Shumaker

JULY 30, 2026 | PUBLICATION

Client Alert: The New AI Coverage Fight: Exclusions, Endorsements, and Denied Claims

INDUSTRY SECTOR

Technology

SERVICE LINE

Data Privacy
Technology

RELATED PROFESSIONALS

Brian C. Focht

MEDIA CONTACT

Wendy M. Byrne
wbyrne@shumaker.com

Many businesses are adopting artificial intelligence (AI) under the familiar assumption that if something goes wrong, insurance will help absorb the loss.

That assumption may need revision.

As with cybersecurity and cyber liability insurance a decade ago, the insurance industry is doing something it doesn't like doing: paying for claims caused by rapidly evolving threats relying on insurance policy language that has been standard for decades.

So, when it comes to AI, insurance companies have started doing what they're famous for. They're saying "no."

The issue is not that insurance companies have suddenly discovered AI, and it's also not malicious or a money grab. It's actually far more mundane. Insurers are now deciding how much of the risk they want to keep, how much they want to price separately, and how much they want to exclude entirely.

That change is happening through new AI-specific exclusions, narrower endorsements, sublimits, underwriting questions, and old policy language that was never written with generative AI or AI agents in mind. In other words, insurance companies are no longer waiting until you file a claim to fight coverage for AI-related problems.

The Real Issue: "Silent AI" Coverage Is Starting to Disappear

Until recently, most AI-related risks weren't addressed as a unique category of risk. A cyber policy might cover a data breach involving an AI tool. A technology errors and omissions policy might respond to an AI-enabled service failure. A media or general liability policy might apply to certain defamation, privacy, or advertising injury claims involving AI-generated content. That kind of unspoken coverage is sometimes called "silent AI" coverage.

The problem for insurers should be obvious: AI can turn ordinary risks into larger, faster, harder-to-measure risks. A bad employee email reaches one customer. A bad chatbot answer may reach thousands. A mistaken professional judgment affects one matter. An AI system embedded in a workflow may repeat the same mistake across an entire customer base before anyone notices. An AI agent may not just recommend an ill-considered approach; it may implement it.

Insurers have noticed. The market is moving away from "silent AI" coverage as carriers introduce AI-specific exclusions, revised forms, and narrower coverage structures ahead of renewals. This matters because business owners may think they are buying the same insurance program they bought last year, when the renewal actually moves AI risk into a different box: excluded, sublimited, subject to special conditions, or only covered under a separate AI endorsement.

1. **Insurers are Adding Express AI Exclusions**

The most direct approach is also the easiest to understand: insurers are adding policy language that says certain AI-related claims are not covered.

One major development is the emergence of exclusions for generative AI in commercial general liability policies. Insurance Services Office (ISO) form CG 40 47 01 26 excludes bodily injury, property damage, and personal and advertising injury arising out of, or attributable to, generative AI. Other reporting similarly found that new general liability endorsements were developed to allow carriers to exclude generative AI exposures, with forms effective in January 2026. The identified forms include CG 40 47, CG 40 48, and CG 35 08, with CG 35 08 directed to products/completed operations liability coverage.

Why does that matter to a business client?

Commercial general liability (CGL) coverage is often the form of coverage companies expect to apply to ordinary third-party claims: bodily injury, property damage, defamation, certain privacy claims, and advertising injury. Ten years ago, if you were the victim of a cyberattack, you were most likely to find coverage in your CGL policy.

However, under a CGL policy with exclusions arising out of the use of AI, if an AI-generated instruction, recommendation, image, advertisement, chatbot statement, or product-related output causes harm, an insurer may argue the claim is excluded because the alleged injury "arises out of" generative AI.

In an insurance contract, your rights are based almost entirely on the words on the page. The specific terms and usage matter to the courts interpreting them. Insurance companies use words and phrases with specific intentionality. The phrase "arising out of" matters and is often broader than businesspeople expect. The insurer does not necessarily need the AI system to be the only cause of the loss. The fight may become whether the AI system was connected closely enough to the alleged harm.

