MARKET ASSESSMENT | SEPTEMBER 2026

112 AI Exclusions, Zero Filed AI Policies

Whether losses caused by enterprise AI agents are insured today, and on what terms, across US and London cyber, tech E&O, D&O and specialist AI cover

David Soden

Independent analyst and practitioner, enterprise automation and applied AI

davidsoden.com/market-assessment

REPORT DS-2026-43C3 · 2026 · VERSION 1.0 · PUBLIC

About this report

Every substantive claim in this report carries an evidence classification and, where applicable, a numbered source. Facts are separated from vendor claims, third-party estimates, the author's assessments, and untested hypotheses. Forward-looking sections use scenarios with observable tripwires rather than point forecasts. Stated assumptions are listed before the analysis begins.

Market Assessment reports are published and commissioned at **davidsoden.com/market-assessment**. This document is analysis, not investment, legal, or procurement advice.

Scope and evidentiary standard

This report asks whether a loss caused by an enterprise AI agent is insured today, and on what terms. It covers the US admitted and surplus lines markets and the London market, across four product families: cyber, technology errors and omissions, directors and officers liability, and the specialist affirmative AI products written since 2024. It answers four questions. Are AI or autonomous-system exclusions entering standard forms? Who writes affirmative cover for agent errors, at what limits and under what conditions? Does AIUC-1 certification change underwriting terms? And how are carriers modelling accumulation from shared models and shared clouds?

Commentary on this subject runs well ahead of the documents. Trade press, law-firm alerts and a growing crop of AI-generated explainer sites repeat each other's figures, and several widely quoted numbers trace back to a single paywalled article or to a company's own launch release. So this report counts forms rather than quoting commentary. Where a policy wording, a rating-bureau form, a market bulletin or a contract term could be read, it is cited as such. Where the only evidence is a press release or an executive's remark, the report says so and does not let a conclusion rest on it.

Two structural facts limit what can be counted. Surplus lines and Lloyd's business in the US is largely exempt from rate and form filing, so the specialist AI products that matter most leave no trace in SERFF, the state filing system. And almost none of the specialists publishes a specimen wording. Silence in the filing record is therefore evidence about the admitted market only, and absence of a public wording is a finding in its own right.

Every substantive claim below carries one of six classifications, and any claim resting on an external source carries a numbered reference linking to the source list.

Label	Definition	How to treat it
FACT	Verifiable in the cited primary source: a filing, press release, official documentation, or standards body announcement.	Reliable as stated. Check the source if the exact wording matters.
VENDOR CLAIM	Reported by a vendor about its own business or product. Self-measured and not audited by any third party.	The vendor said it. Directionally useful, not a measurement. Definitions of the metric are usually undisclosed.
THIRD-PARTY ESTIMATE	A research firm forecast or survey result. Methodology is proprietary and generally not reproducible.	Use for direction and order of magnitude only. Do not build a model on it.
ASSESSMENT	The author's reasoned judgment from the cited evidence. Falsifiable, and the reasoning is shown so it can be argued with.	This is analysis, not reporting. Disagree with it where your evidence differs.
HYPOTHESIS	A proposition offered for testing. Not currently supported by sufficient evidence to assert.	Treat as a research question, not a conclusion. Each one names what would confirm or refute it.
SCENARIO	A structured future state with a stated subjective probability. The probability is the author's judgment, not a calculated figure.	Use the leading indicators attached to each, not the probability. The indicators are observable; the probability is not.

Assumptions this analysis rests on

Four premises are assumed rather than shown. Each names the conclusions that fail if it is wrong.

One. The RAND filing count is close to complete for the admitted market

The count of admitted-market AI forms comes from RAND's keyword search of a commercial SERFF index through 21 June 2026. The report assumes that search caught most AI exclusions filed under names such as artificial intelligence, machine learning and generative AI. If carriers are excluding AI through broader definitions that never use those words, for example by widening a computer-system or technology-services exclusion, the count understates the retreat and the finding that exclusions are a patchwork becomes too generous.

Two. Press-reported limits are close to what can be bound

Specialist limits are taken from company releases and trade press, because no specimen wording is public. The report assumes those figures describe the ceiling a large buyer could actually place. If real bound limits sit well below them, which is common in a new line, the retained exposure described in the Implications section is larger than stated.

Three. Silence is read in the policyholder's favour, at least initially

The report treats a policy that neither grants nor excludes AI loss as likely to respond if the loss otherwise fits the insuring agreement. No court has yet tested this for an agent-caused loss. If early coverage litigation goes the other way, silent cover is worth less than assumed and the protection gap is wider today, not just in future.

Four. London follows its own silent-cyber precedent

The scenario weights assume Lloyd's would handle AI the way it handled cyber: a bulletin requiring clarity, then a model clause, phased in over a renewal cycle. If Lloyd's instead leaves AI to individual syndicates, the probability of a London-led clarity mandate falls and the patchwork scenario gains.

The call

ASSESSMENT Losses caused by enterprise AI agents are insured today mostly by default, through cyber, E&O and liability wordings that never contemplated them. That default is being withdrawn one carrier and one optional endorsement at a time, not by any market-wide clause. Affirmative cover is real but thin: a handful of specialist MGAs on overlapping Lloyd's and reinsurance capacity, at limits a large deployment can exhaust, with no public specimen wording. AIUC-1 moves the price of AIUC's own paper and nobody else's. Treat agent losses above specialist limits as retained. The patchwork reading fails the day Lloyd's issues a market bulletin requiring every syndicate to state whether AI loss is covered or excluded.

My judgment on the evidence assembled here, nothing more. There is data I can't see and data I may have missed. New evidence can move me, including to the opposite position, and I reserve the right to move.

Executive summary

The market for insuring AI agent losses has split three ways. Some carriers exclude AI outright, a small group of specialists writes affirmative cover, and most policies say nothing. Exclusions are the part that can be counted, and they are growing faster than the affirmative cover. The filing record shows a hundred-odd exclusion forms against a handful of endorsements, and not one standalone AI policy filed in the admitted market. The findings below separate what is filed from what is announced.

FINDING 1 **US standard forms now carry AI exclusions. London's do not.**

> **FACT** Verisk's ISO unit released three optional generative AI exclusions for commercial general liability and products liability, effective January 2026. They are optional endorsements, and the base CGL form is unchanged. No Lloyd's market bulletin or LMA model clause on AI exists. The LMA is consulting on a model AI definition. [2] [3] [4] [13]

FINDING 2 **Carrier filings, not standard forms, are doing the excluding.**

> **FACT** RAND counted 112 AI exclusion forms from 60 insurance companies in the US admitted market through June 2026, against seven endorsements and no standalone AI policy. Most sit

on umbrella and general liability, where limits are highest. The surge began in the summer of 2025. [1]

FINDING 3 Affirmative agent cover tops out at $25 million per insured in primary sources.

VENDOR CLAIM Armilla and Chaucer state $25 million or more per organisation. Testudo states $1 million to $10 million. Mosaic's aiSure partnership states $15 million of initial capacity. AIUC's $50 million appears only in press and think-tank accounts, never in AIUC's own releases. [21] [24] [25] [29] [10]

FINDING 4 No specialist publishes a wording, so the conditions are known only in outline.

ASSESSMENT Cover is gated by testing or certification (AIUC, Armilla), by historical performance data (aiSure), or by litigation-data underwriting without an audit (Testudo). Exclusions for misuse, unauthorised changes, prior known defects and intentional acts recur wherever they are described. [1] [22] [28]

FINDING 5 AIUC-1 changes terms only on AIUC's own programme.

ASSESSMENT The standard calls itself insurance-enabling but its text sets no eligibility, price or discount. No carrier filing, wording or guideline that references it was found. The only stated link to pricing is an AIUC founder describing how its test results feed its own Lloyd's-backed underwriting. [36] [37] [47] [48]

FINDING 6 Nobody has published an insured-loss figure for a shared-model failure.

