AON

AI Fact Sheet 2026

Artificial Intelligence Risk Management 2026

Artificial intelligence ("AI") technology, products and services have developed much faster than AI risk management. As a result, to the extent that evolving AI perils might be covered by existing insurance policies, such coverage falls over ninety percent (90%) under "Silent AI" insurance.[1] Silent AI coverage means that the underlying policy neither affirmatively covers AI nor excludes AI – the policies don't mention AI. Therefore, we are in a similar position to that of cyber perils in the early 2000's whereby insureds attempted to fit cyber losses into property, general liability, professional liability, employment practices liability, crime, etc. policies – none of which mentioned cyber ("Silent Cyber" insurance).

Insurers responded by adding specific language to legacy policies to clarify that most cyber perils were not intended to be covered, which led to the birth of standalone cyber insurance policies. It took the Lloyds of London Market Association eighteen (18) years of coverage ambiguity to finally mandate that all policies must either affirmatively specify cyber coverage grants or include specific cyber exclusions (no more "Silent Cyber" insurance). If we leverage the lessons learned from cyber, we can expedite the AI risk management improvement curve.

Artificial Intelligence (AI) Historical Timeline

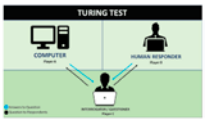

Alan Turing ("Turing Test")

1950

1954

Term **"artificial intelligence"** coined at the Dartmouth Summer Research Project on Artificial Intelligence. Organized by John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon

2012

Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton introduced AlexNet, a deep convolutional neural network that dominated the ImageNet competition, launching the modern era of deep learning

2013

First commercial artificial intelligence policy placed

2016

DeepMind's AlphaGo defeated world champion Lee Sedol in the game of Go, a massive achievement in artificial intelligence and reinforcement learning

2017

A group of researchers at Google published a paper called "Attention is All You Need" that introduced the "transformer" (AI parses a sentence all at once instead of sequentially)

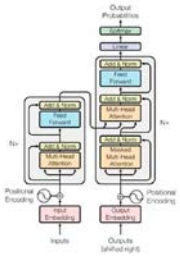

2022

ChatGPT 3.5 Released

2024

2024 Nobel Prize in Physics "for foundational discoveries and inventions that enable machine learning with artificial neural networks." **Geoffrey Hinton & John Hopfield**

2024 Nobel Prize in Chemistry "for protein structure prediction" (AI model called AlphaFold2). **Demis Hassabis & John Jumper**

2025

- **EU AI Act** Effective Feb 2, 2025
- **U.S. AI Action Plan** Jul 2025
- **U.S. Executive Order** Dec 11, 2025

2026

Insurance addresses AI in one of three manners:

Silent	Affirmative Coverage	Specific Exclusions
Creates ambiguity like the early days of cyber	Via endorsement to existing or new standard AI policy	

The insurance industry plays a pivotal role with respect to such novel perils. Insurance does more than transfer risk after the fact; it rewards stronger risk controls, clarifies standards of care, validates operational discipline and translates risk into transparent price signals. While cyber insurance carriers have accumulated sufficient data to build accurate cyber loss models, we do not yet have the same amount and quality of AI frequency and loss information.

Therefore, Aon, via development of its Aon AI Risk Diagnostic Tool, is leading the effort to rapidly collect and analyze AI loss data correlated to risk management. Insurance can act like a real time, dynamic and fluid "prediction market"[2] regarding the frequency and severity of AI perils as legacy and new policies clarify the scope of AI coverage and premium. As AI adoption accelerates, risk managers and brokers are being asked a harder question: What actually responds when an AI system causes harm?[3] Clarity beats ambiguity — especially at scale.[4]

"Roadmap To Address AI Exposures"

AI risk management must be dynamic and fluid because AI perils are dynamic and fluid.

This fact sheet sets forth an initial roadmap to help organizations become:

Better Informed

About AI risk management scaled to AI growth; AI perils by industry; and AI case law trends.

Better Advised

Regarding sources of AI liability; risk mitigation best practices; how legacy and new AI specific insurance policies may respond.

Better Decisions

In lowering the Total Cost of AI Risk and optimal capital allocation around AI governance.

