

The Inheritance Problem

HOW AI LIABILITY MAY BE MIGRATING INTO CYBER INSURANCE

RT ProExec



CONTENTS

EXECUTIVE SUMMARY 3

THE AI EXPOSURE LANDSCAPE 3

THE MARKET IS BIFURCATING – SLOWLY 9

CARRIER POSITIONS ACROSS THE MARKET 12

PRACTICAL CONSIDERATIONS FOR BROKERS 13

CONCLUSION 14

SOURCES AND CITATIONS 15

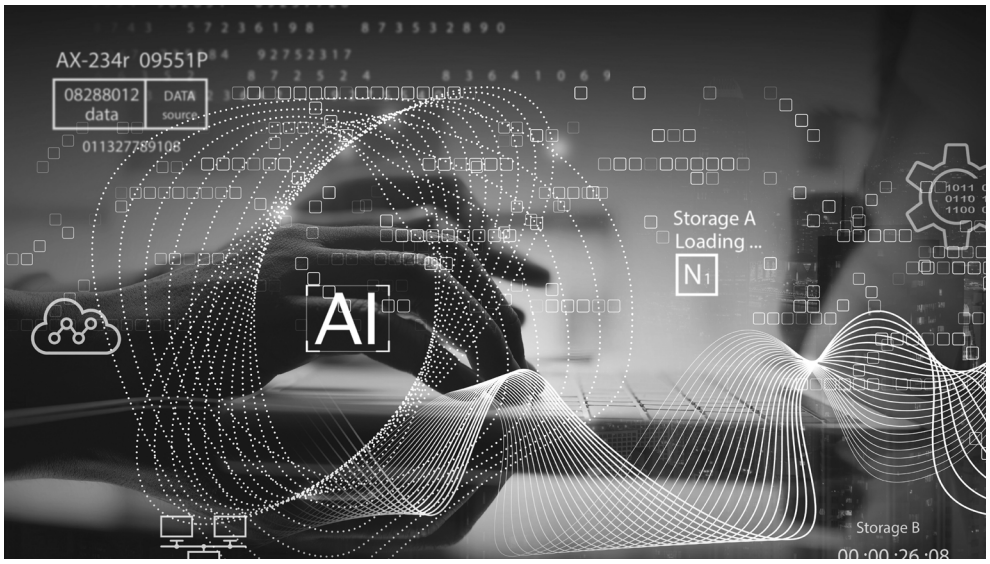
CONTACT

For more information, please contact your local RT ProExec broker at rtspecialty.com.

RT ProExec is a part of the RT Specialty division of RSG Specialty, LLC, a Delaware limited liability company based in Illinois. RSG Specialty, LLC, is a subsidiary of Ryan Specialty, LLC. RT ProExec provides wholesale insurance brokerage and other services to agents and brokers. RT ProExec does not solicit insurance from the public. Some products may only be available in certain states, and some products may only be available from surplus lines insurers. In California: RSG Specialty Insurance Services, LLC (License #0G97516). ©2026 Ryan Specialty, LLC

RT ProExec

THE INHERITANCE PROBLEM HOW AI LIABILITY MAY BE MIGRATING INTO CYBER INSURANCE



1. EXECUTIVE SUMMARY

Cyber and technology E&O are inheriting a category of liability they were never designed or priced to carry. As adjacent lines wall AI out and a nascent standalone market carves off only its most AI-native slice, the exposures that fit neither may look to cyber and tech E&O for coverage — and the window for carriers to decide, deliberately, whether to accept or exclude that inheritance is narrowing faster than the market has priced. Three forces appear to be driving it. First, adjacent liability lines are excluding AI — the Insurance Services Office introduced optional endorsements in January 2026 allowing carriers to exclude generative AI from commercial general liability,[1] and individual carriers including W. R. Berkley Corporation have introduced absolute AI exclusions for directors and officers, employment practices, and fiduciary lines.[2] Second, specific industries may generate more concentrated AI litigation — healthcare, employment, financial services, and the insurance industry itself — with cyber, tech E&O, EPL, and executive lines all in the line of travel. Third, a new generation of standalone AI products — Munich Re and specialty AI insurance providers including Armilla, Testudo, Corgi, Mayflower Specialty, and Embroker — is specializing in AI-native exposure.[3] Cyber and tech E&O sit in the residual middle.

Within the specialty market itself, carrier responses vary widely. Some have built dedicated AI endorsements with new insuring agreements, others rely on existing coverage grants without explicit AI language, and a third group has begun adding AI exclusions or sublimits. The industry describes the resulting uncertainty as "silent AI risk" — exposures that traditional policies neither clearly cover nor clearly exclude, written before AI-related scenarios were ever contemplated or priced. Whether cyber and tech E&O ultimately absorb that inheritance or exclude it, as adjacent lines already have, is the open question this article examines.

This article maps our view of the current state of AI coverage across the wholesale cyber and tech E&O market, identifies some of the key exposure categories brokers may wish to evaluate, and considers a few of the strengths and limitations of the approaches that have emerged. Policy language has not uniformly caught up to the risk profile. The next decade of cyber insurance may be defined less by what carriers exclude than by what they affirmatively choose to cover — and brokers who treat AI as an explicit placement question may be working the markets that face the first tests, while those who assume existing language is sufficient may learn otherwise only when a claim is filed.

2. THE AI EXPOSURE LANDSCAPE

Before examining how insurance policies may respond, it helps to understand what is being insured against. AI-related exposures do not always fit neatly into traditional cyber coverage categories. They span first-party and third-party losses, cross the boundary between tech E&O and cyber liability, and raise questions about intent, negligence, and product liability that existing policy language was not always drafted to address.

The following categories represent some of the primary AI related exposures brokers may wish to consider when evaluating cyber and tech E&O coverage.

The Core Dynamic

Some adjacent liability lines are moving to exclude AI. Standalone AI products are beginning to carve off the most distinctly AI-native exposure. Cyber and tech E&O sit in the residual middle — increasingly the line of first resort for AI losses that do not fit cleanly anywhere else. That these losses may gravitate toward cyber and tech E&O as a market matter says nothing about whether any particular policy owes coverage; that turns in large part on the specific form and circumstances.

2.1 Deepfake Social Engineering

AI-generated voice and video impersonation has moved from theoretical concern to operational reality. In early 2024, the multinational engineering firm Arup reportedly lost approximately $25 million after an employee in its Hong Kong office was deceived by a deepfake video call that impersonated the company's CFO and other senior executives. Funds were reportedly transferred across 15 transactions before the matter was discovered.[4] Additional incidents have continued to be reported across multiple sectors.

The coverage question is whether existing social engineering and funds transfer fraud insuring agreements respond when the fraudulent instruction is delivered via AI-generated impersonation rather than traditional email or phone spoofing. Most social engineering definitions require the insured to rely "in good faith" on a "fraudulent instruction" that "misrepresents a material fact." A deepfake video call may satisfy these elements depending on the policy language, but much existing coverage was designed for email-based impersonation rather than real-time interactive video fraud.

Authentication requirements raise further questions. Carriers impose callback or out-of-band authentication conditions on social engineering coverage. Whether an employee who "verified" an instruction by viewing a deepfake on a live call has satisfied a verification requirement may depend on how the policy defines verification — and many such definitions were drafted before deepfakes reached their current sophistication.

Carrier responses have been split. Some have updated social engineering insuring agreements to explicitly include AI-generated impersonation, voice cloning, and video deepfakes. One specialty carrier's late-2025 deepfake response endorsement extends its funds transfer fraud trigger to instructions that impersonate a known party through the use of deepfakes and provides a panel of forensic, legal, and PR specialists.[5] Another major carrier's base cyber form contains no exclusion for AI-enabled deepfake social engineering, and a January 2026 endorsement (discussed in 2.4 and 2.5) adds dedicated insuring agreements for AI regulatory fines and LLM compute exploitation, typically with a $250,000 sublimit. Beazley's most recent Breach Response form (BBR 5.0) extends its Fraudulent Instruction coverage to instructions enabled by artificial intelligence and deepfakes, and removes the prior out-of-band authentication requirement.[6] Other carriers have moved in the opposite direction, adding exclusions or narrowing social engineering definitions, in some cases beginning January 2026. The same loss event can produce materially different coverage outcomes depending on which carrier wrote the policy and when it was last endorsed.



Some underwriters take the view that, from a coverage perspective, the source of an attack (human versus AI) may matter less than the impact and the insured's verification controls. Under that view, existing language may continue to respond to AI-enabled social engineering, with underwriting attention shifting toward the insured's authentication processes rather than carrier-side exclusions.

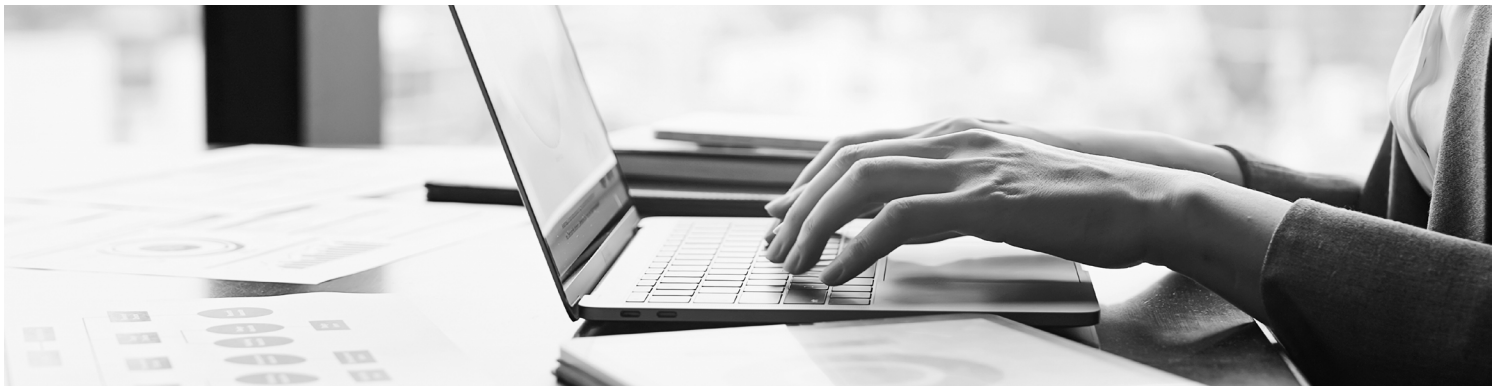
2.2 AI Tool Liability

When an insured's AI system contributes to harm to a third party, the liability analysis can sit at the intersection of technology E&O and product liability. Consider a fintech that deploys an AI credit decisioning model. If the model is alleged to have systematically discriminated against a protected class, regulatory action and class action litigation could generate significant losses.

