



ACTIVANT RESEARCH

Insuring the Agent

Who Owns the Risk Layer for AI?

October 2026

AUTHORS:



Scott Watson

Vice President



Avi Gulati

Intern

About [Activant](#)

[Activant](#) is a research-led global investment firm that partners with high-growth companies. Since 2015, we have invested in category-defining businesses during their most critical phases of growth, partnering with founders who have won their initial battles and are ready for the next challenge.

Our approach pairs deep, proprietary research with patient, flexible capital and hands-on operational partnership. We work alongside founders and leadership teams to refine strategy, strengthen operations, and accelerate sustainable growth.

[Activant](#) Research is dedicated to uncovering the most exciting emerging technologies, sectors, and companies we believe will shape the future. Our research-driven perspective informs everything we do, helping us invest at meaningful inflection points and support founders in building enduring, category-leading businesses.

You can find out more about [Activant](#) and our research at <https://activantcapital.com/>.

Our recent research keeps leading us back to insurance. Our [Agentic Commerce series](#) mapped software that shops; [Agentic Procurement](#) followed B2B spend toward machine-to-machine negotiation; and [Robots: Valuations and The Scaling Question](#) gave the agent a body. Each piece explored a different part of the emerging agent economy, but all raised the same unanswered question: when the agent gets it wrong, who pays?

An industry is now forming around that question, and its answer will help determine how quickly the technologies we have written about can be deployed. It is hard for a business to trust an agent with consequential work if nobody will stand behind its mistakes.

The Market Split

In January 2026, US commercial insurers gained a new way to exclude generative AI risks. [Verisk](#)'s Insurance Services Office (ISO) introduced optional general liability exclusions: form CG 40 47 removes coverage for bodily injury, property damage, and personal and advertising injury arising out of generative AI.¹

In March 2026, [HSB](#), a [Munich Re](#) company, launched dedicated AI liability coverage for small and medium-sized businesses, designed to fill some of those gaps.² It joined [Armillia](#) and [Testudo](#), which had already launched standalone AI liability products.^{3,4} **One part of the industry is writing AI out of coverage. Another is building products around the same risk.**

Both responses are reasonable. A distinct AI liability market is forming, alongside performance warranties and extensions to existing policies. It is still small: Deloitte projects AI-related premiums of about \$4.7 billion a year by 2032, an estimate modeled on cyber's early adoption curve rather than AI loss data.⁵ But we do not believe the lasting value will lie with those providing the capital to pay claims. Uncertainty makes capacity harder to price, but better evidence and tighter coverage terms could attract more of it. Our thesis is that the most defensible businesses in this market may sit above the balance sheet: **the standards that define an acceptable level of agent risk, the proprietary data linking agent failures to claims, and the distribution channels that attach coverage when an agent is deployed or transacts.** Many entrants rely on someone else's underwriting capacity. That need not stop them from building a business around data or distribution they own.

The Risk Underwriting Has to Adapt To

Three Centuries, Three Assumptions

Underwriting is how an insurer prices risk. The insurer collects premiums from many policyholders, holds capital for years when claims exceed expectations, and transfers some of its largest exposures to reinsurers. Staying solvent depends on pricing that coverage correctly. The model has a long history. Edmond Halley published an influential mortality table in 1693, and [Equitable Life](#) put age-based life assurance on an actuarial footing in 1762, drawing on James Dodson's work.⁶ Three centuries of refinement rest on three assumptions:

1. **The past predicts the future.** Actuaries combine a policyholder's loss history with that of the wider pool, using a method called credibility theory. They use that record to estimate next year's claims, adjusting for changes in exposure.
2. **The insured thing** holds still. A factory, driver, or company changes slowly enough that its risk can be assessed over a policy term, with material changes disclosed and reflected in the cover.
3. **Losses are** diversified. Pooling many exposures can make aggregate losses more predictable. Insurers already manage correlated risks such as hurricanes; the challenge is to identify concentrations and hold or transfer enough capital against them.