Better Informed[5]

Growth of artificial intelligence by AI deployers[6]

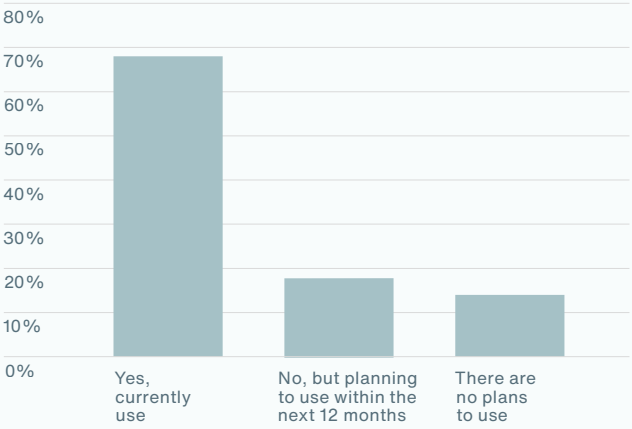

Most organizations **(86 percent of respondents)** use or intend to use artificial intelligence (AI)products or services.

- Average enterprise AI spending per organization hit $7 million in 2025, a 180% jump from $2.5 million in 2024, with projections of $11.6 million for 2026.[7]
- Hyperscalers are planning to pour $1.3 trillion into AI infrastructure through 2027[8] As a percentage of GDP, the projected 2026 spending of four tech giants rivals the most momentous capital efforts in U.S. history.[9]
- In order to position itself as a leader in artificial intelligence, a world-wide leading consulting firm is linking use of AI tools to promotion potential and will exit employees who are not adopting the use of AI at work.[10]
- KPMG threatened to find a new auditor if Grant Thornton did not pass on savings from AI, leading to a 14% reduction in the audit fee, dropping from $416,000 in 2024 to $357,000 in 2025.[11] Theoretically, if AI increases accuracy, efficiency and productivity in the insurance industry, then it could reduce insurance market costs and could enable insurance carriers to reduce premiums over time (All other things being equal).
- The standalone artificial intelligence insurance market is expected to become a $4.8 billion annual market by 2032.[12]
- More than 90% of organizations see a need for insurance coverage tailored to AI/Generative AI threats.[13]
- Over two thirds would pay at least 10% more in premiums for explicit insurance policy extensions that cover Gen/AI related risks.[14]

> “
>
> AI data centers are the insurance market’s biggest-ever opportunity.
>
> **Joe Peiser, Aon Risk Capital CEO**

Monitor evolving AI litigation relative to your organization’s AI use

Aon aggregated AI related claims data from multiple sources (250+ cases as of December 2025). We then map the AI cases by category to existing insurance lines of insurance (there can be overlap and gaps so the following does not add up to 100%):[15]

- Approximately 15% of current AI claims can be addressed by standalone cyber insurance.
- Approximately 80% of current AI claims can be addressed by Errors & Omissions/Professional Liability/Media Liability Insurance – a lot of copyright infringement and Advertising Injury and Personal Injury claims to date, but less than 10% of organizations purchase E & O (typical “named risk” risk E & O policies cover approximately 58% of AI claims compared to 80% addressed by customized “all-risk” E & O).
- Approximately 20% of current claims can be addressed by other lines of insurance, such as Director’s & Officer’s (“AI washing”); Employment Practices Liability Insurance (employment related discrimination); Crime (deepfake financial fraud/ funds transfer); General Liability; Intellectual Property; Property; etc. Note that Insurance Services Office (ISO), a Verisk business and major provider of standard policy forms, introduced new optional endorsements to exclude Generative Artificial Intelligence exposures from Commercial General Liability policies.[16]

- Potential AI claims addressed by General Liability/ Product Liability/etc. insurance may grow faster than all other categories within the next few years due to:
  - Increase in AI in robotics/autonomous vehicles/ drones/Internet of Things appliances, etc. (Jury Finds Tesla Autopilot Defective in Landmark $243 Million Verdict – What It Means for Future Cases).[17]
  - Copyright and other AI/PI issues will increase in predictability once the courts establish liability – behavior will change (even if there are a few big verdicts/settlements).
  - Bodily injury and tangible property damage injuries are currently infrequent, but we anticipate will grow materially relative to intangible injury claims; (Tesla factory technician sues for $51 million after assembly-line robot knocks him unconscious)[18] Compare to growth of "Ransomware" claims in cyber insurance relative to privacy/data breach claims.