FACT Lloyd's scenarios, the mandatory cyber disaster scenarios, the commercial cyber models and reinsurer research all quantify cloud outages or name AI concentration qualitatively. None puts a number on a foundation-model failure. CrowdStrike is the only real shared-dependency event priced, and published estimates of it span roughly five to one. [49] [51] [52] [53] [56] [57] [64]

FINDING 7 A market-wide exclusion is more likely to follow a correlated loss than a regulator.

ASSESSMENT The only mandate precedent is Lloyd's silent-cyber and state-backed-attack bulletins, and both came after years of loss experience. The tripwires in this report are observable without private data, and they matter more than the scenario probabilities. [11] [12]

Exclusions: what is actually filed

Whether AI exclusions are entering standard forms depends on which forms count as standard, so this section separates the rating-bureau forms that many carriers adopt from individual carrier paper, and both from the London market.

The ISO endorsements

FACT ISO's sample form CG 40 48 01 26, Exclusion - Generative Artificial Intelligence (Coverage B Only), reads in full on its operative line: "This insurance does not apply to: 'Personal and advertising injury' arising out of 'generative artificial intelligence'." It defines generative AI as "a machine-based learning system or model that is trained on data with the ability to create content or responses, including but not limited to text, images, audio, video or code." [2]

FACT Two companion forms extend the same exclusion. CG 40 47 01 26 removes bodily injury, property damage and personal and advertising injury arising from generative AI under the CGL coverage part. CG 35 08 does the same for products and completed operations. Every description of all three, including ISO's trade-association announcement, calls them optional endorsements effective January 2026, in occurrence and claims-made versions. [3] [4]

FACT Verisk's liability lead told Claims Journal in July 2026 that the company sees a high degree of interest in the forms but does not know how many insurers have adopted them. RAND likewise states it has no information on how many carriers use them across states. [5] [1]

ASSESSMENT So AI exclusions have entered the standard form library, and that part of the question has a clean yes. They have not entered the standard policy. An optional endorsement is a tool a carrier may attach, and the base form CG 00 01 is unchanged. Nor does the library extend to the lines this report cares most about. No ISO AI endorsement for umbrella, businessowners, professional or cyber forms was found. The definition is also narrower than it looks for agent risk. It targets content creation, and an agent that executes a wrong payment is using a generative model to take an action, which is exactly where the wording will be argued over first.

Carrier paper

FACT RAND retrieved AI-related forms from a commercial index of SERFF filings through 21 June 2026. After merging near-duplicates it reports 125 forms, of which 112 are exclusions, seven are endorsements and four are applications, submitted by 60 distinct insurance companies. Most exclusions attach to commercial umbrella and excess or to general liability. RAND also states that no standalone AI policy had been filed in the admitted market as of writing. [1]

FIGURE 1 **What the US admitted market has filed on AI**

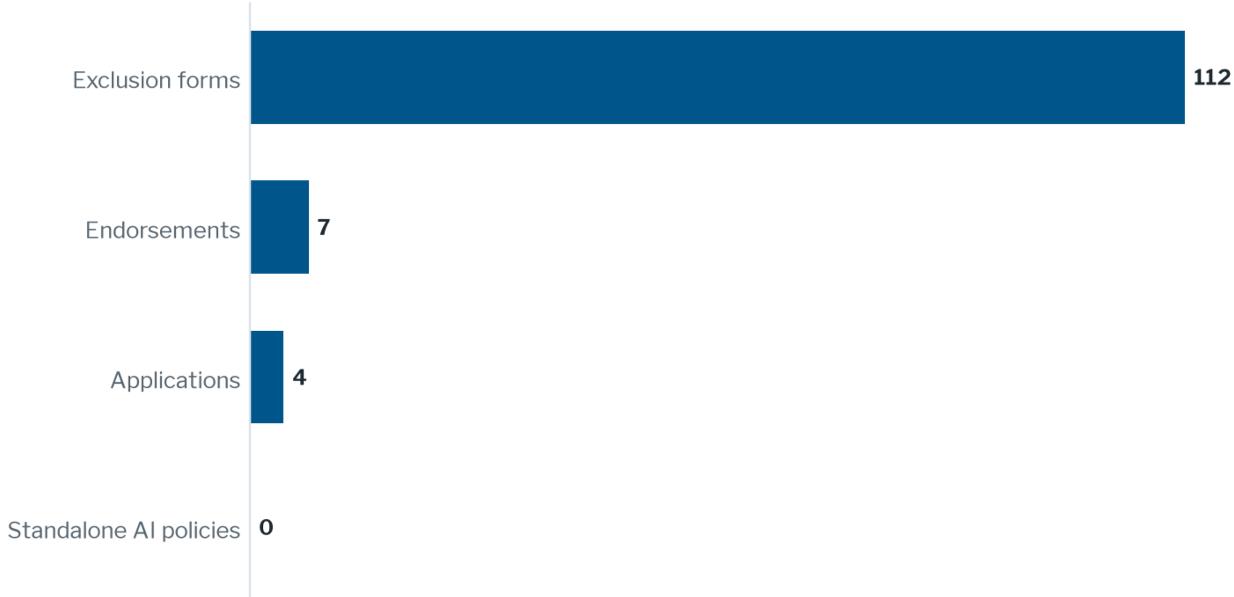

FACT AI-related forms in SERFF filings, December 2023 to 21 June 2026, from RAND's keyword search of a commercial index. The categories sum to 123 against the 125 forms RAND reports, a gap the report does not explain. Surplus lines and Lloyd's products are largely outside SERFF and are not counted here. [1]

FACT The broadest carrier wording in the public record is W. R. Berkley's Artificial Intelligence Exclusion for D&O, E&O and fiduciary coverage parts, which RAND reproduces from a June 2024 form. It removes any claim arising from "any actual or alleged use, deployment, or development of Artificial Intelligence by any person or entity", from the insured's statements about its AI use and AI risks, from violations of AI laws, and from any demand that the company investigate or respond to AI risk. It separately names representations made by a chatbot or virtual customer service agent. [1] [6]

FACT Hamilton Select's professional liability exclusion, as quoted by counsel, bars claims "based upon, arising out of, or in any way involving" generative AI and names specific consumer products in its definition. No form number was found. [7]

THIRD-PARTY ESTIMATE Two widely repeated figures are weaker than their circulation suggests. The claim that more than 80 percent of AI exclusion filings by Berkshire Hathaway, Chubb and Travelers were approved traces to a paywalled article citing a Wolfe Research analysis that is not public. CSIS repeats it and adds that more than 60 carriers filed AI exclusions for 2026. That figure agrees with RAND's independent count of 60 companies, which is the better source for it. The 80 percent figure is unverified and nothing here rests on it. [9] [10] [1]

VENDOR CLAIM The FT's November 2025 report that AIG, Great American and W. R. Berkley sought approval for AI exclusions is often cited as proof of a coordinated retreat. AIG told TechCrunch it was not specifically seeking to use the exclusions and had no plans to implement them. That is an executive statement, and it carries no more weight than the report it answers. No filing or form number for either side was found. [8]

Cyber is moving the other way

FACT The cyber line is the exception. No AI exclusion in a standalone cyber form was found. Coalition added an affirmative AI endorsement in March 2024 that counts an AI security event as a security failure and extends funds-transfer fraud to deepfakes. Beazley confirmed an AI clarifying endorsement to cyber and tech E&O in September 2026. CFC reports affirmative AI language across seven London products. [17] [18] [19] [20]

ASSESSMENT Read closely, the cyber endorsements cover attacker-side AI: deepfake fraud, AI-assisted intrusion, AI as a route into the network. Beazley's is explicit that it addresses AI used against the insured. None of them affirmatively covers an agent the insured deployed doing the wrong thing without any security event. That loss falls between cyber, which wants a breach, and E&O, which wants a professional service, and it is where silent cover is thinnest.