Consider how underwriting prices familiar risks: a commercial account comes with five years of claims data; a driver has a record, a credit score, and, increasingly, a telematics feed; a physician has a claims history, a specialty classification, and board certification. The underwriter prices the individual record against the class, adjusts for what is observable, and often re-rates at renewal, while allowing for changes during the year. Many inputs in that process are data assets somebody built: the mortality table, the loss bureau, the credit score, the certification. **Underwriting looks like a capital business; underneath, it is a records business.**

Four Pressure Points


AI agents challenge those assumptions and make the underlying risks harder to observe:

1. **Experience rating** becomes harder. An agent substantially updated over a weekend may no longer behave like the previous version. Its loss history has to be read in the context of the version and use case.
2. **The policy term** becomes harder to manage. A twelve-month contract can cover many versions of a system assessed when the policy was written.
3. **Diversification** can weaken. Thousands of deployments relying on the same model can share a failure mode, exposing apparently separate insureds to correlated losses.

4. **Observability** is incomplete. Evaluations, logs, and agent registries exist, but there is no widely accepted record connecting an agent's behavior to the probability and cost of an insured claim.⁷

The underwriter of a human sees a mature record. The underwriter of an agent sees a much shorter one.

Limited evidence makes AI risk harder to price, while shared exposures and uncertain liability make it harder to contain. Better observation can help underwriters decide what to cover and on what terms. That creates an opportunity to supply useful evidence throughout the insurance process.



What Goes Wrong

Start with the losses themselves, many of which are already familiar. If an agent is hacked or makes an unauthorized charge, existing cyber, fraud, and payment-dispute rules may apply. US consumer protections include Regulation Z for credit cards and Regulation E for electronic transfers, subject to their scope and conditions.⁸ If a vendor ships defective software, product liability or errors and omissions (E&O) cover may respond, depending on the nature of the loss and the policy wording.

The harder question is what happens when an agent acts within its authority but gets something wrong. A support agent promises a refund the company never offered. That happened to [Air Canada](#): in February 2024, a British Columbia tribunal held the airline liable for its chatbot's misleading bereavement-fare information and ordered it to pay approximately C\$800, including interest and fees. The ruling was narrow, involving a negligent-misrepresentation claim on specific facts, but its logic has broader implications: the deployer may be held responsible for what its agent says.⁹

The wider legal exposure is growing. [Testudo](#)'s January 2026 report records a 137% increase in US generative AI lawsuit filings in 2025 in its proprietary dataset,

following 59% growth in 2024.¹⁰ But patent and copyright infringement account for the two largest claim categories, while hallucinations feature in only 4.9% of suits. These figures show the breadth of generative AI litigation, not the prevalence of agent errors or the scale of resulting insured losses.¹¹

What You Actually Insure


In most established lines of insurance, it is clear what is being insured. With agents, even that basic question is harder to answer. Agents can be cloned, deleted, rolled back, or silently updated overnight. The policyholder is generally the company developing or deploying the system; the covered exposure may be defined by its AI use, a specified application, or a performance promise. Certification is one underwriting route among several. [AIUC](#), [Armilla](#), and [Testudo](#) take materially different approaches, as the comparison below shows. These approaches test related questions: can technical evaluation predict failures, can deployment context predict legal exposure, and how do those risks translate into claims?

ACTIVANT RESEARCH			
	Provider & Customer	What is covered	Underwriting Approach
	AI builders and deployers	AI-agent exposures linked to certified systems	AIUC-1 testing and audit evidence
	AI vendors and enterprise users	AI liability; separate performance-warranty offering	Certification and risk reporting informed by 500+ evaluations
	US enterprises deploying generative AI	Third-party liability from generative AI use	Litigation, regulatory and deployment signals; no required technical audit or certification
	Startups and technology businesses	Selected AI/algorithmic modules added to technology E&O	Integrated platform; carrier and policy terms depend on placement
Public product descriptions as of October 7, 2026. Coverage remains subject to underwriting and policy terms; this is a selected comparison, not an exhaustive market count.			