AI deployment and exposures are reaching every industry and geography

AI Risk Perils Across Various Aon Industry Verticals

Construction and Real Estate
- AI-driven safety prediction failures
- Autonomous equipment malfunction
- Model errors in project scheduling or cost estimation
- Algorithmic flaws in property valuation models
- Cyber vulnerabilities

Financial Institutions
- Algorithmic bias and unfairness in credit and pricing
- Model governance and regulatory compliance failure
- Algorithmic trading instability from model errors
- Synthetic identity and deepfake fraud

Financial Sponsors
- Model-driven investment blind spots
- Portfolio data misuse
- AI risk concentration
- Reputational spillover

Food, Agribusiness, and Beverage
- Yield and quality model errors
- Food safety surveillance gaps
- Environmental and resource mis-management
- Farm data ownership disputes

Healthcare Providers and Services
- AI-assisted diagnosis
- Biased diagnosis and triage
- Safety-critical clinical errors
- Sensitive PHI data leakage
- Regulatory and documentation gaps

Hospitality, Travel, and Leisure
- Discriminatory pricing and access
- Guest surveillance and privacy
- Service automation failures
- Reputation and review manipulation

Industrial and Manufacturing
- AI-driven quality control failures
- Predictive maintenance errors
- Adversarial attacks on robotics or autonomous machinery
- Process-automation failures

Insurance
- Algorithmic underwriting bias
- AI-driven claims-automation errors
- Model governance failures
- Fraud detection inaccuracies
- Regulatory non-compliance

Life Sciences
- Biased clinical algorithms
- Regulatory non-compliance
- Data integrity and provenance issues
- IP and trade secret exposure

Natural Resources
- Incorrect geological or seismic predictions
- Autonomous equipment failure
- Cyberattacks on operational technology (OT)
- Model bias or drift in resource-yield forecasting

Professional and Business Services
- AI hallucinations in research, legal analysis, or advisory
- Unintentional exposure of confidential client data
- Intellectual property risks
- Automated financial or analytical model errors

Public Sector
- Biased public services
- Opaque automated decisions
- Civic misinformation and trust erosion
- Critical infrastructure attacks

Retail and Consumer Goods
- Bias in personalized recommendation engines
- AI-powered fraud or synthetic review manipulation
- AI forecasting causing inventory shortages, supply chain disruption

Sports and Entertainment
- Fan data exploitation
- Deepfake and likeness misuse
- Integrity and match-fixing risk
- Content moderation reputation

Technology, Media, and Communications
- AI-generated misinformation
- Content moderation failures
- Security model exploitation
- Surveillance and privacy beyond legal boundaries

Transportation and Logistics
- Autonomous system failures
- Real-time decision errors
- Navigation and sensor spoofing
- Workforce impact and disputes

Better Advised[19]

Identify and Mitigate the Sources of AI Exposures

Insurance underwriting, combined with statutes, litigation precedent, contractual allocation of liability, and evolving standards, such as the National Institute of Standards and Technology ("NIST") AI Risk Management Framework,[20] set thresholds for what are considered acceptable risks. While national security and geopolitical competition are driving approaches to national AI regulation,[21] there are additional key policy issues confronting leadership at both the federal and state level, which can be addressed by private sector AI risk management if properly dosed.[22]

As AI adoption accelerates, organizations face a dual imperative: maximizing AI-driven ROI while managing emerging risks.

Should AI product/service providers be liable for anything that goes wrong, such as an AI user suicide or the AI development or production of chemical, biological, radiological, and nuclear threats? Should deployers of AI technology be responsible for AI Frontier Model product/service hallucinations, deepfake fraudulent funds transfers and bodily injury or tangible property damage from IoT/robotic/autonomous vehicle deployments?

