insurers continue to [treat AI as an extension of traditional cyber perils](#), such as unauthorized access, data breaches, and social engineering fraud.

Indeed, some carriers have reinforced their intent to cover AI-enabled threats, such as deepfake fraud or AI-driven cyberattacks, through clarifying endorsements.

That said, cyber policies were not designed to address all AI-related risks, particularly those involving liability arising from AI outputs, intellectual property issues, or regulatory exposure.

As a result, cyber coverage alone is unlikely to provide a complete solution.

B. Tech E&O and Professional Liability

Tech E&O policies have historically provided a critical layer of protection for technology companies, including coverage for negligent acts, errors, or omissions in the provision of products and services.

However, insurers are increasingly scrutinizing AI-related exposures in this context. Reports indicate that some carriers are declining to cover liabilities associated with AI-generated outputs or are introducing exclusions targeting AI-related errors or decision-making.

At the same time, other insurers are experimenting with affirmative AI endorsements that expressly cover defined categories of AI risk. These endorsements may provide clarity, but they also risk narrowing coverage by limiting protection to specifically enumerated scenarios.

C. Management Liability (D&O, EPLI, Fiduciary)

The most significant recent developments have occurred in management liability lines, including D&O and EPLI.

Several insurers have introduced broad “absolute” [AI exclusions](#) in these policies, purporting to exclude coverage for claims arising out of any use, development, or deployment of artificial intelligence.

These exclusions are often drafted in expansive terms and may capture a wide range of ordinary business activities involving AI. Notably, similar exclusionary language is now appearing across multiple coverage lines, including D&O, EPLI, and fiduciary liability policies.

This trend is particularly significant given the increasing volume of AI-related litigation affecting management liability exposures. Examples include:

- Securities class actions alleging misstatements regarding AI capabilities
- Regulatory investigations and enforcement actions
- Employment-related claims arising from alleged bias or discrimination in AI-driven decision-making

The Central Coverage Issue: Nature of the Claim vs. Nature of the Technology

A recurring issue in AI-related coverage disputes is whether coverage turns on the involvement of AI as a technology or on the legal nature of the claim being asserted.

This issue is not limited to companies that develop, license, or sell AI products. A retailer, manufacturer, financial institution, professional services firm, or other non-AI-native business may create enterprise AI exposure when employees use third-party generative AI tools to draft code, summarize contracts, screen applicants, analyze customer data, prepare marketing materials, or support customer communications.

In practice, insurance coverage is typically determined by the characterization of the claim, not the underlying technology. For example:

- A claim alleging a software defect or system failure may fall within Tech E&O coverage.
- A claim alleging discrimination arising from an AI-driven hiring tool is more likely to be treated as an employment practices claim and fall within EPLI (or be excluded entirely).

- A claim alleging that employees uploaded confidential customer, employee, or proprietary data into a third-party AI tool may implicate cyber, privacy, or professional liability coverage.
- A claim alleging that AI-generated marketing copy, code, or product content infringed third-party intellectual property may implicate Tech E&O, media liability, or commercial general liability coverage.
- A claim alleging that AI-assisted hiring, promotion, or workforce analytics produced discriminatory results may implicate EPLI or D&O coverage, even where the company merely used a vendor tool.

This distinction is critical. The same underlying AI system may give rise to different claims depending on how plaintiffs frame their allegations.

As a result, AI creates a unique form of “classification risk,” where coverage depends less on the facts of the underlying technology or conduct and more on how the claim is pleaded. This classification risk can arise from enterprise adoption of third-party AI tools just as readily as from proprietary AI development.

Coverage Fragmentation and ‘Gap Risk’

AI-related risks do not align neatly with traditional insurance categories. A single AI-driven event may implicate multiple lines of coverage, including:

- Cyber (data breach or security incident)
- Tech E&O (product or service failure)
- D&O (disclosure or governance failures)
- EPLI (employment-related claims)

At the same time, insurers are narrowing coverage within each of these lines independently. The combined effect is the emergence of “gap risk,” where:

- Each policy responds only to a subset of the risk.
- Exclusions overlap or “stack” across policies.
- No single policy provides comprehensive coverage.