That is where the comfortable assumption breaks down, as a company may believe it has general liability coverage, and the carrier may respond that it does, except for the portion of the claim at issue.

2. **Some Exclusions are Much Broader than "AI Content"**

A narrow AI exclusion might apply only to generative AI output. That is still important, but at least the business can identify the target. The broader exclusions are more troubling.

Insurance company Berkley Insurance recently introduced a so-called "Absolute" AI exclusion in several specialty liability lines, signaling a broader effort to compartmentalize AI risk. The concern with this type of language is that it may not be limited to a hallucinated chatbot response or an AI-generated image. It may reach claims involving the actual or alleged use, deployment, development, integration, or failure of AI. For businesses, that distinction is critical.

AI is no longer a single product sitting in a lab; it's everywhere. It's been shoe-horned into most of the software that businesses use every day. It may be inside customer service tools, applicant screening workflows, fraud monitoring systems, marketing platforms, contract review tools, software development tools, call summaries, sales enablement systems, and internal analytics. It is already in your Microsoft Office products. If an exclusion is broad enough, the insurer may have an argument whenever AI appears anywhere in the factual background.

That is not how most businesses think about insurance; rather, they tend to think in terms of the harm: Was there a lawsuit? Was there a breach? Was there a professional error? Was there a customer injury? Insurers may increasingly focus on the mechanism: Did AI contribute to the loss?

That shift matters.

3. **Existing Policy Language may be Used Against AI Claims**

New exclusions are only part of the issue. Insurers do not always need new AI wording to contest coverage. They can also rely on existing definitions, exclusions, and coverage limitations, which complicates things.

AI-related claims often do not fit neatly into one traditional policy bucket. A deepfake payment fraud event may look like a crime loss, a social engineering loss, a cyber incident, an internal control failure, and a management liability problem. A hallucinated customer service response may appear to be a consumer protection issue, a professional services error, a false advertising claim, or a regulatory problem. Numerous coverage lines are likely impacted by AI and deepfake risks, including commercial general liability, cyber/media, directors and officers (D&O), errors and omissions, professional liability, intellectual property (IP), and commercial crime coverage.

The fact that several policies might be relevant does not mean any one policy will actually pay.

Below are some ways insurers may avoid coverage:

- The claim is not an "occurrence" (another one of those key insurance terms) because the injury resulted from intentional deployment of AI
- The use of AI may trigger a professional services exclusion
- The policy covers cyber incidents but not model failures, hallucinations, or poor business output

While cyber liability insurance may provide clear coverage for AI-related events, such as hacks exposing private information or data poisoning attacks, cyber insurance policies vary widely and are renowned for including electronic data, cyber, AI-specific, or intentional/wrongful-act limitations. AI does not just create new risk. It creates fights over the specific classification of those risks.

4. **D&O and Professional Liability May Have Hidden Gaps**

Directors and officers often assume AI risk is primarily a technology issue, which is a dangerous assumption.

AI adoption can become a management issue. For example, did leadership approve the use case? Did the company accurately disclose material AI risks? Did it oversee vendors? Did it test the system? Did it ignore known failure patterns? Did it claim AI capabilities that the product did not actually have? Did it rely on automated decisions without appropriate human review? These questions can become shareholder claims, customer claims, regulatory inquiries, employment claims, or professional negligence claims.

Directors and officers may be operating with unrecognized liabilities under the false impression that traditional D&O policies fully insure AI-related risks, especially as AI exclusions begin to appear and professional-services definitions or exclusions may limit coverage.

A D&O policy may not be designed to cover the company's delivery of professional services. An errors and omissions policy may not cover all product-related or regulatory claims. A cyber policy may not cover every AI-driven business failure. The result can be an ugly round of insurance finger-pointing after the claim arrives, and that is a very expensive time to discover that no one mapped the risk before renewal.

