Generative AI and Cyber Liability: Recent U.S. Legal and Insurance Developments

By Richard S. Dukes, Jr.

Shareholder, Turner Padget Graham & Laney

With assistance from Artificial Intelligence.

Introduction

Generative AI systems like large language models and image generators have seen explosive adoption across industries, bringing unprecedented capabilities – and novel liability risks. Unlike traditional software that follows predefined rules, generative AI can produce unpredictable outputs ("hallucinations") that may be false, infringing, or harmful (Judge Denies Motion to Dismiss AI Defamation Suit | Alerts and Articles | Insights | Ballard Spahr) (Judge Denies Motion to Dismiss AI Defamation Suit | Alerts and Articles | Insights | Ballard Spahr). Businesses using these tools face potential legal exposure in areas ranging from data privacy breaches and intellectual property (IP) infringement to defamation and professional malpractice. U.S. regulators and courts have begun grappling with these issues, while insurers and policyholders are evaluating how existing coverage (cyber, E&O, D&O, etc.) applies to AI-related risks. This report surveys recent U.S. legal and regulatory developments involving generative AI and cyber liability, and examines how insurance policies might respond or exclude these emerging exposures. It is intended for in-house counsel at insurance carriers to understand the evolving landscape and coverage implications.

Regulatory and Legislative Developments in the U.S.

Government agencies have stepped up scrutiny of AI to protect consumers and markets. The Federal Trade Commission (FTC) in 2024 launched an enforcement sweep dubbed "Operation AI Comply" targeting companies that allegedly used AI for **deceptive or unfair practices** (New FTC Initiative Targets Deceptive AI Claims and a Generative AI Service | Insights | Skadden, Arps, Slate, Meagher & Flom LLP) (New FTC Initiative Targets Deceptive AI Claims and a Generative AI Service | Insights | Skadden, Arps, Slate, Meagher & Flom LLP). For example, the FTC pursued firms for overstating the capabilities of AI products and even took action against an AI startup, DoNotPay, for claiming its chatbot could replace human lawyers when it **did not work as promised** (New FTC Initiative Targets Deceptive AI Claims and a Generative AI Service | Insights | Skadden, Arps, Slate, Meagher & Flom LLP) (New FTC Initiative Targets

Deceptive AI Claims and a Generative AI Service | Insights | Skadden, Arps, Slate, Meagher & Flom LLP). The FTC signaled that making **unsupported claims about AI** or providing AI tools that facilitate fraud can violate Section 5 of the FTC Act (New FTC Initiative Targets Deceptive AI Claims and a Generative AI Service | Insights | Skadden, Arps, Slate, Meagher & Flom LLP). In another first-of-its-kind action, the U.S. Equal Employment Opportunity Commission (EEOC) settled its **first AI bias lawsuit** in 2023, involving a recruitment algorithm that allegedly **discriminated against older job applicants** (Tutoring firm settles US agency's first bias lawsuit involving AI software | Reuters) (Tutoring firm settles US agency's first bias lawsuit involving AI software | Reuters). The AI-driven hiring software had automatically filtered out women over 55 and men over 60, in violation of age discrimination laws. The employer paid $365,000 to settle the EEOC's claims, without admitting wrongdoing (Tutoring firm settles US agency's first bias lawsuit involving AI software | Reuters) (Tutoring firm settles US agency's first bias lawsuit involving AI software | Reuters). This enforcement trend highlights that companies deploying AI must ensure compliance with consumer protection and anti-discrimination laws, or face regulatory action.