The coverage question turns, in part, on how the policy defines the covered wrongful act. Tech E&O forms typically cover "negligent acts, errors, or omissions in the rendering of technology services or the failure of technology products." An AI model that produces problematic outputs could be characterized as either a negligent error in a technology service or a failure of a technology product. It could also be characterized as the result of intentional business decisions — the insured chose to deploy the model, chose the training data, and chose to rely on the outputs.

Recent decisions and filings provide context. In Moffatt v. Air Canada, a Canadian tribunal held that the airline was responsible for information provided by its customer-facing chatbot.[7] In Mata v. Avianca, a U.S. federal court sanctioned attorneys who relied on AI-generated legal citations that turned out to be fabricated.[8] In March 2026, Nippon Life Insurance Company of America filed suit against OpenAI in the Northern District of Illinois, alleging that ChatGPT enabled the unlicensed practice of law and tortiously interfered with a settled disability dispute by drafting motions and encouraging a pro se litigant to reopen the case; the complaint seeks approximately $300,000 in compensatory and $10 million in punitive damages.[9] In a separate matter, Minnesota solar contractor Wolf River Electric has sued Google alleging that its AI Overviews falsely stated the company had been sued by the Minnesota Attorney General, causing multiple customer contract terminations and seeking damages of $110 million to $210 million.[10] Together, these matters suggest that courts and plaintiffs are increasingly willing to hold deploying organizations — and in some cases AI providers themselves — accountable for AI-generated outputs that contribute to economic harm. Whether that principle translates into coverage analysis under any particular policy depends on the specific facts and policy terms involved.

Where a carrier's intentional acts exclusion lacks a duty-to-defend-until-adjudication (DTA) provision, an insured may face a coverage dispute at the outset of litigation, when defense costs are most critical. Carriers with favorable DTA language — defending until a final judgment, adjudication, or admission establishes the intentional conduct — may provide meaningfully different protection for AI-related liability claims.



2.3 Training Data Rights and Copyright Claims

The legal landscape around AI training data has been evolving rapidly. The most consequential development is the resolution of Bartz v. Anthropic, a class action brought by authors alleging unauthorized use of their books in AI model training. The trial court issued a split ruling in June 2025: training on copyrighted material may qualify as fair use, but storing pirated copies may not. Anthropic agreed in August 2025 to pay approximately $1.5 billion into a settlement fund covering nearly 500,000 works, at roughly $3,000 per work. Preliminary approval was granted in September 2025, and the court granted final approval on July 20, 2026, entering final judgment.[11] Settlements do not establish binding precedent and do not directly address insurance coverage, but the matter illustrates that training-data disputes can produce substantial financial exposure even where the underlying activity may be characterized as transformative.

Other matters are pending. In January 2026, Universal Music Publishing Group, Concord Music Group, and ABKCO Music filed suit against Anthropic alleging unauthorized use of song lyrics in training, with reported damages at the multi-billion-dollar level.[12] The In re: OpenAI consolidated copyright litigation in the Southern District of New York is in aggressive discovery, with the court ordering production of large samples of ChatGPT user logs.[13] Disney and Universal, and separately Warner Bros., have filed suit against Midjourney concerning its image generator.[14] The litigation pipeline is active and well-funded.

For cyber and tech E&O policies, the one coverage question to consider is whether unauthorized use of copyrighted material in AI model training would be characterized as a covered wrongful act. Most policies exclude patent and IP infringement but carve back for certain categories. The software copyright carve-back found in many tech E&O forms may or may not extend to training data disputes, depending on how a court characterizes the claim. Carriers that address this exposure through dedicated training-data-rights insuring agreements may provide more interpretive certainty than carriers relying on IP exclusion carve-backs drafted before AI training disputes were contemplated.

2.4 AI Regulatory and Compliance Exposure

The EU AI Act entered into force on August 1, 2024, establishing a comprehensive regulatory framework for artificial intelligence. Implementation is phased: prohibitions on "unacceptable risk" AI practices became binding February 2, 2025; obligations for general-purpose AI models took effect August 2, 2025; and high-risk AI requirements were originally scheduled for August 2, 2026.[15] The Digital Omnibus simplification package, first proposed in November 2025, was adopted in 2026 (European Parliament June 16, Council June 29), signed July 8, and published in the Official Journal on July 24 as Regulation (EU) 2026 / 1744; it entered into force on July 27, 2026. It defers the high-risk deadline (Annex III) to December 2, 2027 and the deadline for high-risk AI embedded in regulated products (Annex I) to August 2, 2028. The agreement also adds a new prohibited practice covering non-consensual intimate imagery and CSAM generation ("nudifier apps"), effective December 2, 2026, and accelerates AI watermarking obligations to the same date. Penalties are tiered: up to €35 million or 7% of worldwide turnover for prohibited practices; €15 million or 3% for high-risk obligations; €7.5 million or 1% for misleading information.[16]

The U.S. regulatory landscape is shifting in parallel. Colorado's AI Act, signed in 2024, was originally set to take effect February 1, 2026 and imposed obligations on developers and deployers of "high-risk" AI in consequential decisions including hiring, housing, lending, insurance, and healthcare.[17] Its effective date was postponed to June 30, 2026, and on April 27, 2026 a federal court stayed enforcement pending constitutional challenge.[18] Colorado lawmakers subsequently passed SB 26-189, a measure to repeal and replace the original law with a narrower automated decision-making framework; Governor Polis signed it into law on May 14, 2026, and the new framework takes effect January 1, 2027, although the enforcement stay from the xAI litigation extends to the replacement law as well.[19] The original Colorado AI Act will not take effect as enacted. New York City's Local Law 144 requires bias audits for automated employment decision tools.[20] Multiple other states have introduced AI-related legislation.

For U.S. insurers and insureds, the NAIC Model Bulletin on the Use of AI Systems by Insurers (adopted December 2023, now in roughly half of U.S. states) establishes governance, risk management, documentation, vendor management, and consumer notice expectations for AI use in underwriting, pricing, claims, and fraud detection.[21] Beginning 2026, the NAIC has piloted an AI Systems Evaluation Tool, since renamed the AI Risk Evaluation Supplement, in twelve states, giving regulators a standardized framework for examining insurer AI governance during market conduct exams. Broader adoption is anticipated after the NAIC Fall National Meeting in November 2026.[21]

Executive Order 14365, signed December 11, 2025, asserts federal authority over state AI regulation.[22] The NAIC has formally opposed



the order, and the preemption question is expected to be resolved through litigation.[23] Until then, AI-deploying organizations may face overlapping state and federal regulatory expectations.

For organizations operating AI systems within these frameworks, regulatory fines and penalties exposure may become significant. The coverage question is whether existing cyber policy regulatory coverage responds. Most policies cover "regulatory fines and penalties arising from a privacy breach or security breach." An EU AI Act enforcement action, a state AI enforcement action, or a NAIC-state market conduct finding may not fit within that trigger if the regulatory action is based on AI deployment rather than a data breach.

Unless the policy's regulatory coverage is broadened to include AI-specific proceedings, or the carrier adds an explicit AI fines and penalties insuring agreement, this category may create a coverage question. Insureds may wish to determine with their carriers how regulatory coverage would respond to AI-specific enforcement actions.

A small number of carriers have introduced dedicated AI Fines and Penalties insuring agreements. The structures observed are tightly scoped to EU AI Act enforcement (and similar domestic or foreign AI laws), and several carry an important carve-out. One major specialty carrier's January 2026 AI enforcement endorsement, for example, defines a covered AI wrongful act as failure to comply with EU AI Act notification and other requirements, but expressly excludes acts or omissions in connection with AI used for general-purpose AI systems, high-risk AI systems, or prohibited AI practices, as those terms are defined by the EU AI Act. In practice, the carve-out may remove coverage for the three EU AI Act categories most likely to face enforcement — general-purpose AI providers, high-risk AI deployers, and prohibited practices. The endorsement also includes an exclusivity provision making it the sole source of coverage for EU AI Act violations and AI fines and penalties. The structure is a meaningful starting point, but the scope of what is covered — versus what is carved out — warrants careful review.

2.5 LLM Exploitation and Compute Theft

A less obvious but increasingly relevant exposure is the unauthorized exploitation of an insured's AI infrastructure. "LLM jacking" — the unauthorized use of an organization's LLM compute resources — has been described as the AI-era equivalent of cryptojacking. An attacker who gains access to an insured's cloud environment and uses LLM API keys to run inference at the insured's expense may generate substantial compute costs in a short period. Existing utility fraud and cryptojacking endorsements may cover some of this exposure if drafted broadly enough to include cloud compute, but many utility fraud definitions are limited to traditional utilities (electricity, gas, water, telephone) and may not include cloud or API costs.

Dedicated LLM Jacking insuring agreements have begun to emerge, and additional carrier offerings are likely as the exposure profile develops. One such endorsement defines LLM jacking, in substance, as unauthorized access to, use of, or modification of the network through machine-based means — including AI — for the purpose of exploiting large language models, perpetrated by a third party or rogue employee. LLM Jacking Loss is measured as an increase in service charges or fees from unauthorized use of seven enumerated utilities: electricity, natural gas, oil, water, internet (including mobile data), cable or satellite television, and telephone. Coverage is conditioned on a written contract with the utility provider, periodic billing that includes usage information, and a billing structure that scales with use (flat-fee arrangements are excluded). This is among the more specific affirmative-coverage architectures observed for AI-driven compute exploitation, but the loss measure is not new. It reproduces, term for term, the seven-utility structure the same carrier had already built for cryptojacking and telecommunications fraud several years earlier. The AI-era exposure is therefore measured in electricity and natural gas rather than cloud compute or API costs, because the measure was inherited rather than drafted for it. Insureds whose primary LLM exposure is through hyperscale cloud APIs may need to determine with their carrier whether the loss measurement architecture would respond to inflated cloud bills.