AI forecasters predict that up to 20% of large enterprises will see lawsuits and fines due to inadequate AI agent controls by 2030.[23] Today's state of AI governance makes that a likely reality (or worse) unless we improve AI governance immediately. Case law on these questions in the context of AI is developing, but a forthcoming article in the Stanford Journal of Law, Economics & Business suggests juries may treat AI-enabled conduct more harshly in legal disputes, resulting in what the authors call an "AI penalty." In an experiment inspired by a case involving automated-bot scraping, the authors found that factfinders may be more inclined to assign liability and grant higher damages when defendants use AI.[24]

To test whether organizations are implementing risk management best practices related to AI liability and insurance coverage, insurance carriers are developing AI insurance application questions, such as the following representative sample:[25]

- Do you have an AI Governance Framework with Board/Management oversight, which includes specific use cases, restrictions and authorization policies?
- Is the Insured developing and providing Generative AI solutions to customers, or is the use for internal purposes only?
- Are you confirming whether any key vendors are utilizing AI, and if so, what safeguards are in place to prevent errors? Have you conducted pre-contract diligence?
- Disclosure — are notices provided to consumers on the use of AI?
- How do you prevent unintended bias when utilizing AI?
- Does a human need to verify accuracy before AI takes action?
- Are AI actions logged so potential errors can be reviewed and remediated?
- Contracting:
  - Who owns and has what rights to use the prompts (inputs) and outputs?
  - Are the prompts and outputs subject to confidentiality obligations?
  - What is the allocation of liability, including hold harmless and indemnity?
  - Are there Service Level Agreements?
- Are there separate agreements for plugins and APIs?
- Are there requirements for evidence of insurance, including limits and representations that the insurance adequately covers the AI services/products being provided?
- Representation that AI provider is following applicable laws?[26]
- Intellectual Property:
  - Please provide an overview of controls or guardrails in place to prevent infringement of IP when using GenAI or LLMs in the Insured's business.
  - Are the GenAI or Large Language Models (LLMs) used all in-licensed from third party suppliers?
  - If yes, please confirm whether indemnities are provided in respect of IP infringement by the licensor, and whether such indemnities carry insurance backing.
  - Have the GenAI or LLMs used been developed or enhanced in-house by the Insured?
    - If yes, please confirm whether any data that is not owned by the Insured has been used for input/training (incl. any personal data)?
    - If any other data has been used, please provide details.

Map Your Organization's AI Risks to Existing Insurance as AI Exclusions and sub-limit provisions proliferate[27]

	Peril	Media Liability	Tech E&O, MPL, PI	General Liability	Intellectual Property	Cyber	Crime	D&O	EPL	AI Specific
Legacy Perils	Professional Services/technology errors, omissions or negligent acts	Limited to Media orgs								
	Copyright (Output) or Trademark		Can be added	Except media orgs		Requires cyber trigger				
	Patent Infringement and Trade Secrets	Not in base form								
	Discrimination and Bias (employment related – AI Deployer)									
	Defamation, Libel, Slander		Can be added	Except media orgs						
	Bodily Injury		Base forms exclude	Products only		Contingent BI can be added				By special endorsement
	Tangible Property Damage		Base forms exclude	Products only		Contingent PD can be added				By special endorsement
	Privacy and Security breaches liability									
	Loss of Financial Assets/ Market Manipulation									
	Product Liability/Recall									
	Business Interruption		Only when cyber added							
New AI Perils	Copyright (Input AI Training)	Limited to media perils								
	AI Regulation Violations		Can be added			Only for privacy/ security		Some coverage for FTC/ SEC, etc.		To extent allowable by law
	Breach of AI Warranty/ Performance Guaranty		Only if due to negligence,	Only via customization						Varies by solution
	Breach of DandO Duties (AI Washing)									Partially included
	Discrimination and Bias (AI Provider)					Requires cyber trigger			Not for non-EPLI	
	Autonomous Vehicles/ Robotics/IoT		Requires bespoke			Requires cyber trigger				
	Deepfake	Limited to media perils	Excludes funds transfers			Excludes funds transfers	For funds transfers			Excludes funds transfers
	Hallucinations	Limited to media perils				Limited to privacy/ security				