Regulators are also pushing for transparency and accountability in AI use. The **Securities and Exchange Commission (SEC)** has emphasized that public companies should disclose material risks related to AI and avoid "AI-washing" (misrepresenting AI capabilities) in statements to investors (Protecting Your Business: AI Washing and D&O Insurance) (Protecting Your Business: AI Washing and D&O Insurance). In Congress and federal agencies, policymakers have discussed frameworks for AI governance (such as risk management guidelines and bills addressing AI accountability), but as of this writing no comprehensive federal AI law has passed. Nevertheless, an "expanding regulatory landscape aimed at protecting shareholders and consumers" is creating new compliance challenges for businesses using generative AI (Insurers Explore New AI Coverage Options, Potentially Filling Coverage Gaps for Policyholders Developing Generative AI | The Policyholder Perspective). The White House's October 2023 Executive Order on safe and trustworthy AI and the NIST AI Risk Management Framework provide guidance that, while not legally binding, indicate best practices (e.g. security testing of AI models, data privacy protections) that regulators may expect organizations to follow. In sum, U.S. regulators are warning that **if AI tools cause harm – whether through faulty outputs, bias, or misuse of data – the responsible companies can be held to account** under existing laws.

Litigation Trends Involving Generative AI

Multiple lawsuits in the past two years illustrate how courts are beginning to address liability arising from generative AI. These cases span a range of claims – from copyright and trademark infringement to defamation, privacy, and even securities fraud – highlighting the diverse risks posed by AI-generated content and decisions.

Intellectual Property Disputes

Copyright infringement is a major flashpoint. A wave of class action lawsuits by authors and artists accuses AI developers of using copyrighted works without permission to train generative models (OpenAI, Microsoft defeat US consumer-privacy lawsuit for now | Reuters). For example, in *Andersen v. Stability AI*, a group of visual artists sued the makers of image generator Stable Diffusion, alleging the system was trained on billions of online images (including their artwork) scraped without consent (Takeaways from the Andersen v. Stability AI Copyright Case | Copyright Alliance) (Takeaways from the Andersen v. Stability AI Copyright Case | Copyright Alliance). In October 2023, the federal court overseeing that case largely **granted the defendants' motions to dismiss** many claims, but allowed the core claim of direct copyright infringement to proceed (Takeaways from the Andersen v. Stability AI Copyright Case | Copyright Alliance). After the plaintiffs amended their complaint, the court in August 2024 issued a detailed order providing early insight into how courts may analyze generative AI's use of training data. Notably, Judge Orrick **denied dismissal of the direct infringement and inducement claims**, finding the plaintiffs plausibly alleged that their copyrighted works are "contained, in some manner" within the AI model's data structure (Takeaways from the Andersen v. Stability AI Copyright Case | Copyright Alliance) (Takeaways from the Andersen v. Stability AI Copyright Case | Copyright Alliance). The court rejected defense analogies to familiar technologies (like VCRs), emphasizing that generative AI is unique and must be evaluated on its own facts (Takeaways from the Andersen v. Stability AI Copyright Case | Copyright Alliance) (Takeaways from the Andersen v. Stability AI Copyright Case | Copyright Alliance). Crucially, the judge held that if a model effectively embeds protected expression as mathematical representations, it could still infringe – a significant win for copyright owners at this stage (Takeaways from the Andersen v. Stability AI Copyright Case | Copyright Alliance) (Takeaways from the Andersen v. Stability AI Copyright Case | Copyright Alliance).

At the same time, courts are setting some limits on IP theories against AI developers. In the *Stable Diffusion* case, all claims under the Digital Millennium Copyright Act (DMCA) were **dismissed with prejudice** because the plaintiffs could not show that any AI outputs were *identical* to their works (Takeaways from the Andersen v. Stability AI Copyright Case | Copyright Alliance) (Takeaways from the Andersen v. Stability AI Copyright Case | Copyright Alliance). The court adopted an "identicality" requirement for DMCA §1202 claims (echoing a ruling in an AI code case, *Doe v. GitHub*), reasoning that output which merely resembles or remixs training data is insufficient (Takeaways from the Andersen v. Stability AI Copyright Case | Copyright Alliance). The judge also found that various state law claims (like unjust enrichment and negligence) were preempted by the Copyright Act, since the harm alleged boiled down to unauthorized copying (Motion To Dismiss Ruling Provides Further Insight Into How Courts View AI Training Data Cases | Insights | Skadden, Arps, Slate, Meagher & Flom LLP). Similarly, in the AI coding assistant case *Doe 1 v. GitHub* (involving Microsoft's GitHub Copilot tool), the court initially required plaintiffs to demonstrate concrete injury by identifying instances where the AI **reproduced their code**. By early 2024, some plaintiffs were able to allege specific examples of Copilot output matching their code, which the court found sufficient to confer standing for a copyright claim (Motion To Dismiss Ruling Provides Further Insight Into How Courts View AI Training Data Cases | Insights | Skadden, Arps, Slate, Meagher & Flom LLP) (Motion To Dismiss Ruling Provides Further Insight Into How Courts View AI Training Data Cases | Insights | Skadden, Arps, Slate, Meagher & Flom LLP). However, that court then agreed that the DMCA claims must be dismissed because the outputs were usually modified and not verbatim copies – the plaintiffs had "pleaded themselves out" of those claims by admitting Copilot rarely outputs identical code (Motion To Dismiss Ruling Provides Further Insight Into How Courts View AI Training Data Cases | Insights | Skadden, Arps, Slate, Meagher