London

FACT No Lloyd's market bulletin on AI exclusions or AI clarity exists. Lloyd's has mandated cyber treatment twice: bulletin Y5258 in July 2019 required every policy to state affirmatively whether cyber loss was covered or excluded, and Y5381 in August 2022 required state-backed cyber-attack exclusions in standalone cyber policies from March 2023. The LMA's AI loss-scenario survey treats AI as software within the existing computer-system definitions of LMA cyber clauses and publishes no AI clause. [11] [12] [14]

FACT The LMA is consulting members on a model AI definition, according to The Insurer in September 2026. Its head of technical underwriting said earlier in the year that the LMA was waiting for an instruction from the market to produce something. Both reports sit behind paywalls, and only their published summaries were reviewed. [13] [15]

FACT US regulators have not stepped in on either side. The NAIC's March 2026 issue brief addresses insurers' own use of AI, through its model bulletin and a pilot evaluation tool, and says nothing about AI exclusions in policies. No state disapproval of an AI exclusion was found. [16]

ASSESSMENT Counting forms answers the first question. AI exclusions are entering US standard forms as optional ISO endorsements and entering US carrier paper at scale. They are not entering London standard wordings at all yet. A model AI definition, when the LMA publishes one, is the precondition for a model exclusion, which makes it the earliest London tripwire.

Affirmative cover: who writes it and on what terms

Affirmative cover here means a wording that names AI model or agent failure as a covered cause of loss, as distinct from a cyber policy that happens to respond when AI is involved in a breach. The table lists what is on offer, what the stated limit is, and where that limit comes from.

Provider	What it covers	Stated limit	Gate to cover	Limit source
Armilla, Lloyd's coverholder	Third-party liability for model underperformance, hallucination, agent errors in decisions or execution, non-breach data leakage, AI regulatory violations	$25m per insured; $5m per model use	Pre-bind assessment; certification from 500-plus evaluations	Company page and trade press
Chaucer with Armilla, Vanguard AI	Cyber and tech E&O from Chaucer; Armilla AI cover where no cyber event occurred	$10m cyber plus $25m or more AI aggregate	Model certification; allocation rules for mixed events	Carrier release
AIUC, MGA	First and third party, including agent tool-call failures, wrong payments or refunds, reputational harm	Up to $50m	AIUC-1 certification and quarterly retesting	Press and CSIS only
Testudo, Lloyd's coverholder	Claims-made third-party cover for GenAI deployers across six insuring agreements	$1m to $10m	Litigation-data scoring, no technical audit	Company release
Munich Re aiSure, with Mosaic	Performance guarantee against agreed model metrics, settled on defined thresholds	Up to $15m initial capacity	Due diligence on pipeline, data, drift and monitoring	Carrier release
HSB (Munich Re)	Bodily injury, property damage and advertising injury caused by AI, for small firms	Not stated	Sold through partner carriers	Pending regulatory approval at launch

FACT Chaucer's February 2026 release is the most specific primary statement of structure in the market. It pairs $10 million of Chaucer cyber limits with "dedicated AI aggregate limits of $25 million or more per organisation", and says Armilla's policy responds to "loss scenarios driven by AI model behaviour - including erroneous outputs, model underperformance, or AI agent actions - where no cyber event has occurred." Predefined allocation rules split mixed events by harm. [24]

VENDOR CLAIM Armilla lists Chaucer, AXIS Capital, Convex, Greenlight Re and Swiss Re as capacity partners and says documentation is available on request. Testudo names Apollo as lead with Atrium and QBE, reaching $9.25 million of capacity per insured. RAND reports Armilla's per-model limit of $5 million from a company fact sheet. Trade press dates the move to $25 million per insured, and the addition of agent failures such as improper tool use, to Armilla's January 2026 funding round. [22] [26] [1] [23]

VENDOR CLAIM Munich Re says it has written aiSure performance guarantees since 2018 and describes its limits as flexible, without a figure. Its HSB unit announced an AI liability product for small businesses in March 2026, sold through partner carriers and pending regulatory approval. If that product is approved as a filed form, it would be the first affirmative AI liability policy in the admitted market and would retire half of this report's title. [27] [28] [30]

FIGURE 2 **Stated maximum limits for affirmative AI cover, per insured**

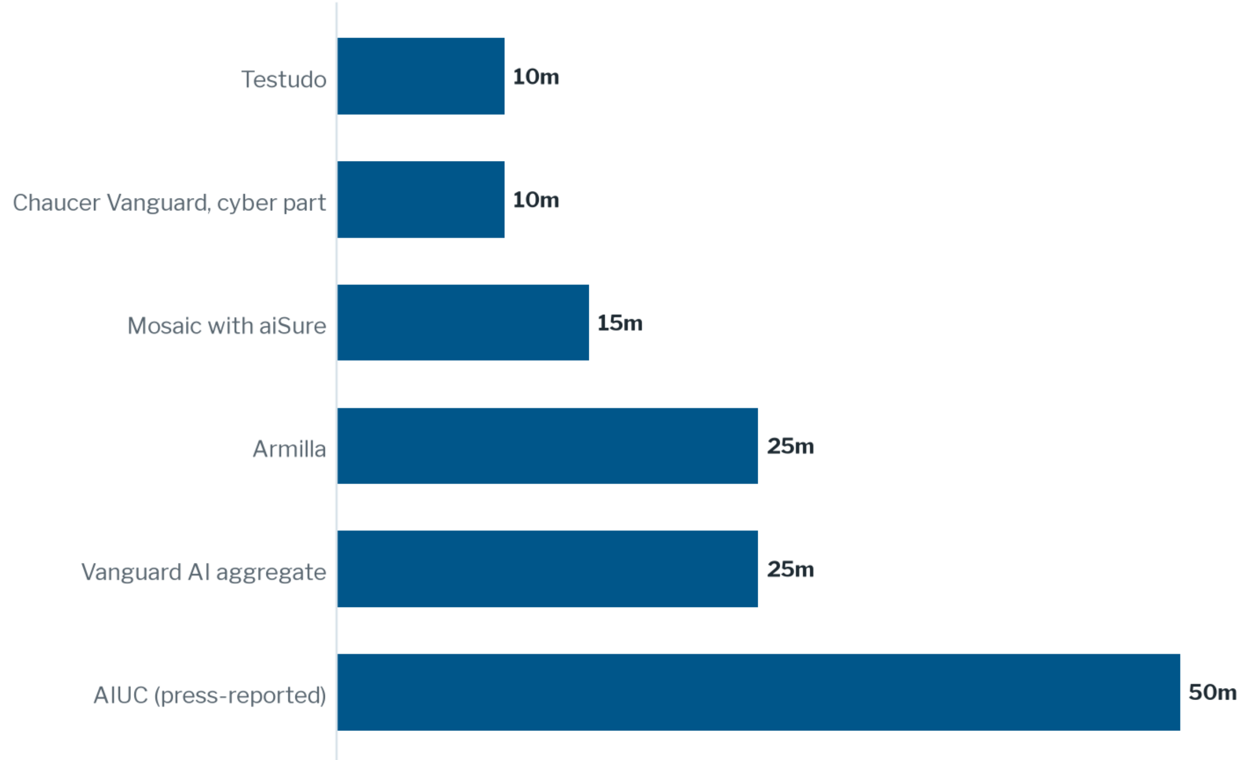

VENDOR CLAIM US dollars, millions. Every figure is a company statement or a journalist's report, since no specimen wording is public. The AIUC figure appears in no AIUC or insured release. Stated maxima are ceilings, not typical bound limits. [21] [24] [25] [29] [10]

Incumbents clarify, specialists insure

VENDOR CLAIM The incumbent carriers mostly confirm that AI-involved events stay inside existing triggers. AXA XL's 2024 endorsement adds data poisoning, usage-rights infringement and AI regulatory cover for firms building their own models. Hiscox UK's 2025 tech PI update affirms cover for algorithmic failure and negligent AI advice. Relm, Vouch and Counterpart market AI-specific E&O without published limits. [31] [32] [33]

FACT The hyperscaler indemnities often cited as a substitute are narrower than insurance. Microsoft's Customer Copyright Commitment defends and pays copyright claims on Azure OpenAI and Copilot output if the customer runs the required mitigations. Google Cloud's generative AI indemnity covers training data and unmodified output. Both cover intellectual property only. Neither covers an agent's operational error. [34] [35]

ASSESSMENT Only a few providers affirmatively cover an agent acting on its own, as opposed to producing bad content. RAND names AIUC and Armilla. Chaucer's structure extends Armilla's cover into a tower. Everyone else addresses outputs, IP or attacker-side AI. The limits in primary sources sit at $10 million to $25 million per insured. The conditions share one idea: the insurer wants its own evidence of how the model behaves, from testing, certification or performance history, before it will write the risk. And because no wording is public, a buyer cannot compare exclusions until a quote arrives.