For investors, the distinction matters. Certification can help underwriters understand a system's behavior, while litigation and claims data can reveal losses that technical tests miss. Either approach must still show that its information improves risk selection or pricing. To see where each can earn a return, it helps to follow how a policy is assessed, sold, and funded.

Five Hands, One Policy

An AI insurance policy depends on five functions, from assessing the risk to paying a claim. The market map below shows the companies pursuing those functions and the infrastructure supporting them. Some businesses perform several of these functions. For investors, the question is which gives them a lasting advantage.



1. Capital

At the base is the balance sheet that absorbs the loss. Reinsurers and specialty carriers already provide meaningful capacity: [Munich Re](#) says it wrote its first AI performance policy in 2018, covering an anti-fraud model; [Chaucer](#) backs [Armilla](#), while [Apollo](#), [Atrium](#), and [QBE](#) support [Testudo](#).^{12,13,14} The publicly announced limits illustrate the scale: [Armilla](#)'s standalone product offers up to \$25 million per organization, while [Testudo](#)'s February 2026 capacity announcement specifies \$10 million per insured.¹⁴ Those are policy limits, not company capital or premium revenue. Building a licensed, adequately capitalized carrier is possible, but it is a different undertaking from launching a managing general agent (MGA).

2. Origination

Working with that capital are the companies that select risks, price them, and sell policies. This is where the new entrants cluster, but their customers and methods differ. [AIUC](#) announced a \$40 million Series A in September 2026, bringing disclosed funding to \$55 million.¹⁵ [Armilla](#) combines evaluation and liability underwriting; in February 2026 it joined [Chaucer](#) in Vanguard AI, coordinating AI liability with cyber and technology E&O. [Testudo](#) launched its standalone enterprise product in January 2026 and clarified in September that its covered AI exposures are intended to respond on a primary basis, subject to policy terms.¹⁶ [Corgi](#) takes a more integrated approach, with an affiliated carrier as well as partner-carrier arrangements and a modular AI endorsement to technology E&O. It announced a \$106 million Series B1 at a \$2.6 billion valuation in May 2026. That funding backs its whole insurance platform, of which AI liability is one line.¹⁷

An MGA exercises underwriting authority delegated by an insurer, typically earning commission and sometimes profit-related compensation. The insurer

carries the contractual claim obligation; an MGA can still own the customer relationship, underwriting process, and data. [Mount](#) and [Klaimee](#) also combine assessment with financial protection for agents, while [Mosaic](#)'s partnership with [Munich Re](#) offers up to 15 million in euros, US dollars, or Canadian dollars for defined AI performance failures. The products differ enough that a simple count would overstate the market.^{18,19,20}

3. The Standard and the Record

To underwrite these risks, insurers need evidence of how agents fail. Building that evidence is becoming a business in its own right, with three approaches emerging: certification, identity and trust, and large-scale testing of agent failures.

Certification sets a standard. AIUC-1 covers six domains, uses quarterly technical evaluations and draws on MITRE ATLAS (Adversarial Threat Landscape for Artificial-Intelligence Systems) and the Open Worldwide Application Security Project (OWASP). [Schellman](#) provides independent audit evidence and reporting, while [AIUC](#) conducts technical evaluations and issues certification.²¹

Identity and trust establish who is authorized to act. Credit bureau [Experian](#) launched Agent Trust on April 30, 2026, linking people to authorized agents through a trust token and a registry that evaluates behavior over time. This is an identity and fraud-risk framework, not yet a proven insurance loss score.²²

Evaluators generate evidence of how agents fail at scale. [Gray Swan](#) says its community of more than 15,000 red-teamers has generated over three million attack attempts.²³ A March 2026 analysis by the National Institute of Standards and Technology (NIST) of a Gray Swan-hosted competition examined more than 250,000 attacks against thirteen frontier models.²⁴ Large security companies are already buying these capabilities: [Cisco](#) acquired [Robust Intelligence](#) in 2024, and [Palo Alto Networks](#) acquired [Protect AI](#) in 2025.^{25,26}