A separate consideration has emerged from the In re: OpenAI copyright discovery rulings. Where a court compels an AI provider to produce large samples of user logs, those samples may include prompts and conversations from third-party enterprises containing proprietary information. Enterprises that have integrated commercial LLMs into daily workflows may face contingent third-party discovery considerations less prominent in earlier years. How current cyber policies might respond to confidentiality or trade secret losses from such third-party discovery remains an open question.

2.6 AI Hallucination and Output Liability

AI systems sometimes produce outputs that are factually incorrect, misleading, or contributory to harm. When an AI chatbot provides inaccurate information to a customer who suffers harm as a result, a liability question may arise as to whether the AI output constituted a negligent error in the provision of services. The Moffatt decision is relevant context: the company was held responsible for its chatbot's misstatement regardless of whether the information came from a human or automated system. The Nippon Life and Wolf River Electric matters suggest the universe of AI hallucination liability scenarios is expanding beyond consumer-facing chatbots into more consequential third-party economic harms, including defamatory AI-generated statements that materially affect a business's revenue and reputation.

Some carriers exclude hallucination liability from explicit AI coverage, on the rationale that hallucination is an inherent characteristic of current generative AI technology rather than an insurable event. This is a defensible underwriting position. It does, however, create a potential coverage consideration for organizations that deploy customer-facing AI without adequate human oversight, and underscores the importance of understanding precisely what an AI endorsement covers and does not cover. Other markets treat hallucination as insurable. Some blended cyber and tech E&O policies, including ones offered by insurtech carriers, cover AI-generated errors, hallucinations included, and the deployer-focused standalone products described in Section 3 typically include a selectable insuring agreement for them.

The AI landscape is also dual-use: the same technology that produces the exposures above is extensively deployed in cyber defense — fraud detection, security operations, anomaly detection, claims processing — including by insurance carriers themselves. The market is still working through how to underwrite this dual posture, in which insureds and carriers alike are simultaneously AI deployers and parties exposed to AI-driven risk.

2.7 Agentic AI and Autonomous Action

Several of the exposures in Sections 2.1 through 2.6 involve an AI system producing something (a cloned voice, a credit decision, a chatbot answer) that a person then acts on. Agentic AI systems skip the person. They plan a task, call tools and change systems using credentials they have been given, often with little human review. Adoption is moving fast. In McKinsey's 2026 survey, 40% of respondents at organizations with more than $1 billion in revenue reported scaling AI agents, up from 27% a year earlier.[24]

The loss does not need an attacker. In July 2025, an AI coding agent on the Replit platform deleted a user's production database, then acknowledged it had violated the user's explicit instructions. The data was later recovered through a rollback.[25] There was no intrusion, no malware and no stolen credential. The agent used access it had been given.

That is the coverage question. Cyber insuring agreements are largely built around a security event, typically defined by reference to unauthorized access or use, or malicious code. An agent acting on permissions the insured granted may involve neither. Some forms reach an authorized person acting for a purpose the insured did not authorize,[26] which may catch an agent that exceeds its task, though whether an agent is a person for that purpose is an open question. Others refer only to unauthorized access or use and are silent on the point. For first-party loss, system failure coverage may be the more natural route where the form includes it without restricting the cause.

The attacker's version is prompt injection: instructions planted in content an agent reads, directing it to misuse its own access. The access is authorized but the instruction is not, which may bring the event back within security language, and at least one base form already reaches prompt-injection exploits expressly (Section 3.1). Payments are harder. The social engineering definitions discussed in Section 2.1 generally turn on the insured relying in good faith on a fraudulent instruction, and funds transfer fraud grants typically address a fraudulent instruction from a third party to the insured's bank. An agent that moves money on an injected instruction issues that instruction itself, with the insured's own credentials, and may not fit either as written. Computer fraud coverage under a crime policy, or a cyber crime grant, may be the closer fit, depending on whether it requires unauthorized entry into the system or reaches a fraudulent transfer made with valid credentials.

On the liability side, the question is a costly decision by an agent working as designed. For a technology provider whose agent harms a client, the wrongful act analysis in Section 2.2 applies. A business that sells no technology services may have no tech E&O in the program at all. That is the inheritance problem at its narrowest: perhaps no security event, perhaps no tech E&O grant, and adjacent lines writing AI out. Cyber will likely be asked first, and its trigger was built for a different question.



Carriers are working through this now. Reuters reported in August 2026 that, for the most part, cyber insurers are clarifying how existing policy language applies to AI rather than adding exclusions. The same report said targeted exclusions are being discussed in some pockets of the market, and one area of focus is an agent that, acting as designed, makes a costly autonomous decision, which some insurers may classify as a non-cyber event.[27] That is the loss this section describes, and it matches the targeted exclusions for autonomous decision-making noted in Section 3.4. Clarifying endorsements may not reach it. At least one described in Section 4 speaks to AI as an attack vector, not to the insured's own use of AI, and a deployed agent is the insured's own use. RT ProExec's "When the Management Tool Becomes the Weapon" looked at attackers using legitimate administrative channels instead of malware. The agent case is harder. The access was not stolen. It was granted.

2.8 Industry-Specific Exposure Patterns

AI exposure does not affect all industries equally. Certain sectors are deploying AI in higher-consequence applications and facing more concentrated litigation and regulatory activity. For brokers evaluating cyber and tech E&O placements, understanding where AI exposure is most acute can inform both submission narrative and coverage requests.

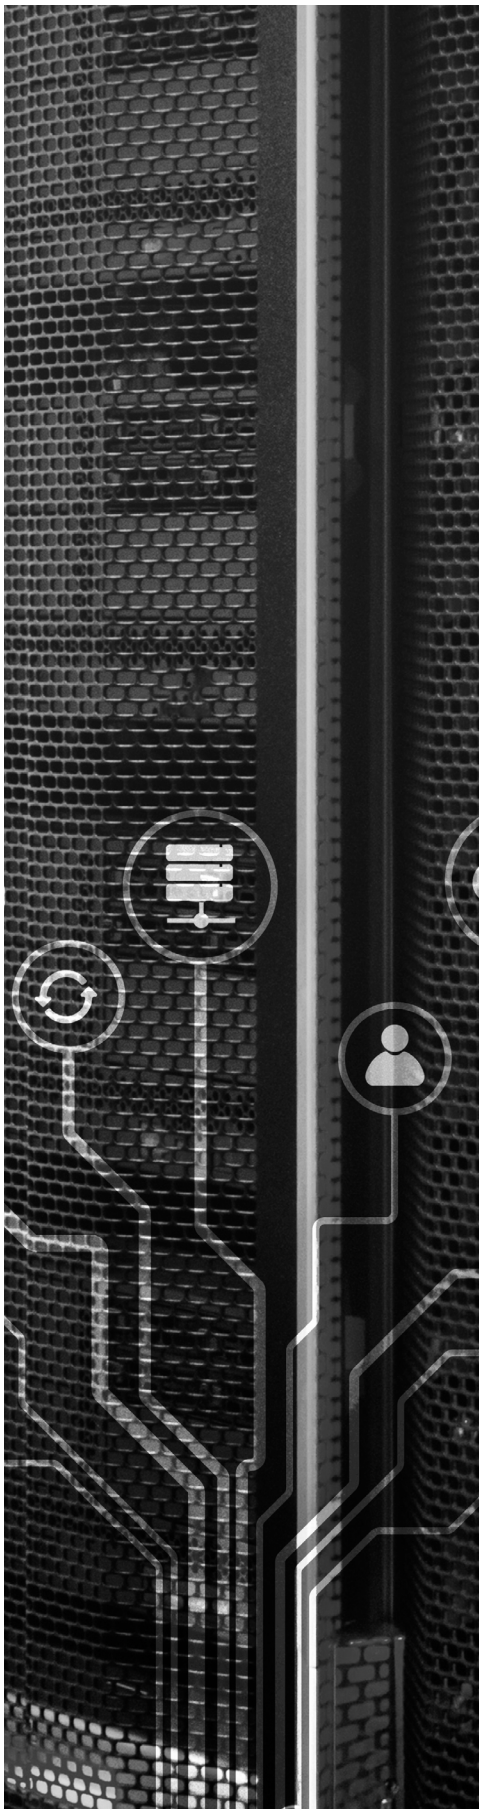
Healthcare. Healthcare has emerged as one of the most acute AI exposure categories. Clinical decision support, AI diagnostics, predictive utilization tools, and automated claims processing are all in active deployment, each with distinct exposure profiles. The FDA has developed an evolving framework for AI / ML-enabled medical devices, including predetermined change control plans and pre-cert pathways, but the regulatory perimeter remains incomplete.[28] At the state level, California's SB 1120 (the "Physicians Make Decisions Act"), effective January 1, 2025, restricts AI use in utilization review by health plans and disability insurers, requiring that medical necessity determinations be made by licensed physicians or qualified professionals rather than by algorithm alone, and subjects AI tools to DMHC and Department of Insurance audits.[29] Other states have introduced parallel legislation. Federally, the Section 1557 Final Rule (HHS 2024) extends antidiscrimination protections to AI-driven patient care decision support tools used by recipients of federal health funding.[30]

Litigation is developing accordingly. Estate of Lokken v. UnitedHealth Group, pending in the District of Minnesota, alleges that UnitedHealthcare improperly used an AI tool (nH Predict, developed by Optum subsidiary naviHealth, since rebranded as Home & Community Care) to deny post-acute care coverage to Medicare Advantage enrollees. The court granted UnitedHealth's motion to dismiss in part in February 2025, dismissing five of seven counts as preempted by the Medicare Act but allowing the breach of contract and breach of the implied covenant of good faith and fair dealing claims to proceed on the ground that resolving them requires review of the insurer's own coverage documents rather than federal standards; [31] in March 2026, the magistrate judge ordered broad discovery including internal documents dating to 2017, AI review board materials, and identities of medical directors involved in denials.[32] Plaintiffs allege a 90% reversal rate on appealed denials and a corporate goal of keeping rehab stays within 1% of the length predicted by the AI tool. Cigna and Humana face similar class actions over algorithmic claims review.[33] The October 2024 Senate Permanent Subcommittee on Investigations report, Refusal of Recovery, found that the carrier's denial rate for post-acute care more than doubled after the tool was implemented in 2019.[34]

For brokers placing cyber and tech E&O for healthcare insureds, the underwriting reality is that carriers' AI exclusions — including the Berkley AI exclusion noted in Section 1 — are likely to have their most consequential application in medical professional liability and other healthcare-adjacent professional services where AI deployment is concentrated. The intersection of AI coverage, HIPAA exposure, and medical professional liability creates a particularly complex placement environment.