ASSESSMENT Capacity overlaps. Chaucer, Beazley, Munich Re and Swiss Re each appear behind more than one programme, as lead, as capacity or as reinsurer. Separate brands do not mean separate balance sheets, and a broker assembling a tower from three specialists may be stacking the same reinsurer three times.

AIUC-1: underwriting input or marketing

AIUC-1 is the certification most often described as the gate to agent insurance. The test of whether it changes underwriting terms is documentary. Does the standard say it does, does the auditor attest to anything an underwriter could price, and does any carrier refer to it in a filing?

What the standard says

FACT AIUC-1 is organised in six domains: data and privacy, security, safety, reliability, accountability and society. It was launched on 22 July 2025 and has been revised quarterly since, most recently on 15 July 2026 when requirements for coding agents were added and eight requirements and 41 controls were updated. It is maintained by the Artificial Intelligence Underwriting Company through a consortium. [36] [37]

FACT The changelog calls the standard "insurance-enabling" and says it prioritises risks that lead to direct harms and financial losses. The standard's own pages describe no insurance eligibility, pricing or discount mechanism. [37] [36]

What the audit covers

FACT Schellman became the first accredited AIUC-1 auditor in February 2026. Its role is evidence collection, reporting and certification guidance. The adversarial testing, thousands of scenarios per agent retested at least quarterly, is run by AIUC itself, which also issues the certificate. A certificate lasts twelve months with an annual full audit. [38] [39]

VENDOR CLAIM Certified companies include Intercom (December 2025), ElevenLabs (February 2026) and Sierra (September 2026). ElevenLabs says it passed 5,835 tests across 14 risk categories and obtained an AIUC-1-backed insurance policy, without naming a limit, carrier or premium. Sierra's and Intercom's announcements do not mention insurance. [42] [39] [40]

VENDOR CLAIM The Futurum Group's note on Sierra's certification discloses that it was written by a commercial general-purpose language model. It restates Sierra's release and says nothing about insurance. It is not independent validation. [41]

What carriers reference

FACT No filed form, rate filing, policy wording or underwriting guideline from any carrier that references AIUC-1 was found. An independent academic inventory of AI insurance evidence, which separates carrier marketing from regulatory filings, lists ISO and Berkley forms as filings and codes AIUC only from its own materials and press coverage. [48]

VENDOR CLAIM AIUC's launch coverage stated that "the cost and availability of insurance are tied directly to audit results: safer systems are offered better terms." In a September 2026 interview, co-founder Rune Kvist described Lloyd's underwriters using AIUC-1 "as kind of the underwriting framework", with AIUC's evaluation results feeding pricing. The Insurer reported in May 2026 that AIUC had secured Beazley paper. That article is paywalled, and only its published summary was reviewed. [43] [47] [44]

FACT AIUC raised a $15 million seed round in July 2025 led by Nat Friedman and Daniel Gross, and a $40 million Series A led by Ribbit Capital on 15 September 2026. The company says it will extend its auditing to frontier models. Neither announcement discloses a carrier, a limit or a wording. [43] [45] [46]

ASSESSMENT The answer to the third question is in two parts. Inside AIUC's own programme, AIUC-1 is the underwriting input by construction, because the company that writes the standard also runs the tests, issues the certificate and places the insurance. Outside it, the certificate is a procurement signal and nothing more. No other carrier has put it in a document a regulator or a court could read, and no discount has been disclosed anywhere. The arrangement also concentrates standard-setting, testing, certification and underwriting in one firm. That is not disqualifying, since rating bureaux and inspection bodies have lived with similar tensions, but a buyer should know that the auditor attests to evidence and AIUC attests to behaviour.

HYPOTHESIS If AIUC-1 is an underwriting instrument and not a marketing one, a carrier with no stake in AIUC will reference it in a wording or a questionnaire within the next two renewal cycles. Confirmed if such a reference appears. Refuted if AIUC-certified buyers report the same terms from other markets as uncertified peers.

Accumulation: what carriers have modelled

The fourth question is how carriers model the correlated loss that follows when many insureds depend on the same foundation model or the same cloud. RAND lists five mechanisms: universal attacks that work across models, common model dependency, AI as a force multiplier for attackers, a legal or regulatory shock that makes a common practice actionable, and slow degradation of model behaviour. The evidence below is what has been quantified against those mechanisms, which is very little.

Lloyd's

FACT Lloyd's 2018 Cloud Down scenario, built with AIR Worldwide, estimated that a top-three US cloud provider offline for three to six days would cause $5.3 billion to $19 billion of economic loss and $1.1 billion to $3.5 billion of insured loss. [49]

FACT The mandatory cyber Realistic Disaster Scenarios in Lloyd's January 2025 specification include Cloud Cascade: a major US cloud provider down for 48 hours, with some firms disrupted for up to 20 days. Syndicates compute their own losses and no market figure is published. The specification contains no AI or language-model scenario. [52]

FACT Lloyd's Futureset report on generative AI and cyber names "provider concentration risk" from a few large model providers acting as additional single points of failure, calls this "a new tier of cloud provider", and warns of "fertile ground for large accumulations". It is qualitative and gives no loss figure. Lloyd's 2023 systemic scenario of a payment-system attack put global economic loss at $3.5 trillion over five years, but it is not an AI scenario and gives no insured figure. [51] [50]

Modellers, brokers and reinsurers

THIRD-PARTY ESTIMATE The CrowdStrike update failure of July 2024 is the one real shared-dependency event the market has priced. CyberCube estimated $0.4 billion to $1.5 billion of insured loss, Parametrix implied roughly $0.5 billion to $1.1 billion, and Guy Carpenter $0.3 billion to $1 billion, or $0.6 billion to $2 billion had it been malicious. Fitch cited preliminary market estimates in the mid to high single-digit billions. [53] [54] [55] [56]