Those deals show demand for AI security tools. Their value to insurers still depends on whether the test results help predict costly claims. That is harder than it looks, because the two records describe different risks. Red-teaming measures security and behavioral failure. The claims arriving today are mostly something else: in Testudo's data, patent and copyright claims lead, while hallucinations feature in under 5% of suits. Testudo interprets this as evidence that model performance has limited impact on a firm's liability exposure.¹¹ Our view is that the durable record is the one tied to paid claims. Technical evaluation earns its value where agent behavior drives the loss: authorized transactions, customer commitments, and operational errors. We expect those losses to account for a larger share as agents gain authority. Litigation and claims data help price the other exposures. The strongest business will bring both kinds of evidence together, linking failures found in testing to the costs they produce in court.

Model change creates a test for this data advantage. A failure observed in one version may say little about the next. Our hypothesis is that more durable

information concerns the deployment: the authority an agent receives, the controls on its actions and the obligations of the business using it. For a purchasing agent, that means recording spending limits and approval rules alongside model versions, then linking those conditions to errors and resulting claims. The advantage depends on showing which relationships remain predictive across updates and customers. A growing archive alone is not enough.

4. Distribution

The policy still has to reach a buyer. A promising enterprise channel is the AI vendor itself. In February 2026, [ElevenLabs](#) announced what it described as the first live AIUC-1-backed policy covering AI voice agents, following more than 5,000 adversarial simulations. ElevenLabs says its ElevenAgents platform powers more than three million voice agents, which shows the reach of the vendor channel. Its customers can pursue certification and coverage, though insurance is not yet bundled into every sale.²⁷ [Google Cloud](#)'s Risk Protection Program offers access to partner cyber policies from [Beazley](#), [Chubb](#), and [Munich Re](#), with AI-related terms varying by product.²⁸ [Coalition](#) added a Deepfake Response Endorsement in December 2025.²⁹ On the consumer side, [American Express](#)'s announced commitment to protect eligible registered-agent purchases points to the payment rails as another channel. Its published materials still describe the protection in future terms.³⁰ **An authorized mistake, such as a wrong flight or an overpaid invoice, can fall outside card fraud protections. We expect specialist underwriting to reach consumers through platforms and wallets rather than through a new consumer insurance brand.**

5. Identity

Underneath many agent transactions is identity, the passport layer proving an agent is authorized: [Skyfire](#), [Crossmint](#), [Nekuda](#), and the network protocols. As we discussed in [Authenticating AI Agents](#), connecting an action to an accountable principal is a prerequisite for effective controls.³¹ Identity tells an underwriter who acted, but says little about what the loss will cost. Our concern is that platforms will bundle protocol-level identification into their services, leaving independent providers to prove their value in fraud prevention, permissions, or cross-platform trust.

Which of these functions lets a business build something its customers and capital providers cannot easily replace? That is where we would look to invest.

What Does the Underwriter Own?

Being close to the customer is not enough if an originator cannot secure capital to back its policies. While capital is scarce, its providers hold the bargaining power. Clearer policy terms, limits on shared exposures, and reinsurance could attract more capacity, eventually drawing in capital-market investors. If that happens,

leverage can shift toward businesses with the data or distribution that makes the risk worth funding. Our thesis depends on that shift: the businesses we favor gain power only if capital becomes plentiful. If long-tail losses cannot be contained, capital stays scarce and its providers keep their leverage. That is the risk we return to in the first bear case.

[Verisk](#) offers a precedent: a business built on pooled insurance information, analytics, and policy forms.³² An insurer commits capital to absorb claims; a data business can sell the same information to many customers without funding their losses. An originator can therefore rent a balance sheet while owning a record that remains valuable when its capital provider changes. For an AI underwriting business, that means building data that improves risk selection and becomes part of customers' decisions. The rights to retain and reuse it matter as much as collecting it. Verisk shows what that model can become, not what an emerging AI certifier has already earned.