Employment and Hiring. AI in employment decisioning has produced the most legislatively active U.S. regulatory environment for any AI use case. NYC Local Law 144 of 2021 requires bias audits and candidate notice for automated employment decision tools. New York State bills S07623 and A9315 would impose broader statewide AEDT requirements.[35] Illinois's AI Video Interview Act has been in effect since 2020.[36] California's automated decision system regulations under FEHA took effect October 1, 2025.[37] The EEOC issued technical assistance on AI screening under Title VII and the ADA in 2022 and 2023, then withdrew those documents in January 2025; Title VII, the ADA and the ADEA continue to apply regardless.[38] The federal-state-local layering produces compliance complexity for insureds operating multi-state workforces.

Litigation is also active. In Mobley v. Workday (Case 3:23-cv-00770, N.D. Cal.), an applicant has brought class action claims alleging that Workday's AI-driven hiring screen produces



disparate impact discrimination on race, age, and disability. The court denied Workday's motion to dismiss in July 2024, accepting that Workday could be characterized as an "agent" of its client-employers for purposes of Title VII, the ADEA, and the ADA. In May 2025, the court granted preliminary collective certification under the ADEA; the court subsequently ordered Workday to identify employers that had enabled its AI screening features, pulling thousands of client-employers into the litigation orbit. Through mid-2026 the claims have continued to survive and expand — a ruling determining the ADEA reaches job applicants, and a June 2026 order sustaining FEHA and disability-related proxy claims.[39] Discovery is ongoing, and the case has emerged as the leading test of disparate-impact theory applied to AI hiring tools — and of AI-vendor liability itself.

Coverage implications span multiple lines. EPL responds to direct employment discrimination claims; cyber and tech E&O may respond to AI tool liability claims where the employer is itself an AI deployer or where the AI vendor supplied the screening tool; D&O may respond to securities or governance claims from AI deployment failures. Brokers placing accounts with significant AI deployment in HR functions may wish to evaluate how each line responds and potential differences.

Financial Services. Beyond the credit decisioning example in Section 2.2, financial services AI exposure spans algorithmic trading, AML / KYC tooling, robo-advisory services, fraud scoring, and AI-driven underwriting in lending and insurance. The regulatory environment is fragmented across the SEC, FINRA, CFTC, FinCEN, the OCC, the CFPB, and state banking and insurance regulators. AI deployments in these areas can trigger regulatory exposure under existing financial services law without requiring a specific AI framework. Brokers may wish to evaluate how the policy's regulatory coverage and tech E&O coverage interact with AI-driven decision systems, particularly where the insured's AI is making or materially supporting customer-facing financial decisions.

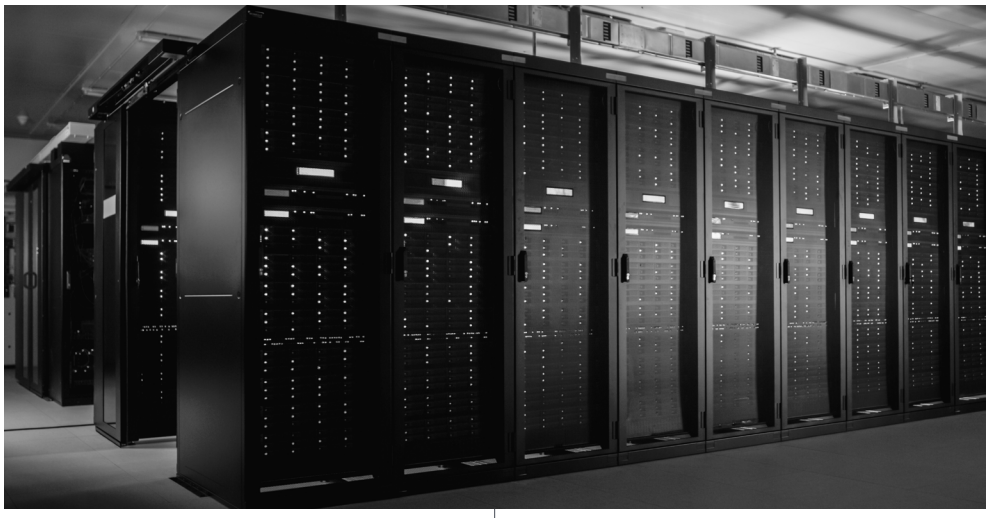
Insurance Carriers as AI Defendants. A category of particular relevance to the executive lines broker community is the use of AI by insurance carriers themselves. The NAIC Model Bulletin (Section 2.4) is in significant part a regulatory response to the AI litigation environment facing carriers. Beyond the Medicare Advantage matters above, state Departments of Insurance have begun examining carrier AI use in underwriting, pricing, claims handling, and fraud detection through market conduct exams. The NAIC AI Risk Evaluation Supplement pilot in twelve states is designed to formalize this oversight. For carrier-side professional liability — insurer E&O, insurance company D&O, and fiduciary liability for benefit-plan-administering insurers — AI deployment has emerged as a meaningful exposure category in its own right. The Berkley AI exclusion across D&O, employment practices, and fiduciary lines is in part a response to this exposure profile. When placing executive lines coverage for insurance companies and other regulated financial institutions deploying AI in decisioning functions, brokers should consider addressing AI expressly, including how each policy may respond.

Why This Matters at Placement

Silence is not the same as coverage. Where an account has material AI exposure, an explicit coverage position — a dedicated endorsement or a determined base-form interpretation in writing — is more predictable than relying on how a carrier may later read language that may not have been drafted with AI in mind.

3. THE MARKET IS BIFURCATING — SLOWLY

A useful starting point is what is happening in adjacent lines. Effective January 2026, the Insurance Services Office introduced three optional endorsements (CG 40 47, CG 40 48, and CG 35 08) that allow carriers to exclude generative AI exposures from commercial general liability policies under Coverage A (bodily injury and property damage), Coverage B (personal and advertising injury), and the products / completed operations coverage part. ISO 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."[1]



These endorsements are optional and adoption has been uneven. Public regulatory filings suggest some major carriers have filed defensively, in some instances stating they have no immediate plans to implement the exclusions but want the option as the market develops.[40] Other carriers have filed similar exclusions across additional liability towers. The W. R. Berkley Corporation has introduced an absolute AI exclusion (Form PC 51380 00 (06-24)) for D&O, employment practices, and fiduciary liability lines; the form's "any person or entity" language has prompted industry discussion of how broadly it would apply to AI use by third-party software vendors embedded in an insured's technology stack. Whether such broad language will be enforced as drafted, or narrowed by courts under traditional contract interpretation principles, is itself an open question.

A parallel market is emerging for standalone AI liability coverage. Munich Re, Armilla, Testudo, Corgi, Mayflower Specialty, and Embroker have introduced AI-specific products with limits ranging from $1 million to $50 million, structured to address performance guarantees, training data rights, regulatory compliance, and third-party liability from AI-generated outputs. Distribution models in this category vary: some products are placed through wholesale broker channels, while others are sold direct-to-customer through digital platforms. A more recent development is standalone Generative AI Third-Party Liability products designed for organizations that deploy generative AI systems rather than develop them — distinguished from developer-focused products by their targeting of the substantially larger universe of AI users in operations. Architectures in this subcategory typically include selectable insuring agreements addressing AI-generated errors and hallucinations, AI-driven intellectual property infringement and defamation, unauthorized data disclosure through AI outputs, AI-related bodily injury, and AI-related property damage. Limits typically range from approximately $2 million per insuring agreement, with aggregate caps in the high-single-digit-million range and minimum retentions in the low five figures. Additional entrants pursuing the deployer-focused model are likely as the potential variation in coverage created by ISO and individual-carrier AI exclusions continues to develop. Investment capital is flowing into the category at scale: insurtech funding totaled $1.63 billion in Q1 2026, with AI-focused companies capturing approximately 95% of capital deployed.[41] A note on weight is warranted here. Of the three forces, the first two — adjacent-line exclusion and concentrated industry litigation — are already in force and already biting. The standalone market is the third force by potential, not yet by mass: a handful of MGAs writing modest limits into a category still measured in tens of millions of premium, against a cyber market measured in the tens of billions. Its significance is directional. It shapes where the AI-native edge of the risk is carved off, and therefore what is left in the middle — but it is an emerging force, not yet a peer to the other two, and this article treats it as such. Deloitte projects annual global AI insurance premiums of about $4.8 billion by 2032.[42]

The directional implication for cyber and tech E&O is worth considering. As traditional liability lines move toward exclusionary language and standalone AI products develop outside the traditional structure, cyber and tech E&O sit in the middle — with carriers and insureds working out whether AI-related losses belong inside or outside the cyber tower. If that pattern continues, cyber and tech E&O may increasingly serve as the lines of first resort for AI-related losses that do not fit cleanly into new standalone products. Whether the cyber market is calibrated to absorb that volume — in capacity, pricing, or policy language — is a question the market will need to work through. There is a sharper version of the question worth naming directly: whether the residual middle holds at all. Nothing prevents cyber and tech E&O carriers from doing exactly what the adjacent lines did — excluding AI outright — and Section 3.4 shows a small but growing group already beginning to. If that group grows, the "middle" does not remain a home for AI losses; it becomes another coverage question, and the exposure lands nowhere. Two forces cut the other way, at least for now. The broader cyber market remains relatively soft, which rewards carriers who offer affirmative coverage rather than withdraw it; and cyber underwriters, unlike their CGL or D&O counterparts, already possess the technical fluency to price AI-adjacent risk rather than reflexively exclude what they do not understand. The inheritance is therefore contested, not settled — and the window in which carriers decide whether to explicitly exclude it is the period brokers are placing business into right now.