5. **Insurers Are Using Sublimits and Narrow Affirmative Coverage**

Not every insurer response is a total exclusion. Some carriers are creating affirmative AI coverage products or endorsements, and others are applying lower sublimits to specific AI-related risks.

This response is entirely predictable to anyone who watched the birth and growth of cyber liability insurance. As cyber risk became too large and too specific to fit comfortably within traditional policies, the market developed dedicated cyber products.

The creation of specialized insurance products can be useful, but businesses should read the details carefully. Remember that the words on the page are your policy, and it's likely each one is there for a reason. A policy that says it covers AI risk may cover only a narrow slice of the risk: hallucination, model performance, IP claims, AI-enabled cyber events, or certain regulatory costs.

However, an AI endorsement is not a perfect shield. Businesses still need to know what is covered, what is excluded, what is sublimited, what notice is required, what controls are conditions of coverage, and whether the endorsement coordinates with cyber, technology errors and omissions, media, crime, employment practices, and D&O policies. Otherwise, the company may be buying the insurance version of theater: comforting, expensive, and entirely performative.

What Businesses Should Do Now

Businesses do not need to panic; they just need to stop treating AI insurance as a renewal footnote.

A practical review should include at least the following steps:

1. **Inventory AI use cases.** The company cannot evaluate coverage for AI risk if it does not know where AI is being used. That includes approved tools, vendor-embedded AI, employee use of public tools (including unauthorized use, or "Shadow AI"), customer-facing AI, internal decision tools, and AI features added to existing platforms.
2. **Map each use case to potential harm.** A chatbot creates different risk than an AI coding assistant, and an AI hiring tool creates different risk than an AI fraud monitoring system. The company should identify whether the likely harm is data exposure, discrimination, bad advice, IP infringement, defamation, financial fraud, bodily injury, product failure, or operational disruption.
3. **Review exclusions at renewal, not after the claim.** AI exclusions should be reviewed across the entire insurance program as part of the renewal process, including commercial general liability, cyber, technology errors and omissions, professional liability, D&O, employment practices, fiduciary, media, crime, and property coverage.
4. **Push back on overbroad wording.** Not every AI-related connection should eliminate coverage. Businesses should negotiate for narrower causation language, carve-backs for ordinary business

operations, exceptions for AI used merely as a tool, and preservation of coverage for cyber incidents, privacy claims, defense costs, and non-AI contributing causes. AI-related insurance language is new, which is when insurance language is most likely to be adjusted.

5. **Coordinate insurance with vendor contracts.** If an AI vendor disclaims liability, caps damages at a small amount, refuses meaningful indemnity, and the insurer excludes AI-related claims, the business may be left holding the bag. Know where your gaps are. Vendor contracts and insurance coverage need to be reviewed together.
6. **Document governance and controls.** Underwriters may increasingly ask about AI policies, approved use cases, testing, human review, access controls, monitoring, incident response, and vendor oversight. Good governance is not just compliance hygiene. It may affect underwriting, pricing, exclusions, and claim arguments. Speak with an AI attorney to help create or update these policies.
7. **Build AI into incident response.** If an AI-related incident occurs, the company should know which policies may be implicated, what notice deadlines apply, what vendors must preserve, what logs exist, and who has authority to suspend the tool.

Bottom Line

Assuming that AI risk is automatically insured because the company has insurance is the new version of "famous last words." It may have been a reasonable assumption when AI was a background feature, a limited experiment, or a tool used by a few employees, which is no longer a reasonable assumption when AI is being built into customer interactions, professional services, products, employment decisions, security systems, finance workflows, and executive strategy.

Insurance companies are not waiting for AI coverage law to develop slowly through litigation. They are rewriting the rules now through exclusions, endorsements, underwriting, sublimits, and narrow affirmative products.

Businesses should respond the same way they should respond to AI generally: define the use case, identify the risk, read the contract, test the control, and document the decision.

The policy that matters is not the one the company thinks it bought. It is the one the insurer will actually apply when something goes wrong.

If you have questions or would like more information, please contact Brian Focht.