Certification is less straightforward. [AIUC](#) is trying to establish a privately owned standard through AIUC-1. SOC 2 (System and Organization Controls 2) offers a different model: professional standards, independent examinations, and businesses selling the compliance workflow.³³ For AIUC, combining evaluations and certification with access to insurance makes audit independence important, even where outside auditors provide evidence. Buyers will decide whether they need a particular certificate, accept equivalents, or simply require an insured outcome. If certificates become interchangeable, the stronger business may be the one that owns the loss record and the customer workflow. A certifier earns that position only if its evaluations and data improve risk selection or pricing.

Cyber Was the Dress Rehearsal

Cyber insurance offers a precedent for how difficult-to-measure digital risk can become fundable. Annual premiums now total nearly \$15 billion, a figure Munich Re expects to reach about \$28 billion by 2030.³⁴ Tighter policy terms, stronger security requirements, and better models of shared exposure helped insurers define and limit what they could cover. Reinsurance spread the risk across more balance sheets, while cyber catastrophe bonds brought in capital-market investors, with five issued for \$575 million in the twelve months to June 2024.³⁵ AI may not follow the same pattern. Still, cyber shows how capacity can expand as risks become easier to measure, constrain, and price.

The July 2024 [CrowdStrike](#) outage showed how one software failure can affect many customers at once: a single faulty update took down an estimated 8.5 million Windows devices.³⁶ Shared dependencies remain a problem even in an established insurance market. AI adds uncertainty about how harmful outputs might generate liability over years. Not every AI exposure will follow cyber into insurability.

We expect capital to arrive in waves as particular AI exposures become better understood and bounded. One instrument to watch for is an AI catastrophe bond, paying investors to absorb a defined layer of catastrophic AI losses, much as catastrophe bonds already do for hurricanes. Its emergence would signal that investors are willing to fund a defined part of that risk beyond insurers' balance sheets.

How We Could Be Wrong

Three arguments could break this thesis.

1. AI Is Asbestos, Not a Hurricane

Cyber business interruption losses may be short-lived, but privacy and liability claims can take years to develop. A flawed frontier model deployed across thousands of enterprises could generate long-tail liability under policies never priced for that exposure. That is the asbestos pattern: insurers collect premiums for years before discovering the scale of the claims they have taken on, potentially across many policies at once. Version-specific evidence, recertification, and per-model sub-limits may help; they do not eliminate the possibility of systemic loss. If insurers cannot contain that long-tail risk at a price customers will pay, coverage limits stay low and the standalone market remains smaller than our thesis requires. This is the case that would break the thesis rather than narrow it, because it keeps capital scarce and leaves the power with whoever holds it.

2. The Incumbents Capture the Durable Layer

[Munich Re](#) sees claims across the insurers and programs it backs, while established data businesses already have distribution and insurer trust. A reinsurer does not need to see the entire market to build a useful dataset, and governance can address conflicts without requiring complete separation from underwriting. For startups, that is a serious competitive threat. A newcomer must assemble broader or better evidence, secure rights to use it, and demonstrate predictive value that an incumbent cannot readily replicate. Public standards may also move the economics toward the compliance workflow rather than a privately owned certificate. This narrows the thesis rather than breaking it: the durable layer still exists, but a startup is most likely to win where incumbent data is thin, such as agent behavior that never reaches a reinsurer's claims file.

3. The Demand Never Arrives as Insurance

An enterprise deploying an agent may demand a vendor indemnity rather than buy standalone insurance. The largest vendors may retain that exposure, while existing cyber or E&O insurers may absorb more AI risk into broader policies. An indemnity still leaves the vendor with a risk-management problem, but one it can partly solve in-house. Demand for controls and data could therefore grow even as

the standalone insurance market remains small. The thesis weakens if external risk pricing and specialist distribution fail to capture a meaningful share of that spend. The opportunity would then lie more in selling data and controls to vendors than in selling standalone policies.

A useful service is not necessarily an attractive standalone business. The functions we favor may prove important without producing strong returns for startup investors.

What Passes the Screen

We apply one screen to every company in this market, and it has three parts:

1. **Own or rent?** Which parts of its data, underwriting process, or distribution does the company own, and how dependent is it on any one capital provider?
2. **Does it compound?** Does each evaluation, deployment, or claim add usable proprietary information that the company has the right to retain and learn from?
3. **What is the unit of risk?** Does the provider price the actual insured exposure, including the customer's use case, authority, model dependencies, and legal obligations, whether or not certification is part of its method?