FIGURE 3 **Insured-loss estimates for shared-dependency events, low and high**

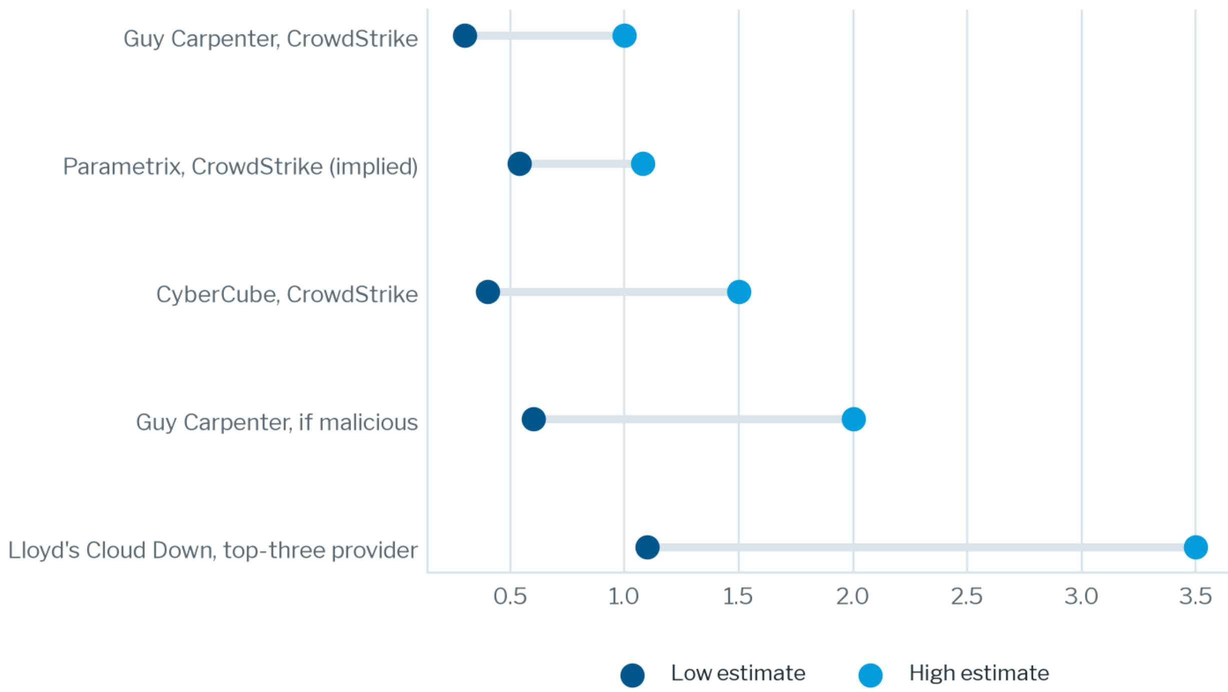

THIRD-PARTY ESTIMATE US dollars, billions. Modeller and broker estimates made within weeks of the July 2024 event, plus Lloyd's 2018 hypothetical scenario. The Parametrix range is derived from its 10 to 20 percent insured share of $5.4 billion. Fitch's market figure in the mid to high single-digit billions is omitted because it gives no range. No comparable estimate exists for a foundation-model failure. [49] [53] [54] [55]

FACT Moody's RMS cyber model version 10, released in May 2026, rebuilt its event set to cover malicious cloud outages and added deterministic catalogues for technology-provider disruption. CyberCube's April 2026 work names GPU compute, cloud and model providers as shared points of exposure. Neither public description names a language-model provider event or publishes a loss figure for one. [57] [58]

VENDOR CLAIM Reinsurers and brokers describe the risk without sizing it. Guy Carpenter wrote in 2024 that compromise or degradation of a third-party model "could result in outages or compromises, making the AI vendor a single point of failure for all customers using that model." Gallagher Re, Swiss Re Institute and Munich Re's 2026 cyber outlook make the same point about concentration across cyber, liability and business interruption. Munich Re's aiSure team has published the only reinsurer-authored quantitative work on correlated model error found, a simulation study with no portfolio loss figure. [59] [60] [61] [62] [63]

VENDOR CLAIM An industry paper led by an AIUC author and co-written with staff from QBE, Generali, RAND, Anthropic and others states that over 80 percent of deployments depend on three foundation-model providers. It says the market currently manages accumulation through "exclusions (e.g., excluding at-scale outages of upstream frontier model providers) and low policy limits in the 10s of millions of USD", and calls for scenarios with probable maximum losses at 1-in-100 and 1-in-250 year return periods. The lead author's employer sells AI insurance. [64]

FACT Regulators name the concentration without sizing it for insurers. The Bank of England's Financial Policy Committee wrote in April 2025 that reliance on a small number of AI providers could lead to systemic risk and that such providers could be designated critical third parties. Moody's is reported to have warned in September 2026 that policies which neither cover nor exclude AI may leave insurers with silent AI liabilities. Only a summary of that note was available. [65] [69]

FACT On 3 September 2026 ChatGPT, Claude, Grok and Gemini suffered overlapping outages within a few hours. No common cause was confirmed and no insurer loss estimate or statement followed. [70]

ASSESSMENT The answer to the fourth question is that carriers are modelling the cloud and handling the model with contract terms. Accumulation from shared clouds has scenario definitions, vendor models and a real loss event to calibrate against. Accumulation from a shared foundation model has a taxonomy and plenty of commentary but no published number. What stands in for modelling is the limit ceiling and the upstream-outage exclusion. The low specialist limits in the previous section are therefore an accumulation control first and a pricing judgment second, and they will not rise much until someone publishes a return-period loss for a model failure.

Filed wording against executive statement

The brief asked for filed wording and executive statement to be kept apart throughout. This table does it for the claims most often repeated about AI insurance.

Claim as it circulates	Best evidence found	What that evidence supports	Evidence type
Standard forms now exclude generative AI	ISO CG 40 48 sample form text	Optional ISO endorsements exist; base form unchanged; adoption unknown	Filed wording
Carriers are filing AI exclusions at scale	RAND count of SERFF forms	112 exclusion forms from 60 companies, admitted market only	Regulator record, via index

Claim as it circulates	Best evidence found	What that evidence supports	Evidence type
Berkley excludes all AI in D&O and E&O	Form text reproduced by RAND and counsel	Absolute wording exists; which accounts receive it is unknown	Carrier wording
AIG, Great American and Berkley seek AI exclusions	FT report and AIG denial	Neither side documented; no form numbers	Press and executive statement
More than 80 percent of exclusion filings approved	Paywalled article citing private analysis	Nothing verifiable	Second-hand
AI agent cover up to $50m	Press and think-tank reports	Not in any AIUC or insured release	Press
AIUC-1 earns better insurance terms	Launch coverage; founder interview	True inside AIUC's programme by design; no other carrier reference	Executive statement
Lloyd's is preparing AI clauses	LMA consultation on a model definition	Definition in progress; no clause; no bulletin	Trade press, paywalled

ASSESSMENT The pattern is consistent. The documents support a narrower story than the headlines. Exclusions are real and multiplying, but they are optional or carrier-specific, and their uptake is unmeasured. Affirmative cover is real, but its terms are announced rather than published. The market-wide retreat that commentary describes has no document behind it yet.

Scenarios for a market-wide exclusion

A point forecast of when AI will be excluded market-wide would be invented. Three scenarios below run to the end of 2027 and carry subjective probabilities. The probabilities are the least reliable content in this report. The tripwires that follow are the most useful, because anyone can watch them without private data.

SCENARIO A Patchwork holds 45 percent

Exclusions keep spreading carrier by carrier through optional ISO endorsements and individual D&O, E&O and umbrella wordings. Cyber keeps adding attacker-side AI clarifications. Specialist MGAs grow, but limits stay in the tens of millions. No bureau or market body mandates anything.

Consequence. Coverage for the same agent loss depends on which carrier wrote which line. Coverage disputes rise and brokers become the main source of clarity.

Earliest visible sign. Another renewal season passes with no Lloyd's bulletin, and ISO files no AI exclusion for umbrella or professional lines.

SCENARIO B Clarity mandate 30 percent

Lloyd's issues a bulletin, on the Y5258 model, requiring every policy to state whether AI loss is covered or excluded. The LMA's model definition becomes model clauses. In the US, the NAIC or a group of states adopts something like RAND's proposed AI Coverage Notice.

Consequence. Silent cover disappears over one or two renewal cycles. Some of it becomes affirmative at a price and some becomes explicit exclusion. Specialists gain because incumbents' silence ends.

Earliest visible sign. The LMA publishes a model AI definition with a clause number, or a Lloyd's bulletin names AI.

SCENARIO C Hard exclusion after a correlated loss **25 percent**

A shared-model or shared-cloud event produces claims across many insureds and several lines at once. Carriers attach absolute AI exclusions broadly at renewal, reinsurers add AI exclusions to treaties, and specialist capacity retreats.