- Affirmative AI coverage
- Available, but silent as to AI
- Ambiguous, unless specifically added
- Excluded

Better Decisions

AI experts estimate a varied potential range of AI frequency and severity losses. A white paper titled: "AI 2027"[28] posits a range of scenarios involving AI systems surpassing human-level intelligence, and actions by policymakers could determine whether the consequences are catastrophic. Geoffrey Hinton, the widely considered "Godfather of AI," sees two main future possibilities for the relationship between humans and AI's: either AI's take over and humanity goes extinct, or AI amplifies human potential and greatly improves healthcare, education, scientific discovery, productivity, and the economy; it democratizes knowledge, helps address climate change, and enhances standards of living. In Hinton's dichotomy, avoiding the catastrophe of the former, and realizing the benefits of the latter depends heavily on governance, ethical design, and avoiding misuse.[29]

During a Federal Reserve conference in July 2025, OpenAI CEO Sam Altman warned banking regulators of a "significant, impending fraud crisis" driven by AI's ability to impersonate voices.[30] While the focus was on financial fraud, Altman also raised broader concerns that AI is poised to disrupt labor markets faster than institutions can adapt. He likened the pace and breadth of potential upheaval to globalization, but compressed into a timeframe of months instead of decades. If Anyone Builds It, Everyone Dies is the title of a September 2025, book by Eliezer Yudkowsky and Nate Soares. The 'it' in question is superintelligence/Artificial General Intelligence ("AGI") built on some form of the current AI paradigm, and the authors very much mean this literally.[31] However, the probability of doom (i.e., the probability of catastrophic, human-extinction-level outcomes caused by advanced AGI) referred to by AI computer scientists as p(doom), is not preordained. There are multiple public-private collaboration initiatives that can, and would, if implemented, lower the p(doom) and facilitate AI growth.[32] In fact, some AI experts believe that we can lower the p(doom) to effectively zero.[33]

From an insurance perspective, we need to leverage the best available quantifiable, objective, fact-based information, which means analyze the measurable range of possible outcomes for each AI scenario. For example, Aon's Resilience Quotient is an analytical framework and a tool that integrates Aon's proprietary Risk Capital and Human Capital data together with sentiment analysis from Gallup, which provides a structured way to understand how risk compounds across systems, how resilience is built and activated and where targeted actions can most effectively influence performance.[34] Specifically, we can apply various risk analyzers[35] to map the frequency and severity of AI-related events across sectors by line of insurance (e.g. Cyber, Property, Casualty and Directors' & Officers'). Through these lenses, organizations can implement a range of best practices, including maintaining an AI model inventory, quantifying exposures through scenario modeling, conducting end-to-end AI system audits, performing third-party vendor due diligence, implementing unacceptable bias detection and validation testing, defining contractual indemnities, and appointing dedicated governance leads for AI deployment. These practices help translate risk into measurable risk mitigation action. For potential aggregated, correlated, systemic risks, insureds may need to consider alternative risk transfer, such as a captive or catastrophe bonds. In fact, insurance linked securities are being explored for hyperscale data center risk and the world's first data center reinsurance treaty has been placed.[36]

Technological innovation throughout history has been maximized when industry and consumers have confidence that novel developments are balanced with appropriate risk mitigation and controls. Superior AI risk management means incrementally calibrating safeguards to match AI's maturity and impact, using feedback signals to avoid under- or over-regulation and governance. AI risk management isn't just "brakes" — it's like shock absorbers on a race car. By managing shocks and maintaining control, organizations can drive faster, more safely, and for longer, maximizing the acceleration of AI innovation. Better access to AI insurance markets translates to lower premiums, which frees up capital for research & development and scaling. Done well, this creates the optimal trust–innovation balance point, accelerating AI growth responsibly.

Insurance does more than transfer risk; it incentivizes behavioral changes, imposes performance standards, and validates operational discipline. If an AI safety implementation is effective at reducing AI risk, then insurers will broaden insurance coverage robustness and adjust premium rates accordingly. This creates a virtuous cycle: good governance leads to better insurability, which in turn supports innovation and insureds' protection.