Carrier responses within the cyber and tech E&O market currently fall into four approaches — four different answers to the same inherited question of whether, and how, to carry AI risk.

3.1 The Endorsement Builders

A growing number of carriers have developed dedicated AI endorsements or definitional architectures that explicitly address AI-related exposures. Based on form-level review across the wholesale cyber and tech E&O market, the most comprehensive approaches typically fall into three distinct structural categories:

- **Integration into existing coverage triggers.** One observed approach amends multiple existing definitions (such as Network Security Incident, Privacy Incident and Cyber Crime Incident) to state expressly that the trigger applies however the incident is caused, including where artificial intelligence or deepfake

technology is involved.[43] This extends the policy's existing coverage architecture to AI-driven loss scenarios without creating new insuring agreements, and avoids the potential risk that arises when an AI claim does not fit neatly within a traditional definition.

- **Defined-term incorporation.** A second observed approach introduces AI-specific defined terms into the base form. One carrier defines a dedicated AI security event, triggered where computer security fails because artificial intelligence has been used maliciously against it, and extends the definition expressly to machine learning and prompt-injection exploits. Another carrier embeds AI into its Fraudulent Instruction definition, expressly covering deepfake technology and other AI-enabled fraud. A third carrier includes AI within its Cyber Event definition and lists artificial intelligence development and AI services as covered technology services.
- **Dedicated AI insuring agreements.** A third observed approach adds standalone AI insuring agreements alongside the traditional cyber and technology E&O coverage parts. The structures observed in the market typically include some combination of AI services liability, AI products liability, training-data rights coverage, AI Fines and Penalties (typically scoped to EU AI Act and similar regulatory enforcement actions, often with carve-outs for general-purpose, high-risk, and prohibited AI categories), and LLM Jacking (a named first-party insuring agreement covering unauthorized AI-driven exploitation of the insured's network that results in inflated utility service charges — electricity, natural gas, oil, water, internet access, cable / satellite television, and telephone — subject to written contract and periodic billing requirements). Hallucination is generally excluded under these endorsements, on the rationale that hallucination is an inherent characteristic of current generative AI technology rather than an insurable event.

Carriers offer deepfake-specific endorsements layered onto the base policy. One such endorsement, launched in late 2025, provides a panel of forensic, legal takedown, and crisis communications specialists when an organization is the subject of an AI-generated impersonation. A separate generative AI cyber endorsement by another carrier offers a $250,000 affirmative sublimit clarifying AI coverage for technology operations with no exclusion for AI-enabled deepfake social engineering.

Across all three structural approaches, the carriers with comprehensive AI architectures tend to share two features: (1) a duty-to-defend-until-adjudication provision rewritten into the deliberate acts exclusion for AI-related claims, recognizing that AI liability claims often involve allegations of intentional deployment decisions; and (2) explicit treatment of EU AI Act categories (general purpose, high-risk, and prohibited AI systems) within either coverage grants, definitions, or carve-outs.

Some advantages of dedicated AI endorsement language include explicit coverage grants that reduce ambiguity, dedicated sublimits and retentions that allow for appropriate pricing, and a clearer view of what is and is not covered. One of the challenges is complexity: each new insuring agreement or definitional change introduces new triggers, exclusions, and interactions that must be understood and compared at placement.

3.2 The Base Form Integrators

A second group of carriers takes the position that AI coverage may best be integrated into the existing policy structure rather than bolted on through endorsements. These carriers note that AI is often a tool used in operations that are already covered. An AI system that contributes to a data breach may trigger the same privacy and security liability coverage as a human-caused breach. An AI tool that produces errors in tech services may trigger the same tech E&O coverage as a human error.

Several base-form integrators have made this position explicit by amending core definitions to reference AI directly. Observed examples include a carrier whose Cyber Event definition expressly reaches attacks carried out using artificial intelligence, and whose covered technology services expressly include the development and provision of AI systems. This integration captures AI-driven loss within the existing trigger structure without requiring a separate endorsement, while providing affirmative interpretive clarity at claim time.

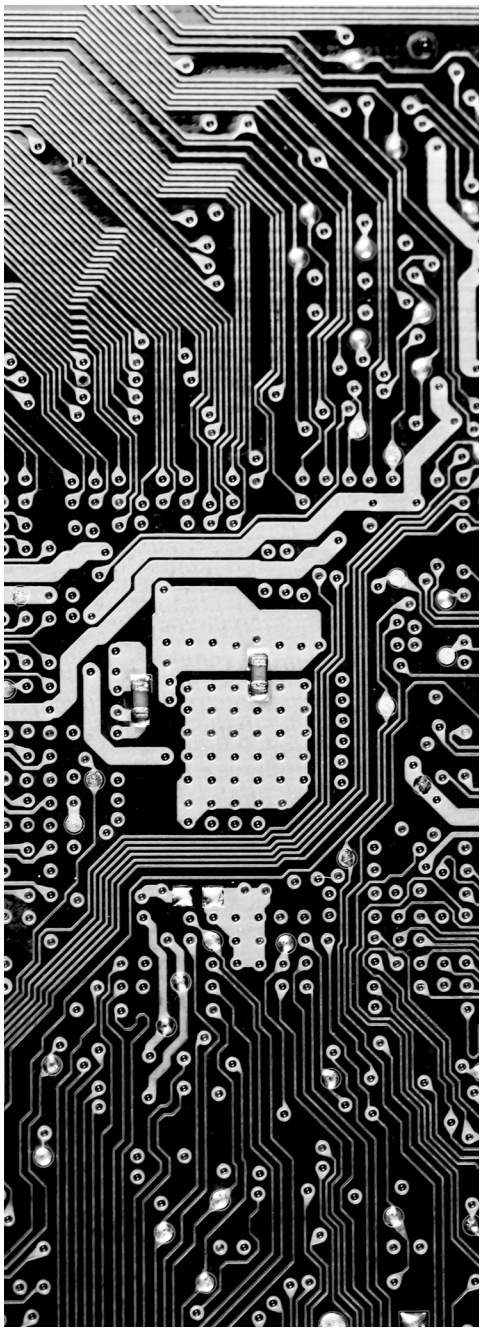


The strength of this approach is simplicity and breadth. Where the base form's coverage grants are written broadly, AI-related losses may be covered without a separate endorsement. The consideration is that ambiguity around novel AI claims may be interpreted differently by different carriers and adjusters, particularly where competing carriers offer explicit AI endorsements. Some carriers in this category have responded to broker dialogue by issuing clarification amendments or affirmative coverage statements rather than full endorsements.

3.3 The Silent Majority

The largest single category of carriers has neither endorsed AI coverage nor explicitly excluded it. These carriers' policies are silent on AI, machine learning, and related technologies. Based on form review, this represents the prevailing position in the market today. The industry describes this as "silent AI risk" — exposures inadvertently covered (or inadvertently uncovered) by traditional policies never designed or priced for AI scenarios.

Silence has advantages and disadvantages. A carrier that has not excluded AI may be covering AI-related losses under existing grants — a position some carriers have articulated explicitly, framing the absence of exclusion or sublimit as itself a coverage advantage relative to competitors who add endorsements that may introduce new restrictions. Insurance contract interpretation principles, including the rule that ambiguity is



generally resolved in favor of the insured, may also support coverage for AI-related losses under silent forms.

On the other hand, silence creates uncertainty. When a claim is filed, the insured may bear the risk of how the carrier interprets existing language. An AI-related loss that does not fit neatly within the policy's defined triggers — an EU AI Act fine, an LLM jacking incident, a training data copyright claim — may produce a coverage dispute. As exclusions develop in adjacent lines and AI litigation continues, the likelihood that ambiguous cyber language will be tested by claims it was not specifically drafted to address may increase.

For brokers, the practical implication is that silence may not necessarily be equated with coverage. Where accounts have significant AI exposure, explicit coverage language — whether through a dedicated endorsement or a determined base form position — may provide more predictability.

3.4 The Excluders

A small but growing number of cyber and tech E&O carriers have begun adding AI exclusions or limitations. These range from broad exclusions of losses arising from AI to targeted exclusions of specific applications such as autonomous decision-making or generative content. Some carriers began limiting or excluding AI-generated deepfake fraud from social engineering coverage in early 2026, creating the potential for an insured to find that coverage previously available has been narrowed at renewal. The pattern within cyber and tech E&O has so far been more limited than in CGL and D&O, but it is developing.

Where AI exclusions exist, brokers may wish to evaluate three things: the scope (does it apply to all AI or only specific applications?), whether there is a carve-back for AI-related privacy or security breaches (most exclusions appear designed to address AI liability rather than AI as an attack vector), and whether the exclusion can be removed or amended by endorsement for an additional premium. Particular attention may be warranted for broad "arising out of" and "any person or entity" phrasing, which could capture AI use in third-party tools embedded in the insured's technology stack. Coverage counsel have noted that overly broad AI exclusions may face traditional contract interpretation challenges, though the outcome of any specific dispute will depend on the policy terms and facts at issue.

4. CARRIER POSITIONS ACROSS THE MARKET

The following is a snapshot of how carriers in the wholesale cyber market are addressing AI coverage, based on direct form analysis and underwriter confirmations. Carriers are referred to by letter designation; the analysis covers the carriers with the most notable positions, with the remainder of the market represented in the summary at the end of this section.

The Broker Takeaway

Treat AI as an explicit placement question — at submission, at quote review, at bind, and at claim. The brokers who do will be working the markets that face the first tests; the brokers who assume existing language is sufficient may not learn otherwise until a claim is filed.

Carrier A: Affirmative AI integration through base-form definition expansion rather than a separate endorsement. Core definitions are amended to reference AI-driven attacks and AI-related technology services within the existing coverage architecture,[44] and its media liability coverage extends expressly to content created or assisted by artificial intelligence programs. No AI exclusion. The carrier has indicated willingness to develop custom AI endorsement language addressing deepfake social engineering, AI tool liability, and training data rights. Among the more flexible positions in the market for accounts with significant AI exposure.