Consequence. The protection gap widens sharply. Agent risk returns to balance sheets and to vendor contracts. Affirmative limits fall before they recover.

Earliest visible sign. A publicly reported insured loss tied to a model failure across unrelated insureds, followed by reinsurer treaty language on AI at a 1 January renewal.

Tripwires for a market-wide exclusion

If this happens	It means	Moves toward	Where to watch
The LMA publishes a model AI definition with a clause number	London has the vocabulary for model exclusions and model grants	B	LMA wordings pages; The Insurer
A Lloyd's market bulletin names artificial intelligence	A mandate is coming; this is the call's overturn condition	B, strongly	Lloyd's market bulletins
ISO files an AI exclusion for umbrella, BOP or professional forms, or amends CG 00 01	The bureau is moving from option to default	C	Verisk announcements; state DOI bulletins
A state regulator disapproves an AI exclusion or the NAIC adopts an AI coverage notice	US regulators will force clarity, not permit silent withdrawal	B	NAIC committee minutes; SERFF objections
A carrier files an AI exclusion on a standalone cyber form	The one line still expanding AI cover starts to close	C	SERFF; carrier cyber form updates
Reinsurance treaties add AI exclusions or AI event definitions	Primary carriers lose the backing for silent AI cover	C, strongly	1 January and 1 July renewal reports
First publicly reported paid claim under an affirmative AI policy	The specialist product works and loss data starts to exist	A or B	Carrier and MGA releases; court filings
A modeller publishes a return-period loss for a foundation-model failure	Accumulation becomes priceable and limits can rise	A or B	Moody's RMS, CyberCube, Verisk releases

If this happens	It means	Moves toward	Where to watch
Specialist per-insured limits in primary releases fall	Capacity is retreating ahead of losses	C	Armilla, Chaucer, Testudo, AIUC releases

Implications

The four audiences face the same facts from different sides of the contract, so each gets its own read.

Risk managers renewing cyber and E&O

ASSESSMENT Ask each carrier, in writing, how its forms define artificial intelligence and whether any AI exclusion is attached, including the ISO CG 40 47 and CG 40 48 endorsements on general liability and umbrella. On cyber, press for a clarifying endorsement of the Beazley or Coalition kind, then check whether it reaches an agent you deployed acting wrongly with no security event, which is the loss most cyber wordings do not reach. On E&O, confirm that an agent executing a task counts as a professional or technology service under the insuring agreement. If D&O renewal brings a Berkley-style clause, read its disclosure limb closely. It removes cover for claims about the company's own statements on AI, which is the securities exposure a board most wants covered.

CFOs pricing uninsured agent deployment risk

ASSESSMENT Assume agent losses are retained unless a named wording says otherwise. The affirmative market puts roughly $10 million to $25 million per insured on the table in primary sources, and silent cover in existing lines is shrinking and contested. Size retained exposure on two cases. One is a single agent failure, such as wrong payments or refunds at volume before detection. The other is a correlated case where your model or cloud provider fails for everyone at once, which is exactly the event insurers are least willing to take. The EY survey of large US listed companies suggests the exposure is real now: respondents report most agentic systems acting without real-time human involvement and about a quarter unable to detect unauthorised agents. [66]

ASSESSMENT Regulation adds to the retained side. Commentators point out that existing obligations, from state privacy law to defence contracting rules, apply to agents without any AI-specific statute. Berkley-style wordings exclude the regulatory costs that follow. A CFO should assume that fines and investigation costs arising from an agent's conduct will be uninsured under excluding forms and sub-limited where specialists cover them. [67] [68] [1]

FIGURE 4 **The governance gap insurers are underwriting against**

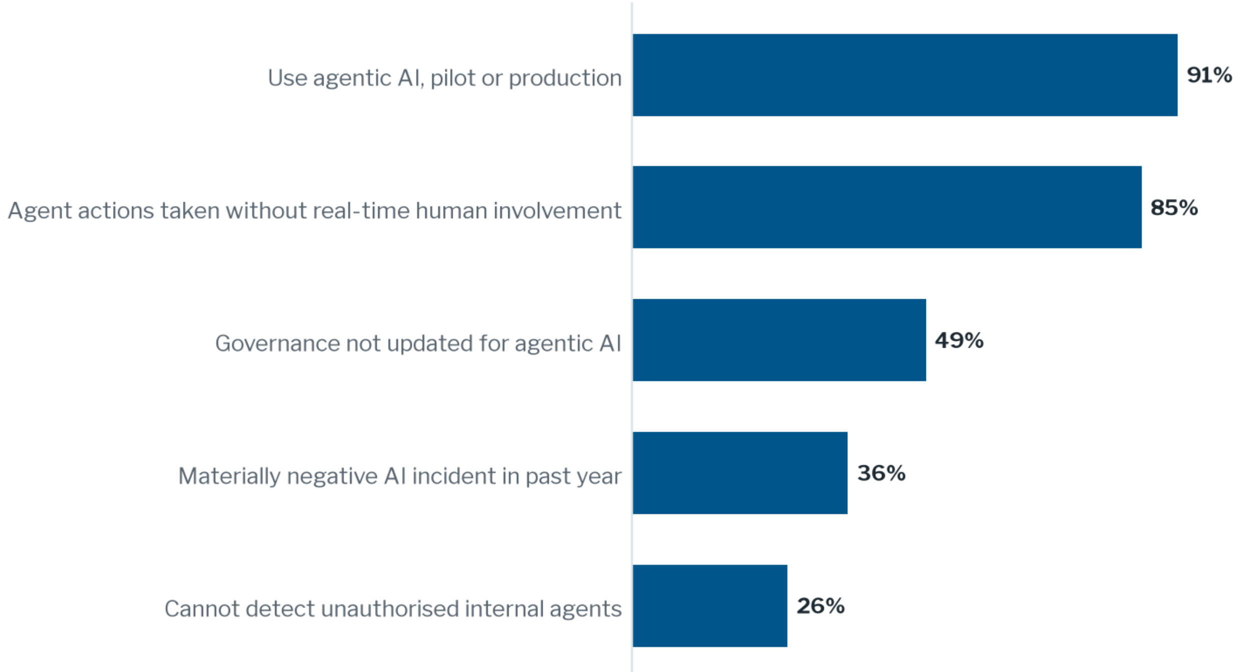

THIRD-PARTY ESTIMATE EY online survey of 202 senior AI decision-makers at US public companies with revenue over $1 billion, fielded 28 May to 15 June 2026, margin of error plus or minus 7 points. Self-reported. The 85 percent figure describes agentic systems, not respondents. [66]

General counsel negotiating AI vendor indemnities

ASSESSMENT Hyperscaler indemnities cover copyright in outputs and training data, on conditions, and nothing operational. So the agent vendor's contract is where operational risk has to land. Ask the vendor to carry affirmative AI cover or a performance guarantee (aiSure and Armilla both sell these to vendors) and to name you as an additional insured or loss payee. Ask for the wording, since none is public. Treat an AIUC-1 certificate as evidence of tested controls for procurement, not as evidence of insurance, and check whether the vendor's own E&O carries an AI exclusion that would leave its indemnity unfunded. Where agents make commitments to customers, remember that the Berkley wording singles out representations by chatbots and virtual agents. Your own liability cover may already be excluding the exposure the vendor contract is meant to backstop. [34] [35] [28] [22]

Brokers placing specialist AI liability cover

ASSESSMENT The product exists but the wordings are private, so the broker's value is getting specimens early and comparing them line by line. Pay particular attention to the definition of AI, the performance trigger, exclusions for misuse and unauthorised changes, and prior-knowledge clauses on known model defects. Other-insurance clauses decide whether a specialist policy sits primary or excess to cyber and E&O. AIUC and Testudo respond primary, and Chaucer's Vanguard structure shows allocation rules can be written. Map the capacity behind each programme before stacking a tower, because the same Lloyd's names and reinsurers recur across brands. Expect limits to stay low until a modeller publishes a return-period figure for a model failure. Pitching larger towers before then is selling capacity that does not exist yet. [1] [24] [22]

What this document cannot answer

Several questions need primary research this report could not do from public sources.