Since artificial intelligence opportunities and perils are dynamic and fluid, insurance underwriting safeguards require constant refining and optimization. We can scale incremental guardrails to the maturity and criticality of AI applications. For example, higher-risk AI like healthcare, defense, & autonomous robotics, such as drones & vehicles should face stricter oversight than low-risk AI like marketing automation. By deploying an iterative risk calibration, insurance underwriters, regulators and organizations can pilot-test AI risk frameworks, observe impacts, and then adjust AI governance in phases. Indicators show when equilibrium is near. Feedback loops (user trust metrics, incident reporting, adoption rates) signal whether risk management is encouraging or choking growth. For example, a rise in AI adoption rates alongside stable risk incidents suggests governance is well-balanced. Conversely, flat AI adoption with rising compliance costs signals over-scrutiny.

AI will remain a defining megatrend for at least the next decade. Organizations that understand their AI perils and risk domains, leverage insurance intelligently, and embed AI into enterprise risk and capital decisions will be better positioned not only to avoid downside, but to capture AI's upside with confidence.

Historical Technology Adoption Across the Titration Curve

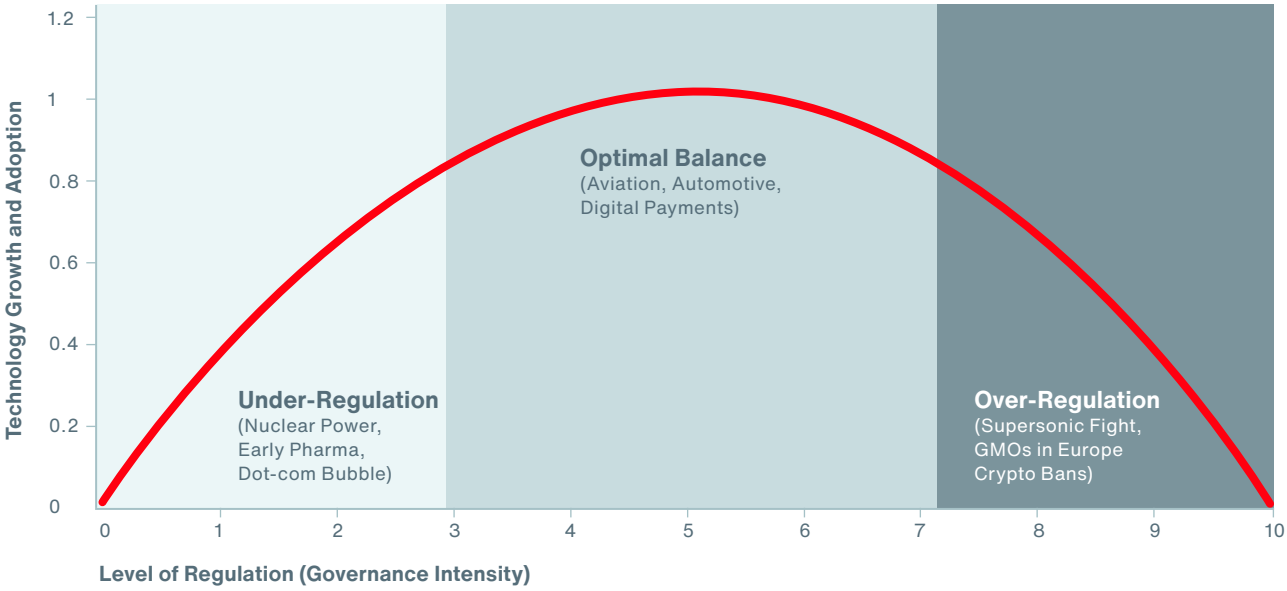

Endnotes

1 As of March 2026.

2 Making Extreme AI Risk Tradeable. If traditional insurance can't handle the extreme risks of frontier AI, catastrophe bonds can cover the gap and compel labs to adopt tougher safety standards. https://ai-frontiers.org/articles/ai-catastrophe-bonds-extreme-risk-tradeable

3 Expect An Uptick in Artificial Intelligence Insurance Related Claims in 2026: https://www.olshanlaw.com/Advertising-Law-Blog/expect-an-uptick-in-artificial-intelligence-insurance-related-claims-in-2026

4 The practical risk for buyers and carriers is not just "non-coverage," it is slow claims resolution driven by interpretive disputes. If AI is business-critical, clarity in triggers, exclusions, and documentation expectations stops being a nice-to-have and becomes basic resilience.