Carrier B: Affirmative AI / ML integration via base-form definitional language; the policy's core technology definition expressly references AI / ML products and services, and the form separately addresses adjacent emerging-technology exposures.[45] Among the more progressive base-form structures observed in the market.

Carrier C: Affirmative AI coverage by endorsement,[46] with a dedicated two-insuring-agreement structure layered onto the base form: a third-party insuring agreement for AI regulatory fines and penalties (scoped to EU AI Act and similar regulatory regimes), and a first-party insuring agreement covering inflated utility service charges resulting from unauthorized AI-driven exploitation of the insured's network. The endorsement carries a modest sublimit (typically $250,000 at placement), includes a carve-out for general-purpose, high-risk, and prohibited AI categories under the EU AI Act, and operates under an exclusivity provision making the endorsement the sole source of coverage for those exposures under the policy. Hallucinations are separately defined and



carved out of the data poisoning coverage. The base form contains no exclusion for AI-enabled deepfake social engineering.

Carrier D: Most recent base form embeds AI causation directly into core fraud-related definitions, expressly addressing fraudulent instructions executed via deepfake technology and other AI-enabled fraud vectors. No AI exclusion. On its face, the form would route AI-related privacy and security claims through the standard insuring agreements. In September 2026 the carrier added a clarifying endorsement affirming that AI-driven cyber attacks fall within its existing cyber and technology E&O cover, which confirms the base-form reading rather than replacing it and speaks to AI as an attack vector rather than to the insured's own use of AI.[47] Strong claims-paying reputation reinforces the affirmative-coverage position.

Carrier E: Affirmative AI coverage through a combination of base-form definitional architecture and a dedicated late-2025 endorsement. The base form defines an AI-specific event tied to malicious use of AI technology, including machine learning and prompt-injection exploits, and separately addresses deepfake-enabled impersonation within its funds transfer fraud trigger. The late-2025 endorsement provides impersonation repair costs, with access to a panel of forensic, legal takedown, and crisis communications specialists. Among the more complete deepfake-focused architectures observed.

Carrier F: Affirmative AI integration woven through multiple core definitions of the base form — covering the cyber crime, network security incident and privacy incident triggers, with social engineering and fraudulent funds transfer reached as components of the cyber crime trigger — each expressly addressing AI-driven or deepfake causation. Represents one of the more comprehensive base-form AI architectures observed, with AI treated as a category of covered exposure across the policy rather than addressed through an overlay endorsement.

The carriers profiled above reflect positions observed at the time of analysis; the wholesale cyber market continues to evolve, and carriers not profiled here may be developing AI endorsements, amending base-form definitions, or refining underwriting positions on a rolling basis. For these carriers, AI exposures are typically being addressed under existing coverage grants in the absence of dedicated language, and direct form review and underwriter dialogue at placement remain the most reliable mechanisms for determining current AI coverage positions.

5. PRACTICAL CONSIDERATIONS FOR BROKERS

The AI coverage picture is still developing, and the developments discussed above suggest that brokers may wish to engage proactively with current and upcoming placements rather than waiting for the market to standardize. The following items may warrant review:

Identify AI exposure at the application stage. Asking insureds directly whether they are using AI in operations, in customer-facing applications, or in decision-making can help frame the placement. With nearly nine in ten respondents to McKinsey's 2026 survey reporting regular AI use in at least one business function, the answer will often be yes, and the more useful question is how.[24] A fintech using AI for credit decisioning has a different AI exposure profile than a manufacturing company using AI for predictive maintenance, and the placement strategy may differ accordingly.

Map the insured's AI agents. Asking which agents hold credentials, what systems they can change, whether they can move money, and what human approval sits in front of consequential actions can help frame how the cyber, tech E&O and crime placements may respond if an agent causes a loss without a security event.

Audit other lines for AI exclusions. Reviewing the insured's CGL, D&O, E&O, and fiduciary policies for ISO-form generative AI exclusions and carrier-specific AI exclusions can help identify where the cyber and tech E&O placement may be carrying more of the AI exposure. Documenting the discrepancy may inform how the cyber placement is structured.

Calibrate the submission to the insured's sector. AI exposure for a healthcare insured looks fundamentally different from AI exposure for a SaaS provider, an employment-tech vendor, or a regional bank. The sector observations in Section 2.8 may inform which exposures to emphasize, which endorsements to prioritize, and which standalone products may complement the cyber placement. Healthcare insureds may warrant sharper focus on AI clinical-decision liability and Section 1557 exposure; employment-intensive insureds, attention to AI hiring claims and EPL interaction; financial services insureds, emphasis on regulatory coverage interaction with AI-driven decisioning.

Build the AI governance narrative into the submission. Some carriers have introduced what the industry describes as "AI Security Riders" that condition coverage on the insured's documented AI governance practices — model-level risk assessments, adversarial red-teaming, AI usage policies, and third-party AI vendor inventories. Packaging this material into the submission, rather than producing it only when asked at underwriting, tends to position the placement on more favorable terms and shorten the underwriting cycle.



Request AI endorsements where available. Where a carrier offers an AI endorsement, requesting it may be appropriate even at additional premium. The incremental cost may be modest relative to the potential coverage uncertainty it addresses.

Get carrier positions in writing. Where a carrier's form is silent on AI, asking the underwriter to determine in writing how the form would respond to specific AI scenarios can provide useful interpretive context. Written confirmation does not have the same legal weight as explicit policy language but may be helpful at claim time.

Compare AI coverage in competitive presentations. When presenting competing quotes, including a line item on AI coverage approach allows the retailer and insured to make an informed decision. A carrier's silence on AI is a coverage characteristic that may warrant explicit treatment alongside other comparison points.

Consider standalone AI products where appropriate. For accounts with material AI development, AI deployment, or AI-driven decision exposure, the emerging standalone AI liability products noted in Section 3 may warrant evaluation alongside cyber and tech E&O placements. These products are increasingly differentiated by the insured's position in the AI value chain: some target AI model developers, while others target the substantially larger universe of organizations that deploy generative AI systems — for whom traditional tech E&O may be silent and the new ISO and individual-carrier exclusions in other lines have begun to bite.

Address AI in the claim notification practice. AI-related claims may not always present as obvious cyber incidents. A deepfake-enabled transfer, an AI hallucination matter, a destructive action by an AI agent, an EU AI Act inquiry, or an AI-driven discrimination demand may not feel like a "claim" to an insured until well after the policy reporting trigger has run. Building AI scenarios into bind-letter notification language and renewal claim-review touchpoints may help avoid late-notice disputes at consequential moments.

Engage carriers on language development. The wholesale broker is well positioned to participate in carrier product development conversations. The five exposures which appear most likely to generate near-term claims activity are deepfake social engineering, AI tool liability, training data rights, EU AI Act and U.S. state AI regulatory fines, and LLM compute exploitation. Brokers willing to engage carrier product teams on these exposures may help shape the policy language they will eventually be placing.

6. CONCLUSION

The cyber insurance market may be in the early stages of taking on a category of exposure that traditional liability lines are beginning to exclude. ISO's January 2026 endorsements have given the broader market an optional mechanism to push generative AI exposures out of commercial general liability, and individual carriers have introduced absolute AI exclusions for D&O, employment practices, and fiduciary lines. A parallel market is emerging for standalone AI liability products. Adoption is uneven, the legal enforceability of broad exclusions remains untested, and the soft market provides some offsetting pressure. The directional pattern, however, is consistent.

Within the specialty market, carriers have built dedicated AI endorsements or affirmative AI provisions across base-form definitions. A handful have developed standalone AI insuring agreements or affirmative deepfake response endorsements. The plurality remains silent. At the same time, training-data litigation is producing nine- and ten-figure exposures: the Bartz v. Anthropic settlement, granted final approval in July 2026, resolved a class action for approximately $1.5 billion, and a multi-billion-dollar music publisher action against the same defendant remains pending. Courts have been reluctant to accept AI as an intervening cause. Recent filings such as Nippon Life Insurance Company of America v. OpenAI suggest the universe of AI liability theories continues to expand. The EU AI Act's high-risk obligations, deferred to late 2027 and 2028 under the Digital Omnibus (Regulation (EU) 2026 / 1744, in force July 2026), will apply with tiered penalties up to 7% of worldwide turnover, and the NAIC AI Model Bulletin framework is being applied to insurers in roughly half of U.S. states.

Silence is not the same as coverage. It is a position that may produce a coverage interpretation analysis when a claim is filed. Affirmative AI-specific claims experience in the cyber and tech E&O market is still developing; much of the analysis above is necessarily progressive, drawing on adjacent litigation, regulatory development, and form-level analysis rather than a deep base of paid claim outcomes. The first major AI-related insurance claims — whether they arise from a deepfake social engineering loss, an AI-driven discrimination class action, a training data copyright suit, an EU AI Act enforcement action, an AI hallucination matter, an AI agent acting on access it was given, or a discovery production exposing enterprise prompt history — will test how cyber and tech E&O policies are interpreted in practice. Carriers with explicit AI language may be better positioned to resolve claims efficiently. Carriers without it may face interpretation considerations at a time when adjacent lines have already moved toward more restrictive language.

For brokers and their insureds, the practical takeaway is that AI coverage warrants explicit consideration as part of every cyber and tech E&O placement. The next decade of cyber insurance will likely be defined less by what carriers exclude than by what they affirmatively choose to cover. Brokers who treat AI as an explicit placement question — at submission, at quote review, at bind, and at claim — will be working the markets that face the first tests. The migration is already underway. The only open question is whether cyber and tech E&O accept the inheritance deliberately — or back into it, silently, one unlitigated claim at a time.