This dynamic is relevant not only for companies operating in complex technology environments, such as digital asset platforms or AI-enabled SaaS providers, but also for businesses that may not view themselves as AI companies. Employee use of publicly available AI tools, vendor-provided AI functionality, automated customer interactions, and AI-assisted internal workflows may cause risk to span multiple coverage categories at once.

Market Responses: Affirmative Coverage and Alternative Risk Transfer

In response to these developments, the market is beginning to explore alternative approaches to managing AI risk.

A. Affirmative AI Coverage

Some insurers have introduced affirmative AI endorsements or standalone products designed to provide explicit coverage for AI-related exposures. While these products may offer greater certainty, they remain in early stages and are not yet widely adopted.

B. Captives and Difference-in-Conditions Structures

Captive insurance vehicles are increasingly being considered as a mechanism to address AI-related coverage gaps.

Captives may be used to:

- Retain risks that are difficult to insure in the commercial market
- Provide difference-in-conditions (DIC) coverage where commercial policies exclude, sublimit, or narrowly define AI exposures; particularly where exclusions across cyber, Tech E&O, D&O, EPLI, or fiduciary policies could otherwise stack
- Finance losses within, above, or alongside self-insured retentions where commercial insurers require higher retentions for AI-related exposures; fund first-party response costs, third-party defense costs, regulatory investigation costs, remediation expenses, or other costs that may not clearly fit within traditional coverage grants.
- Align risk financing with the company's specific AI risk profile, including its internal AI governance, vendor controls, model deployment practices, and tolerance for retaining emerging risks

This approach is particularly relevant for exposures that insurers view as difficult to model, including algorithmic errors, model drift, or governance failures.

DIC coverage may be especially useful where commercial policies appear to cover the same enterprise AI event, but each contains a different exclusion, sublimit, or narrowing interpretation. A captive DIC layer can be manuscripted to respond when underlying coverage is unavailable because of an AI-specific exclusion, a definition that removes AI-enabled conduct from the coverage grant, or a sublimit that is insufficient for the expected loss profile.

Self-insured retentions can also be integrated into the structure. For example, a company may use a captive to reimburse or absorb a defined AI-related retention, to sit excess of an operating-company retention but below commercial excess insurance, or to provide a corridor layer between traditional policies and higher excess limits.

Practical Considerations for Policyholders

In light of these developments, policyholders should consider taking the following steps, with particular attention to 2026 renewals and to AI use that may be embedded in ordinary business operations rather than housed in a dedicated AI product or team:

A. Comprehensive Policy Review

- Identify any AI-related exclusions or endorsements.
- Evaluate how “AI” and related terms are defined.
- Assess how exclusions interact across policies.

B. Focus on Claim Characterization

- Consider how potential claims may be framed by plaintiffs.
- Evaluate whether key risks are more likely to implicate E&O, D&O, EPLI, or other coverage lines.

C. Program Structuring

- Explore opportunities to broaden coverage through manuscript endorsements.
- Consider the role of captives or DIC structures in addressing gaps.
- Coordinate coverage across policies to reduce fragmentation.

Conclusion

The insurance market is undergoing a fundamental shift in how it approaches AI risk. The transition from “silent AI” coverage to explicit risk allocation resembles the earlier transition from “silent cyber” to affirmative cyber coverage, exclusions, and standalone products. The difference is that AI risk may be more diffuse: It can arise from employee workflows, vendor tools, governance decisions, and ordinary business operations, not only from companies that build or sell AI systems.

For policyholders, the primary challenge is not simply the presence of AI exclusions, but the fragmentation and quiet erosion of coverage across multiple policy lines. As insurers continue to refine their approach, proactive program design, including coordinated policy review, AI-use inventories, and alternative risk transfer, will be important for managing AI-related exposures effectively.