5 2026 AON Ponemon Intangible Risk Comparison Report

6 2026 AON Ponemon Intangible Risk Comparison Report

7 https://www.roic.ai/news/jpmorgan-ai-fears-overdone-in-software-02-10-2026?utm_source=substack&utm_medium=email; If traditional insurance can't handle the extreme risks of frontier AI, then catastrophe bonds might be able to cover the gap and compel labs to adopt tougher safety standards: https://ai-frontiers.org/articles/ai-catastrophe-bonds-extreme-risk-tradeable

8 https://fortune.com/2026/02/10/stocks-2-trillion-software-wipeout-ai-bull-market/

9 https://www.wsj.com/tech/ai/ai-spending-tech-companies-compared-02b90046

10 Accenture 'links staff promotions to use of AI tools' | AI (artificial intelligence) | The Guardian

11 https://thefinancestory.com/kpmg-forced-auditor-for-a-14-ai-discount-will-the-billable-hour-die. This, and similar cases, highlight a growing debate in professional services where clients demand that efficiency gains from AI (speeding up document triage, risk assessment, etc.) should be reflected in lower fees.

12 https://www.deloitte.com/us/en/insights/multimedia/videos/ai-insurance-market-potential.html

13 https://www.genevaassociation.org/publication/digital-ai-transformation/gen-ai-risks-businesses-exploring-role-insurance

14 https://www.genevaassociation.org/publication/digital-ai-transformation/gen-ai-risks-businesses-exploring-role-insurance; Businesses Want to Insure Against AI Risks as Adoption Soars - Risk & Insurance : Risk & Insurance

15 Aon Aggregated Artificial Intelligence Litigation Database (GW Dail Database: ~300 AI-related lawsuits extending back to 2010, Stanford Class Action Database: AI-related class action lawsuits not captured in the GW Dail database beginning in 2010, Output provides directional analysis to understand which insuring agreements may address these types of lawsuits)

16 https://insuranceintel.substack.com/p/ai-coverage-gets-cut.
New Endorsement Forms: ISO introduced two primary optional endorsements:
- **CG 40 47** – Exclusion – Generative Artificial Intelligence (for use with both occurrence and claims-made CGL policies).
- **CG 40 48** – Exclusion – Generative Artificial Intelligence (specifically focusing on Coverage B, personal and advertising injury).

These exclusions address liabilities from AI-generated content like defamation or intellectual property infringement. Some insurers are also using "Absolute" AI exclusions for broader AI use, reflecting concerns about AI's unpredictable risks. This trend suggests a move toward dedicated AI insurance products, meaning businesses may need specialized coverage instead of relying on CGL policies for AI-related claims. The following is a sample of new policy wordings introduced or proposed by insurers to limit or exclude coverage for AI-related losses:
- **Hamilton Select Insurance Inc.**: Excludes claims from use of generative AI in professional liability policies
- **Philadelphia Indemnity Insurance Company**: Excludes offenses created by generative AI
- **Cincinnati Insurance Company**: D&O policies exclude coverage for development, deployment, or use of AI
- **Berkley Insurance Company**: "Absolute AI exclusion" broadly bars coverage for any use or development of AI