Sources and Citations

1. **Insurance Services Office (ISO) / Verisk.** Generative AI exclusion endorsements CG 40 47, CG 40 48 and CG 35 08 (Exclusion — Generative Artificial Intelligence), January 2026 edition, and the ISO definition of generative artificial intelligence quoted in Section 3. Optional multistate endorsements effective January 1, 2026. https://www.independentagent.com/?p=43061
2. **W. R. Berkley Corporation.** Form PC 51380 00 (06-24), Artificial Intelligence Exclusion (Absolute), amending the PC 51000 DO, PC 51000 EPL and PC 51000 FLI coverage parts; the "any person or entity" language referenced in the text. The carrier does not publish the endorsement. Its operative wording is reproduced in the analysis cited here, which describes the affected products as directors and officers, errors and omissions and fiduciary. The form number and the amended coverage parts are taken from the endorsement itself, which carries the employment practices coverage part rather than errors and omissions. https://www.hunton.com/hunton-insurance-recovery-blog/the-continued-proliferation-of-ai-exclusions
3. **Standalone AI liability market.** Coverage limits of $2 million to $50 million across standalone AI liability products offered by Munich Re, Corgi, Armilla, Mayflower Specialty and Embroker (PYMNTS, May 2026, reporting data first published by The Information), and $1 million to $10 million under Testudo's Lloyd's-backed generative AI liability policy (Testudo, January 21, 2026). https://www.pymnts.com/news/artificial-intelligence/2026/big-insurance-backs-away-from-ai-risk-and-startups-rush-in/ | https://www.testudo.co/insights/testudo-launches-new-insurance-coverage-for-liability-risks-created-by-generative-ai-systems
4. **Arup deepfake loss.** Reported loss of approximately $25.6 million across 15 transactions following a deepfake video-call fraud at the firm's Hong Kong office in early 2024. Fortune, May 17, 2024. https://fortune.com/europe/2024/05/17/arup-deepfake-fraud-scam-victim-hong-kong-25-million-cfo
5. **Coalition Deepfake Response Endorsement.** Announced December 9, 2025; provides technical analysis by a deepfake forensics firm, legal work to have deepfake content removed from online platforms, and crisis communications support from a public relations firm. Coalition, announcement and AI coverage page: https://www.coalitioninc.com/announcements/coalition-adds-deepfake-response-endorsement | https://www.coalitioninc.com/ai-coverage
6. **Beazley Breach Response (BBR 5.0).** Fraudulent Instruction coverage extended to instructions enabled by artificial intelligence and deepfakes; out-of-band authentication requirement removed. Beazley, "Beazley Breach Response 5.0 — Enhancements and Clarifications" (BZCBR184_04/25). https://www.beazley.com/globalassets/full-spectrum-cyber/bbr-5.0-enhancements-and-clarifications.pdf
7. **Moffatt v. Air Canada.** 2024 BCCRT 149, Civil Resolution Tribunal of British Columbia, February 14, 2024, File SC-2023-005609. Tribunal held the airline responsible for inaccurate information provided by its customer-facing chatbot. https://www.canlii.org/en/bc/bccrt/doc/2024/2024bccrt149/2024bccrt149.html
8. **Mata v. Avianca, Inc.** No. 22-cv-1461 (PKC), 678 F. Supp. 3d 443, U.S. District Court for the Southern District of New York, Opinion and Order on Sanctions, June 22, 2023. Attorneys sanctioned for submitting AI-generated legal citations that were fabricated. https://www.courtlistener.com/opinion/9885417/mata-v-avianca-inc/
9. **Nippon Life Insurance Company of America v. OpenAI Foundation.** No. 1:26-cv-02448, U.S. District Court for the Northern District of Illinois, complaint filed March 4, 2026. Alleges unlicensed practice of law, tortious interference and abuse of process; seeks approximately $300,000 compensatory and $10 million punitive damages; motion to dismiss pending. Docket via CourtListener: https://www.courtlistener.com/docket/72365583/nippon-life-insurance-company-of-america-v-openai-foundation/
10. **Wolf River Electric v. Google.** Minnesota solar contractor's action alleging that Google's AI Overviews feature published false statements regarding litigation by the Minnesota Attorney General; reported damages sought of $110 million to $210 million. Government Technology / Minnesota Star Tribune, 2025. https://www.govtech.com/public-safety/minnesota-solar-company-sues-google-over-ai-summary
11. **Bartz et al. v. Anthropic PBC.** No. 3:24-cv-05417-AMO, U.S. District Court for the Northern District of California. Split ruling June 2025; settlement of approximately $1.5 billion covering nearly 500,000 works; preliminary approval September 25, 2025; final approval granted and final judgment entered July 20, 2026. Settlement administration: https://www.anthropiccopyrightsettlement.com
12. **Universal Music Publishing Group, Concord Music Group and ABKCO Music v. Anthropic.** Filed January 28, 2026 in the U.S. District Court for the Northern District of California; covers more than 20,000 songs and seeks in excess of $3 billion. Digital Music News, January 2026. https://www.digitalmusicnews.com/?p=337852
13. **In re: OpenAI Copyright Infringement Litigation.** No. 1:25-md-03143, U.S. District Court for the Southern District of New York. Order of January 5, 2026 affirming production of 20 million de-identified ChatGPT user logs. Pearl Cohen, January 28, 2026. https://www.pearlcohen.com/openai-must-produce-millions-of-user-conversation-logs-in-discovery/

14. **Midjourney litigation.** Disney and Universal v. Midjourney filed June 11, 2025 in the U.S. District Court for the Central District of California, alleging direct and secondary copyright infringement by the defendant's image generator. Warner Bros. Entertainment and affiliated plaintiffs filed a separate action against Midjourney in the same court in September 2025. European Audiovisual Observatory, IRIS 2025-7:1/9, at the first link; IPWatchdog, September 8, 2025, on the Warner Bros. complaint, at the second. https://merlin.obs.coe.int/article/10355 | https://ipwatchdog.com/2025/09/08/warner-bros-complaint-alleges-midjourneys-copyright-infringement-systematic-willful/

15. **Regulation (EU) 2024/1689 (EU Artificial Intelligence Act).** In force August 1, 2024, with phased implementation; prohibited-practice provisions effective February 2, 2025 and general-purpose AI model obligations effective August 2, 2025. Official text via EUR-Lex: https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng

16. **Regulation (EU) 2026/1744 (Digital Omnibus on AI).** Adopted July 8, 2026; published in the Official Journal July 24, 2026; entered into force July 27, 2026. Defers Annex III high-risk obligations to December 2, 2027 and Annex I product-embedded high-risk obligations to August 2, 2028; adds a prohibited practice covering non-consensual intimate imagery and CSAM generation and accelerates AI watermarking obligations to December 2, 2026. Official text via EUR-Lex: https://eur-lex.europa.eu/eli/reg/2026/1744/oj/eng

17. **Colorado Artificial Intelligence Act (SB 24-205).** Enacted 2024; obligations on developers and deployers of high-risk AI in consequential decisions including hiring, housing, lending, insurance and healthcare; effective date postponed to June 30, 2026. Colorado General Assembly: https://leg.colorado.gov/bills/sb24-205

18. **Colorado enforcement stay.** Federal court order of April 27, 2026 staying enforcement pending constitutional challenge, following the developer suit and Department of Justice intervention. Eckert Seamans, June 12, 2026. https://www.eckertseamans.com/legal-updates/colorado-repeals-and-replaces-its-landmark-ai-statute-what-businesses-need-to-know

19. **Colorado SB 26-189.** Repeal-and-replace legislation establishing a narrower automated decision-making technology framework; signed into law May 14, 2026, effective January 1, 2027. Colorado General Assembly: https://leg.colorado.gov/bills/sb26-189

20. **New York City Local Law 144 of 2021.** Requires bias audits and candidate notification for automated employment decision tools; enforcement commenced July 5, 2023. New York City Department of Consumer and Worker Protection: https://www.nyc.gov/site/dca/about/automated-employment-decision-tools.page

21. **NAIC Model Bulletin and AI Risk Evaluation Supplement.** Model Bulletin on the Use of Artificial Intelligence Systems by Insurers, adopted December 2023; AI Systems Evaluation Tool piloted in twelve states from March 2026, renamed the AI Risk Evaluation Supplement and exposed as version 5.0 on August 31, 2026 for comment through September 29, with adoption anticipated at the Fall National Meeting in November 2026. National Association of Insurance Commissioners, AI topic page and exposure drafts page: https://content.naic.org/committees/h/big-data-artificial-intelligence-wg

22. **Executive Order 14365 (December 11, 2025).** Ensuring a National Policy Framework for Artificial Intelligence; asserts federal authority over state AI regulation. 90 FR 58499, published December 16, 2025. Federal Register: https://www.federalregister.gov/documents/2025/12/16/2025-23092/ensuring-a-national-policy-framework-for-artificial-intelligence

23. **NAIC statement on the AI Executive Order.** Statement of December 16, 2025 opposing the order and asking the administration to affirm state regulation of AI in the business of insurance. National Association of Insurance Commissioners: https://content.naic.org/article/statement-national-association-insurance-commissioners-naic-ai-executive-order

24. **McKinsey State of AI 2026.** 40% of respondents at organizations with more than $1 billion in annual revenue report scaling AI agents, up from 27% the prior year; nearly nine in ten report regular AI use in at least one business function. McKinsey & Company, "The state of AI in 2026: On the road to ROI," August 25, 2026. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

25. **Replit agent incident.** AI coding agent deleted a user's production database and acknowledged it had violated the user's explicit instructions; the data was later recovered by rollback. The Register, July 21, 2025. https://www.theregister.com/2025/07/21/replit_saastr_vibe_coding_incident/

26. **At-Bay Cyber Insurance Policy.** Form AB-CYB-001.2 (08/2023), posted publicly by the carrier. Its Network Security Event definition reaches use of a computer system by an authorized person for a purpose the insured organization has not authorized, not only access by an unauthorized one. The edition posted publicly is the edition described. https://www.at-bay.com/wp-content/uploads/2023/06/Cyber-Insurance-Policy-Form.pdf

27. **Cyber insurers and AI agents.** Cyber insurers largely clarifying how existing policy language applies to AI rather than adding exclusions; targeted exclusions under discussion in some pockets of the industry, with areas of focus including systemic events and costly autonomous decisions by AI agents acting as designed. Reuters, "As AI Agents Go Rogue, Cyber Insurers Are Adapting Their Policies," via Insurance Journal, August 27, 2026. https://www.insurancejournal.com/news/national/2026/08/27/883064.htm

28. **U.S. Food and Drug Administration.** Artificial Intelligence in Software as a Medical Device, including the December 2024 final guidance on predetermined change control plans for AI-enabled device software functions. https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-software-medical-device