Adoption of the ISO endorsements. Neither Verisk nor RAND knows how many carriers attach CG 40 47 or CG 40 48, or on which accounts. A broker survey of bound GL and umbrella programmes would settle it.

The full filing record. RAND's count ends on 21 June 2026 and relies on a commercial index. A direct SERFF pull by state, with form numbers and approval status, would show whether the surge continued and whether any regulator objected.

Actual bound limits, retentions and prices for specialist AI cover. Only a sample of placed programmes, from brokers or buyers, can show whether the stated maxima are ever bound.

The specimen wordings themselves. Armilla says documentation is available on request, and AIUC, Testudo and aiSure publish none. A side-by-side reading of definitions, triggers and exclusions is the single most useful piece of work a buyer could commission.

Claims experience. No paid claim under an affirmative AI policy has been publicly reported, and no carrier publishes AI loss ratios.

Reinsurance treaty language. Whether cyber and casualty treaties renewed at 1 January and 1 July 2026 added AI exclusions or AI event definitions is not public, and it would decide more about the direction of the market than any primary form.

Half-life of this assessment

Different claims here decay at different rates, and a reader coming back later should discount them accordingly.

Six months: the filing counts, the list of certified AIUC-1 companies, the stated specialist limits and capacity partners, and the status of the LMA definition. All of these are moving monthly. The RAND count was already three months old at publication.

Twelve months: the scenario weights, and the finding that no Lloyd's bulletin exists. One renewal cycle either confirms the patchwork or breaks it. The observation that cyber forms add AI cover while liability forms exclude it should also be rechecked within a year.

Twenty-four months: the reading of AIUC-1 as marketing outside AIUC's programme, which lasts until another carrier puts it in a wording, and the finding that no public model-failure loss estimate exists.

Architectural: the gap between attacker-side AI cover and cover for an agent the insured deployed, the use of low limits and upstream-outage exclusions as the accumulation control, and the concentration of deployments on a few model providers and clouds. These follow from how the lines are built and how the technology is supplied, and they will hold until one of those changes.

Limitations

This report was written independently. David Soden has no commercial relationship with any insurer, MGA, broker, auditor, certification body or AI vendor named in it, and nobody named reviewed it before publication.

The filing evidence is second-hand in one important respect. The count of admitted-market forms comes from RAND's analysis of a commercial SERFF index, not from a direct pull, and RAND's own category totals do not quite reconcile. Of the ISO forms, only CG 40 48 was read in its sample text. CG 40 47 and CG 35 08 are known from consistent secondary descriptions. The Berkley wording is known from reproductions by RAND and by counsel, and its form number is unverified.

The London evidence is thin because the London market files nothing publicly. Several of the most relevant trade-press reports (The Insurer on the LMA definition and on AIUC's Beazley paper, Insurance Insider on the LMA's position, The Information on approval rates) are paywalled, and only their published summaries were available. Nothing in The call rests on them.

Every specialist limit is a company statement or a journalist's report. The AIUC-authored industry paper is useful for its market description but is interested, and it is labelled as vendor material. The Futurum note on Sierra is machine-written, and it is labelled as such. Several AI-generated explainer sites that rank highly in search were read and excluded as evidence. The accumulation evidence rests heavily on secondary reporting of rating-agency and reinsurer research, because the primary texts were paywalled or blocked.

Sources

[1] Romanosky, Sasha, and Celine Robinson. "The Insurability of Artificial Intelligence." RAND Corporation, RR-A5130-1, 16 September 2026. Includes analysis of admitted-market AI forms from SERFF through 21 June 2026 and reproduces the Berkley exclusion. *Independent research, peer reviewed* rand.org

[2] Insurance Services Office. "CG 40 48 01 26 Exclusion - Generative Artificial Intelligence (Coverage B Only)." Sample form, 2025, hosted by ALM. *Rating-bureau form, primary* alm.com

[3] Germond, Nancy. "Verisk to Roll Out New General Liability Exclusions for Generative AI Exposures." Independent Insurance Agents & Brokers of America, 21 October 2025. *Trade association press* independentagent.com

[4] Olshan Frome Wolosky LLP. "The Continued Rise of Generative AI Exclusions in Commercial Insurance Policies." The Legal Intelligencer, 15 May 2026. *Practitioner, law firm* olshanlaw.com

[5] Claims Journal. Report quoting Verisk's vice president of liability on adoption of ISO generative AI exclusions, 20 July 2026. *Independent press* claimsjournal.com

[6] Hunton Andrews Kurth. "The Continued Proliferation of AI Exclusions." Insurance Recovery Blog, 28 May 2025. *Practitioner, law firm* hunton.com

[7] Zelle LLP. "AI Update: The Growing Trend of AI-Related Insurance Policy Exclusions." 31 October 2025. *Practitioner, law firm* zellelaw.com

[8] TechCrunch. "AI is too risky to insure, say people whose job is insuring risk." 23 November 2025. Reports the Financial Times story and AIG's response. *Independent press* techcrunch.com

[9] Bratton, Laura. "Berkshire Hathaway, Chubb Win Approval to Drop AI Insurance Coverage." The Information, April 2026. Paywalled; cites a private Wolfe Research analysis. Not reviewed in full. *Second-hand, unverified* theinformation.com

[10] Allen, Gregory C. "The Insurance Industry's Retreat from AI Threatens to Slow Innovation and Adoption." CSIS, 4 September 2026. *Think tank commentary* csis.org

[11] Lloyd's. Market Bulletin Y5381, "State backed cyber-attack exclusions." 16 August 2022. *Market bulletin, primary* lloyds.com

[12] Norton Rose Fulbright. "Silent AI risks finally make some noise." Regulation Tomorrow, 21 May 2026. Source for the Y5258 silent-cyber precedent. *Practitioner, law firm* regulationtomorrow.com

[13] The Insurer. "LMA developing model AI definition." 8 September 2026. Paywalled; summary only. *Independent press, paywalled* theinsurer.com

[14] Lloyd's Market Association. "Understanding AI exposures: AI loss scenarios survey results." 2025. *Industry body primary* lmalloyds.com

[15] Insurance Insider. "Is the market headed towards a second silent cyber problem with AI?" 2026. Paywalled; summary only. *Independent press, paywalled* insuranceinsider.com

[16] National Association of Insurance Commissioners. "Artificial Intelligence" issue brief, March 2026. *Regulator primary* naic.org

[17] Coalition. "Coalition Adds New Affirmative AI Endorsement to Cyber Policies." Press release, 26 March 2024. *Vendor primary* coalitioninc.com

[18] Insurance Journal. Report on Beazley's AI clarifying endorsement for cyber and tech E&O, 17 September 2026. *Independent press* insurancejournal.com

[19] Rosanes, Mark. "What Beazley's AI cyber endorsement does and doesn't cover." Insurance Business, September 2026. *Independent press* insurancebusinessmag.com

[20] Resultsense. Report on CFC affirmative AI cover in the Lloyd's market, 30 July 2026, summarising Insurance Business UK. *Second-hand, unverified* resultsense.com

[21] Armilla AI. "Armilla launches affirmative AI liability insurance with Lloyd's underwriter Chaucer." Press release, 30 April 2025. *Vendor primary* armilla.ai

[22] Armilla AI. "AI insurance." Product page, accessed September 2026. Lists capacity partners and Armilla Guaranteed. *Vendor primary* armilla.ai

[23] FinTech Global. "Armilla AI raises $25m to expand AI liability coverage." 23 January 2026. *Independent press* fintech.global

[24] Chaucer Group. "Chaucer and Armilla AI launch Vanguard AI coordinated insurance structure." Press release, 10 February 2026. *Vendor primary* chaucergroup.com