17 Jury Finds Tesla Autopilot Defective in Landmark $329 Million Verdict – What It Means for Future Cases

18 Tesla factory technician sues for $51 million after assembly-line robot knocks him unconscious

19 When Insurance Won't Cover AI: Why Carriers Are Adding Exclusions, And Why AI Governance Is Now Essential

20 https://www.nist.gov/artificial-intelligence

21 Policy Alert: New U.S. Executive Order on Artificial Intelligence – Aon Tips for Better Risk Capital Decisions

22 Titration of Artificial Intelligence Risk Management to Accelerate AI Growth

23 IDC | AI & GenAI Predictions: Key Insights for 2025 and Beyond - eBook

24 https://assets.law360news.com/2443000/2443289/ssrn-5610270%20(1).pdf

25 In the first quarter of 2026, a familiar pattern sharpened into something more operational, AI risk is increasingly being priced, benchmarked, and litigated in parallel. The Lloyd's Market Association published concrete scenario thinking on AI loss pathways, while researchers and journalists surfaced fresh evidence that frontier model "helpfulness" still cracks in high-stakes contexts like lab safety and health information. At the same time, courts are now explicitly signaling that AI-assisted errors are not just a lawyer problem, they are an access-to-justice problem, and hiring tech is facing a new round of legal scrutiny over opaque scoring. Evaluation artifacts are becoming the lingua franca across underwriting, procurement, and oversight, and teams that cannot produce credible evidence of testing, controls, and governance will find decisions being made for them.

26 Titration of Artificial Intelligence Risk Management to Accelerate AI Growth: 2025+Hatch+Policy+Review.AI+FINAL+DIGITAL.pdf

27 When Insurance Won't Cover AI: Why Carriers Are Adding Exclusions, And Why AI Governance Is Now Essential; The Evolving Landscape of AI Insurance: Empirical Insights into Risks and Policy Gaps; The insurability of EU AI Act Fines is uncertain; How insurers are wording AI exclusions

28 Daniel Kokotajlo, Scott Alexander, Thomas Larsen, Eli Lifland & Romeo Dean, AI 2027, AI Futures Project (Apr. 3, 2025), https://ai-2027.com/

29 Hinton's collaborators, Ilya Sutskever and Yoshua Bengio, share similar views although their estimates of the p(doom) vary.

30 Smarrtif AI, Sam Altman Warns: AI Will Eliminate Jobs and Threaten National Security If Left Unchecked, LinkedIn (Aug. 7, 2025), https://www.linkedin.com/pulse/sam-altman-warns-ai-eliminate-jobs-threaten-national-security-8buvc/

31 If any company or group, anywhere on the planet, builds an artificial superintelligence using anything remotely like current techniques, based on anything remotely like the present understanding of AI, then everyone, everywhere on Earth, will die. See Eliezer Yudkowsky, X, https://x.com/esyudkowsky.

32 Policy Alert: New U.S. Executive Order on Artificial Intelligence – Aon Tips for Better Risk Capital Decisions

33 Yann LeCun, a prominent AI researcher and Chief AI Scientist at Meta, believes the probability of an AI-caused existential catastrophe is effectively zero. His skepticism about existential AI risk contrasts sharply with the views of others in the field, including fellow Turing Award winner Geoffrey Hinton. See Christopher Mims, This AI Pioneer Thinks AI Is Dumber Than a Cat, Wall St. J. (June 20, 2024) (updated Oct. 11, 2024), https://www.wsj.com/tech/ai/yann-lecun-ai-meta-aa59e2f5. Jenson Huang, NVIDIA founder, CEO and President, shares similar views.

34 https://www.aon.com/en/insights/reports/aon-resilience-quotient

35 https://www.aon.com/en/capabilities/risk-analytics/overview/cyber-risk-analyzer

36 https://www.reinsurancene.ws/aon-designed-placed-worlds-first-data-centre-reinsurance-treaty-says-ceo-case/

About Aon

Aon plc (NYSE: AON) exists to shape decisions for the better — to protect and enrich the lives of people around the world. Through actionable analytic insight, globally integrated Risk Capital and Human Capital expertise, and locally relevant solutions, our colleagues provide clients in over 120 countries with the clarity and confidence to make better risk and people decisions that help protect and grow their businesses.

Follow Aon on LinkedIn, X, Facebook and Instagram. Stay up-to-date by visiting the Aon's newsroom and sign up for news alerts here.

aon.com

Contact Us

Jesus Gonzalez
Intangible Assets Global Collaboration Co-leader
+1 312 381 4138
jesus.gonzalez.2@aon.com

Aon Artificial Intelligence Deals Desk
aonartificialintelligencedealsdesk@aon.com