A certification buyers value, a defensible distribution channel, or a failure dataset that demonstrably improves risk selection can pass. Premium volume alone cannot. The test is whether the company is building a transferable advantage alongside its book.

Our view: a certificate offers limited defensibility if buyers readily accept alternatives. We favor businesses that connect detailed deployment evidence to claims outcomes, retain the rights to learn from that record and own a distribution channel through which those insights inform underwriting. Insurers and reinsurers may already hold substantial claims data. A startup's advantage must therefore come from information they lack, or connections they cannot readily reproduce. Writing policies or issuing certificates provides a route to that advantage, but neither establishes it.

Where We Come Out

We will watch three developments to see whether the thesis holds. By the end of 2027, enterprise procurement should reveal whether buyers require a particular AI certificate, accept interchangeable standards, or instead rely on indemnities and insurance. At the end of 2028, we will assess whether third-party capacity has broadened and repeatable risk-transfer structures have emerged, with an AI catastrophe bond one possible signal rather than a prerequisite. And finally, we would track whether buyers renew specialist AI cover and pay for it separately or obtain comparable protection through existing policies and vendor contracts.

Insurers writing exclusions and the providers writing new products are responding to the same uncertainty: how AI exposure translates into loss. The standard, the dataset, and the distribution may outlast the capacity behind them. But that depends on showing that the data predicts losses, not just documents technical failures. Our [Agentic Commerce series](#), [Agentic Procurement](#), and [Robots: Valuations and The Scaling Question](#) all end at the moment an agent acts and someone must answer for it. The durable business in this market belongs to whoever can say, with authority, what that answer costs.

Give the agent a body and the same question becomes physical: who pays when a mistake causes injury or property damage? We examine that next in Insuring the Robot.

If you are building the standard, the dataset, or the distribution layer of this market, we would like to meet.

Disclaimer: The information contained herein is provided for informational purposes only and should not be construed as investment advice. The opinions, views, forecasts, performance, estimates, etc. expressed herein are subject to change without notice. Certain statements contained herein reflect the subjective views and opinions of [Activant](#). Past performance is not indicative of future results. No representation is made that any investment will or is likely to achieve its objectives. All investments involve risk and may result in loss. This newsletter does not constitute an offer to sell or a solicitation of an offer to buy any security. [Activant](#) does not provide tax or legal advice and you are encouraged to seek the advice of a tax or legal professional regarding your individual circumstances.

This content may not under any circumstances be relied upon when making a decision to invest in any fund or investment, including those managed by [Activant](#). Certain information contained in here has been obtained from third-party sources, including from portfolio companies of funds managed by [Activant](#). While taken from sources believed to be reliable, [Activant](#) has not independently verified such information and makes no representations about the current or enduring accuracy of the information or its appropriateness for a given situation.

[Activant](#) does not solicit or make its services available to the public. The content provided herein may include information regarding past and/or present portfolio companies or investments managed by [Activant](#), its affiliates and/or personnel. References to specific companies are for illustrative purposes only and do not necessarily reflect [Activant](#) investments. It should not be assumed that investments made in the future will have similar characteristics. Please see "full list of investments" at activantcapital.com/companies/ for a full list of investments. Any portfolio companies discussed herein should not be assumed to have been profitable. Certain information herein constitutes "forward-looking statements." All forward-looking statements represent only the intent and belief of [Activant](#) as of the date such statements were made. None of [Activant](#) or any of its affiliates (i) assumes any responsibility for the accuracy and completeness of any forward-looking statements or (ii) undertakes any obligation to disseminate any updates or revisions to any forward-looking statement contained herein to reflect any change in their expectation with regard thereto or any change in events, conditions or circumstances on which any such statement is based. Due to various risks and uncertainties, actual events or results may differ materially from those reflected or contemplated in such forward-looking statements.