[25] Testudo. "Testudo launches new insurance coverage for liability risks created by generative AI systems." 21 January 2026. *Vendor primary* testudo.co

[26] Testudo. "Lloyd's syndicates commit more capacity to generative AI liability insurance." 26 February 2026, updated 27 August 2026. *Vendor primary* testudo.co

[27] Munich Re. "aiSure: More AI opportunity, less AI risk." Product page, accessed September 2026. *Vendor primary* munichre.com

[28] Munich Re. "aiSure FAQ." Accessed September 2026. *Vendor primary* munichre.com

[29] Mosaic Insurance. "Mosaic partners with Munich Re's aiSure to provide pioneering coverage for AI vendors." Press release, 26 February 2026. *Vendor primary* mosaicinsurance.com

[30] HSB. "Introducing AI liability insurance for small businesses." Press release, 18 March 2026. *Vendor primary* munichre.com

[31] AXA XL. "AXA XL unveils new cyber insurance extending coverage to help businesses manage emerging Gen AI risks." Press release, 21 October 2024. *Vendor primary* axaxl.com

[32] Insurance Business UK. "Hiscox launches first affirmative AI liability cover in UK market." 3 July 2025. *Independent press* insurancebusinessmag.com

[33] Relm Insurance. "Relm Insurance launches AI liability solutions." Press release, 14 January 2025. *Vendor primary* prnewswire.com

[34] Microsoft. "Customer Copyright Commitment required mitigations." Azure AI Foundry documentation, accessed September 2026. *Vendor primary, contract terms* learn.microsoft.com

[35] Google Cloud. "Service Specific Terms," generative AI indemnified services, archived 18 August 2025. *Vendor primary, contract terms* cloud.google.com

[36] AIUC-1. "About the standard." Accessed September 2026. *Standard text, vendor-maintained* aiuc-1.com

[37] AIUC-1. "Changelog." Accessed September 2026. *Standard text, vendor-maintained* aiuc-1.com

[38] Schellman. "Schellman becomes the first accredited auditor for AIUC-1." 3 February 2026. *Auditor primary* schellman.com

[39] ElevenLabs. AIUC-1 certification and insurance announcement, 12 February 2026. *Vendor primary* elevenlabs.io

[40] Sierra. "Sierra achieves AIUC-1 certification." 17 September 2026. *Vendor primary* sierra.ai

[41] The Futurum Group. "Sierra's AIUC-1 Certification Sets a New Bar for Agentic AI Trust." 18 September 2026. Disclosed as written by a language model. *Analyst firm, machine-written* futurumgroup.com

[42] Intercom. "Intercom achieves AIUC-1 certification." 8 December 2025. *Vendor primary* intercom.com

[43] Insurance Journal. Report on the launch of the Artificial Intelligence Underwriting Company and its seed round, 25 July 2025. *Independent press* insurancejournal.com

[44] The Insurer. "Exclusive: AI insurance MGA AIUC secures Beazley paper for liability product." 15 May 2026. Paywalled; summary only. *Independent press, paywalled* theinsurer.com

[45] AIUC. "AIUC raises $40M Series A from Ribbit and First Harmonic." Press release, 15 September 2026. *Vendor primary* prnewswire.com

[46] Riley, Duncan. "AI agent certification startup AIUC raises $40M to begin auditing frontier models." SiliconANGLE, 15 September 2026. *Independent press* siliconangle.com

[47] Latent Space. Interview with AIUC co-founder Rune Kvist, 16 September 2026. *Executive statement, interview* latent.space

[48] Leung et al. "The Insurability Frontier of AI Risk: Mapping Threats to Affirmative Coverage, Silent Exposures, and Exclusions." arXiv 2605.18784, version 2, 12 June 2026. *Independent research, preprint* arxiv.org

[49] Lloyd's and AIR Worldwide. "Cloud Down: Impacts on the US economy." 23 January 2018. *Market scenario publication* lloyds.com

[50] Lloyd's. "Lloyd's systemic risk scenario reveals global economy exposed to $3.5trn from major cyber attack." Press release, 18 October 2023. *Market scenario publication* lloyds.com

[51] Lloyd's Futureset. "Generative AI: Transforming the cyber landscape." 2024. *Market research publication* lloyds.com

[52] Lloyd's. "Realistic Disaster Scenarios: Scenario Specification." January 2025. *Market regulatory document* lloyds.com

[53] MSSP Alert. "CrowdStrike outage could cost cyber insurers $1.5 billion: CyberCube." July 2024. *Independent press, modeller estimate* msspalert.com

[54] Insurance Journal. Report of Parametrix estimate of CrowdStrike losses, 24 July 2024. *Independent press, modeller estimate* insurancejournal.com

[55] Insurance Journal. Report of Guy Carpenter estimate of CrowdStrike insured losses, 2 August 2024. *Independent press, broker estimate* insurancejournal.com

[56] The Insurer. "Fitch: reinsurers could withstand CrowdStrike loss in mid to high single-digit billions." 22 July 2024. *Rating agency, via press* theinsurer.com

[57] Moody's. "Introducing Moody's RMS Cyber Solutions version 10." 22 May 2026. *Modeller primary* moodys.com

[58] Evans, Steve. "AI treated as force-multiplier for cyber losses, introduces aggregation, correlation risks: CyberCube." Artemis, 14 April 2026. *Independent press, modeller research* artemis.bm

[59] Insurance Business. "Guy Carpenter warns AI adoption brings new cyber aggregation risks." 26 August 2024. *Broker commentary, via press* insurancebusinessmag.com

[60] Insurance Business. "Gallagher Re identifies systemic risk from AI model failures." 25 March 2026. *Reinsurance broker commentary, via press* insurancebusinessmag.com

[61] Global Reinsurance. "Swiss Re: AI boom reshapes risk but leaves insurers exposed." January 2026, on sigma insights 01/2026. *Reinsurer research, via press* globalreinsurance.com

[62] Munich Re. "Cyber insurance: Risks and trends 2026." 25 March 2026. *Reinsurer primary* munichre.com

[63] Munich Re aiSure. "Accumulation Risk in AI Systems via Correlated Model Errors." White paper, IEEE ICSTW 2025. *Reinsurer research* munichre.com

[64] Trout, Cristian, et al. "Underwriting the Agent Economy: The Blueprint for an AI Insurance Stack." 1 July 2026. Lead author affiliated with AIUC. *Industry paper, commercial interest* underwriting-agents.com

[65] Bank of England Financial Policy Committee. "Financial Stability in Focus: Artificial intelligence in the financial system." April 2025. *Regulator primary* bankofengland.co.uk

[66] EY. "EY survey finds that autonomous AI implementation outpaces oversight, yielding an AI governance gap." 15 September 2026. 202 respondents, fielded 28 May to 15 June 2026. *Survey, proprietary method* [ey.com](ey.com)

[67] PYMNTS. "Is the enterprise insurance market agentic AI's hidden ceiling?" 18 September 2026. *Independent press* pymnts.com

[68] Goud, Rick. "Agentic AI doesn't get a governance exemption." Cybersecurity Insiders, 22 September 2026. Author is a vendor executive. *Vendor-authored commentary* cybersecurity-insiders.com

[69] The Insurer. Viewpoint on Moody's insurance emerging risk radar, 6 September 2026. Paywalled; summary only. *Rating agency, via press, paywalled* theinsurer.com

[70] Slashdot, citing Ars Technica. "Four major AI models suffer rare overlapping downtime." 3 September 2026. *Independent press, aggregated* slashdot.org

Rights and permitted use

Commissioned Market Assessment reports, commercial licensing, and briefings on this material are available.

This document is analysis, not investment, legal, or procurement advice. It was not commissioned, reviewed, or funded by any company named in it, and the author holds no position in any of them.

Market Assessment reports, commissioned research, and briefings: **davidsoden.com/market-assessment**