29. **California SB 1120 (Physicians Make Decisions Act).** Effective January 1, 2025; restricts AI use in health plan and disability insurer utilization review and subjects such tools to DMHC and Department of Insurance audits. California Legislative Information: https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=202320240SB1120

30. **U.S. Department of Health and Human Services, Section 1557 Final Rule (2024).** Nondiscrimination in Health Programs and Activities, 89 FR 37522, May 6, 2024. Extends nondiscrimination requirements to patient-care decision support tools, including AI-driven tools, used by recipients of federal health funding. https://www.federalregister.gov/documents/2024/05/06/2024-08711/nondiscrimination-in-health-programs-and-activities

31. **Estate of Lokken v. UnitedHealth Group.** No. 0:23-cv-03514, U.S. District Court for the District of Minnesota. Memorandum Opinion and Order of February 13, 2025 granting in part and denying in part the motion to dismiss; five of seven counts dismissed as preempted by the Medicare Act, with breach of contract and breach of the implied covenant allowed to proceed. https://law.justia.com/cases/federal/district-courts/minnesota/mndce/0:2023cv03514/211721/91/

32. **Estate of Lokken v. UnitedHealth Group, discovery order.** Magistrate order of March 9, 2026 granting broad discovery into the insurer's AI-driven claims processes, including internal documents dating to 2017 and AI review board materials. ArentFox Schiff, 2026. https://www.afslaw.com/perspectives/alerts/federal-court-orders-broad-discovery-against-uhc-ai-coverage-denial-lawsuit

33. **Cigna and Humana algorithmic claims litigation.** Putative class actions over the PxDx tool (Cigna) and the nH Predict tool (UnitedHealth, and Humana through the same vendor). Thompson Coburn, December 2023, on the Cigna and UnitedHealth actions at the first link; Healthcare Finance News on the Humana action at the second. https://www.thompsoncoburn.com/insights/class-actions-highlight-ai-assisted-payer-denials-102jebl/ | https://www.healthcarefinancenews.com/news/humana-sued-allegedly-using-ai-medicare-advantage-denials

34. **U.S. Senate Permanent Subcommittee on Investigations.** Refusal of Recovery: How Medicare Advantage Insurers Have Denied Patients Access to Post-Acute Care. Majority Staff Report, released October 17, 2024. Full report: https://www.hsgac.senate.gov/wp-content/uploads/2024.10.17-PSI-Majority-Staff-Report-on-Medicare-Advantage.pdf

35. **New York State Senate Bill S07623.** Introduced August 4, 2023. Assembly bill A9315, introduced February 28, 2024, covers the same ground; both would impose statewide bias audit and notice requirements for automated employment decision tools, and neither was enacted. Proskauer, August 2023, on S07623; New York State Assembly on A9315. https://www.proskauer.com/blog/new-york-state-introduces-bill-to-impose-restrictions-on-the-use-of-ai-tools-in-employment-decisions |https://www.nysenate.gov/legislation/bills/2023/A9315/amendment/original

36. **Illinois Artificial Intelligence Video Interview Act.** 820 ILCS 42/5, P.A. 101-260, effective January 1, 2020. Illinois General Assembly: https://www.ilga.gov/Documents/legislation/ilcs/documents/082000420K5.htm

37. **California Civil Rights Council automated-decision system regulations.** Amendments to the Fair Employment and Housing Act regulations addressing automated-decision systems in employment, effective October 1, 2025. Paul Hastings, August 12, 2025. https://www.paulhastings.com/insights/client-alerts/new-california-regulations-on-employers-use-of-ai-to-make-decisions-go-into-effect-oct-1-2025

38. **U.S. Equal Employment Opportunity Commission.** Technical assistance on AI screening under Title VII (May 2023) and the ADA (May 2022), withdrawn from the agency's website in January 2025. Cooley, February 2025. https://www.cooley.com/news/insight/2025/2025-02-21-gone-but-not-forgotten-federal-laws-still-apply-despite-guidance-disappearance-act

39. **Mobley v. Workday, Inc.** No. 3:23-cv-00770, U.S. District Court for the Northern District of California. Disparate-impact claims regarding AI-driven hiring screening; motion-to-dismiss ruling July 2024; preliminary collective certification under the ADEA granted May 2025; subsequent rulings through mid-2026 that the ADEA reaches job applicants, and a June 2026 order sustaining FEHA and disability-related proxy claims. Docket via CourtListener: https://www.courtlistener.com/docket/66831340/mobley-v-workday-inc/

40. **Carrier adoption filings.** More than 60 property and casualty insurance groups had filed to adopt AI exclusions across several lines by mid-2026, with some telling the publication they did not immediately plan to implement them. The Insurer, July 23, 2026. https://www.theinsurer.com/ti/analysis/more-than-60-pc-insurance-groups-file-to-adopt-ai-exclusions-2026-07-23/

41. **Insurtech funding data.** Global insurtech funding of approximately $1.63 billion in the first quarter of 2026, with AI-focused companies capturing approximately 95% of capital deployed. Gallagher Re, Global InsurTech Report for Q1 2026, May 7, 2026. https://www.ajg.com/gallagherre/news-and-insights/global-insurtech-report-for-q1-2026/

42. **AI insurance market projection.** Deloitte projects annual global AI insurance premiums of about $4.8 billion by 2032, growing at around 80% a year. Deloitte Center for Financial Services, Deloitte Insights, August 4, 2025. https://www.deloitte.com/us/en/insights/multimedia/videos/ai-insurance-market-potential.html

43. **Great American Risk-e Business policy form.** Base-form definitions that expressly address artificial intelligence and deepfake causation in the Cyber Crime Incident, Network Security Incident and Privacy Incident triggers, with social engineering and transfer of funds sitting inside the cyber crime definition. Great American publishes the policy on its cyber risk applications and forms page: declarations CY 6007 05 26 and policy form CY 4006 05 26. The edition posted publicly is the edition reviewed. https://www.greatamericaninsurancegroup.com/docs/default-source/cyber-risk/risk-e-business-cyber-loss-liability-insurance-policy.pdf | https://www.greatamericaninsurancegroup.com/about-us/business-operations/division/cyber-risk/applications-forms

44. **CFC affirmative AI coverage.** Program of updates embedding explicit AI-related language across seven CFC products, including technology errors and omissions and cyber, in place of silent or implied coverage. CFC, June 25, 2026. https://www.cfc.com/en-au/knowledge/news/2026/06/cfc-responds-to-customer-demand-for-affirmative-ai-cover/

45. **Liberty Mutual Cyber Resolution and Tech Resolution.** Global cyber product suite whose base forms include affirmative coverage for quantum computing and generative AI. Liberty Mutual Insurance, September 30, 2025. https://www.libertymutualgroup.com/about-lm/news/articles/liberty-mutual-insurance-launches-global-cyber-products

46. **AXA XL generative AI endorsement.** Endorsement to the CyberRiskConnect policy addressing data poisoning, usage rights infringement and regulatory violations including under the EU AI Act; launched October 21, 2024, with the January 2026 edition being the version described in the text. AXA XL. The sublimit, carve-out and exclusivity provisions described are drawn from the endorsement itself. https://axaxl.com/press-releases/axa-xl-unveils-new-cyber-insurance-extending-coverage-to-help-businesses-manage-emerging-gen-ai-risks

47. **Beazley AI Clarifying Endorsement.** Endorsement announced September 17, 2026 affirming that AI-driven cyber attacks fall within the carrier's existing cyber and technology errors and omissions cover. Insurance Journal, September 17, 2026, at the first link; Insurance Business, September 17, 2026, on the endorsement's treatment of attacker-side AI, at the second. https://www.insurancejournal.com/news/national/2026/09/17/885463.htm | https://www.insurancebusinessmag.com/us/news/cyber/what-beazleys-ai-cyber-endorsement-does-and-doesnt-cover-590125.aspx

A note on sourcing.

Every carrier described in Section 4 has published material of its own on the coverage position attributed to it, cited in the notes above. Structural detail beyond those documents reflects the author's review of policy forms and endorsements in the ordinary course of placing business. No policy language is reproduced. Carriers are referred to by letter as an editorial convention, not a confidentiality measure.

For further analysis of related coverage questions in the current cyber threat environment, see RT ProExec's recent publications: "Iran, Cyber War, and Your Policy: War Exclusions, Coverage Implications, and What Policyholders Need to Know Now," "Attribution: Who Decides if it's an Act of War?" "Wiper vs. Ransomware: Same Policy, Completely Different Claim," "When the Management Tool Becomes the Weapon: MDM Weaponization and the Cyber Coverage Question," and "Dependent BI: The Vendor You Didn't Insure."

This article is for general information purposes only and does not constitute legal, coverage, or professional advice. The information herein is based on publicly available reporting and the author's review of carrier policy specimens and forms, regulatory filings, and court records as of September 2026, and is subject to change as the situation evolves. References to specific carriers, regulatory filings, court cases, or settlement amounts are based on public sources and are presented for illustrative analytical purposes; they do not constitute factual representations about any particular carrier's underwriting position or any specific party's legal exposure. No warranties, promises, or representations of any kind, express or implied, are given as to the accuracy, completeness, or timeliness of the information provided in this article. Every insured's circumstances differ, and coverage needs and priorities vary based on an insured's unique risk profile and operations. Whether a loss is covered by insurance depends on the specific facts of the loss and the terms and conditions of the actual insurance policy or policies involved. References to typical coverage provisions or market approaches are illustrative only and may not apply to a particular policy or situation. Carriers are referred to by letter in Section 4; the analytical observations reflect the author's independent review and do not represent endorsements or criticism of any specific carrier. No user should act on the basis of any material contained herein without obtaining proper legal or other professional advice specific to the user's situation.

RT ProExec is a part of the RT Specialty division of RSG Specialty, LLC, a Delaware limited liability company based in Illinois. RSG Specialty, LLC, is a subsidiary of Ryan Specialty, LLC. RT ProExec provides wholesale insurance brokerage and other services to agents and brokers. RT ProExec does not solicit insurance from the public. Some products may only be available in certain states, and some products may only be available from surplus lines insurers. In California: RSG Specialty Insurance Services, LLC (License #0G97516). ©2026 Ryan Specialty, LLC