Endnotes

- ¹ Independent Insurance Agents & Brokers of America, [Verisk to Roll Out New General Liability Exclusions for Generative AI Exposures](#), 2025.
- ² HSB / Munich Re, [HSB Introduces AI Liability Insurance for Small Businesses](#), 2026.
- ³ Armilla, [Lessons from a year of underwriting AI liability insurance](#), 2026.
- ⁴ Testudo, [Testudo Launches Standalone AI Liability Insurance for Generative AI Systems](#), 2026.
- ⁵ Deloitte, [Insurance for AI risk](#), 2024.
- ⁶ Society of Actuaries, The Actuary, [The History of Actuarial Science](#), 2024.
- ⁷ NIST, [Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile](#), 2024.
- ⁸ Consumer Financial Protection Bureau, [Regulation Z, 12 CFR 1026.12; Regulation E, 12 CFR 1005.6](#), 2026.
- ⁹ American Bar Association, [BC Tribunal Confirms Companies Remain Liable for Information Provided by AI Chatbot](#), 2024.
- ¹⁰ Testudo, [The State of Play: Generative AI Litigation Market Overview](#), 2026.
- ¹¹ Testudo, [The State of Play: Generative AI Litigation Market Overview](#), 2026.
- ¹² Munich Re, [Insure AI Frequently Asked Questions](#), 2026.
- ¹³ Armilla / Chaucer, [Chaucer and Armilla Launch Vanguard AI to Clarify Cyber, Technology and AI Liability in a Single Coordinated Structure](#), 2026.
- ¹⁴ Testudo, [Lloyd's Syndicates Commit More Capacity to AI Liability Insurance](#), 2026.
- ¹⁵ AIUC, [Series A to build the confidence infrastructure for frontier AI](#), 2026.
- ¹⁶ Testudo, [A Small but Mighty AI Liability Insurance Product Update](#), 2026.
- ¹⁷ Corgi, [Corgi Reaches \\$2.6B Valuation With \\$106M Series B1 Three Weeks After \\$160M Unicorn Round](#), 2026.
- ¹⁸ Mount, [Welcome to Mount](#), 2026.
- ¹⁹ Klaimée, [Insurance for AI agent mistakes](#), 2026.
- ²⁰ Mosaic Insurance, [Mosaic partners with Munich Re's aiSure to provide pioneering coverage for AI vendors](#), 2026.
- ²¹ AIUC-1 / Schellman, [Schellman becomes the first accredited auditor for AIUC-1](#), 2026.
- ²² Experian, [Experian Announces Agent Trust to Power Trusted AI Driven Commerce](#), 2026.
- ²³ Gray Swan, [Enterprise Security for AI-Powered Applications; About Gray Swan](#), 2026.
- ²⁴ NIST, [Insights into AI Agent Security from a Large-Scale Red-Teaming Competition](#), 2026.
- ²⁵ Cisco, [Fortifying the future of Security for AI: Cisco Announces intent to acquire Robust Intelligence](#), 2024.
- ²⁶ Palo Alto Networks, [Palo Alto Networks Completes Acquisition of Protect AI](#), 2025.
- ²⁷ ElevenLabs, [ElevenLabs secures first-of-its-kind AI Agent insurance](#), 2026.
- ²⁸ Google Cloud, [Risk Protection Program](#), 2026.
- ²⁹ Coalition, [Coalition Adds Deepfake Response Endorsement to its Cyber Insurance Policies Globally](#), 2025.

³⁰ American Express, [Agentic Commerce Experiences](#), 2026.

³¹ Activant Research, [Authenticating AI Agents](#), 2025.

³² Verisk, [2018 Form 10-K, Corporate History](#), 2019.

³³ AICPA & CIMA, [System and Organization Controls: SOC Suite of Services](#), 2026.

³⁴ Munich Re, [Global Cyber Risk and Insurance Survey 2026](#), 2026.

³⁵ Aon Securities, [Insurance-Linked Securities Annual Report 2024](#), 2024.

³⁶ Microsoft, [Helping our customers through the CrowdStrike outage](#), 2024.