& Flom LLP). These early decisions indicate that while direct copyright claims and contributory infringement claims (e.g. for distributing an infringing model) may survive, courts remain skeptical of more attenuated theories absent clear evidence of **verbatim appropriation**.

Generative AI has also raised **trademark and publicity rights issues**. In a high-profile suit, Getty Images accused Stability AI of not only copying 12 million Getty photos without a license, but also **reproducing Getty's watermark** on some AI-generated images (Getty Images lawsuit says Stability AI misused photos to train AI | Reuters) (Getty Images lawsuit says Stability AI misused photos to train AI | Reuters). Getty argues this could confuse consumers about the images' source, and it asserts trademark infringement alongside copyright claims (Getty Images lawsuit says Stability AI misused photos to train AI | Reuters). No court rulings have yet been issued on the merits in that case (which is pending in Delaware federal court), but it spotlights another novel exposure: AI outputs inadvertently replicating logos or other protected marks. Likewise, generative models that mimic a person's likeness or voice without consent could face **right of publicity** lawsuits. While we have not yet seen major U.S. litigation over AI-generated "deepfakes" or voice clones using a private individual's identity, celebrities and content creators are increasingly wary of such uses. Companies deploying generative AI must be mindful that training data often contains intellectual property, and outputs can implicate rights ranging from **copyright** to **trademarks and likenesses**. As one court observed, generative AI models are "unlike any technologies" in past IP cases, so analogies are imperfect and outcomes will vary case by case (Takeaways from the Andersen v. Stability AI Copyright Case | Copyright Alliance) (Takeaways from the Andersen v. Stability AI Copyright Case | Copyright Alliance). The flurry of ongoing suits (by authors, artists, photo agencies and others) will be closely watched as courts continue to define the boundaries of AI-related IP liability.

Defamation and Misinformation

Generative AI's tendency to produce false information ("hallucinate") has already led to at least one defamation lawsuit testing who is responsible when an AI maligns someone. In **Walters v. OpenAI**, a Georgia radio host sued OpenAI after its ChatGPT model falsely accused him of embezzling funds from a non-profit (Judge Denies Motion to Dismiss AI Defamation Suit | Alerts and Articles | Insights | Ballard Spahr) (Judge Denies Motion to Dismiss AI Defamation Suit | Alerts and Articles | Insights | Ballard Spahr). The incident arose when a third party asked ChatGPT to summarize a legal complaint; ChatGPT **fabricated a non-existent lawsuit** that described Walters as a defendant who had defrauded an organization, even though Walters had no connection to the real case. The defamatory summary, complete with details of supposed financial misconduct, was entirely AI-generated (Judge Denies Motion to Dismiss AI Defamation Suit | Alerts and Articles | Insights | Ballard Spahr) (Judge Denies Motion to Dismiss AI Defamation Suit | Alerts and Articles | Insights | Ballard Spahr). Walters sued OpenAI in June 2023, and in January 2024 the court **denied OpenAI's motion to dismiss**, allowing the case to move forward (Judge Denies Motion to Dismiss AI Defamation Suit | Alerts and Articles | Insights | Ballard Spahr) (Judge Denies Motion to Dismiss AI Defamation Suit | Alerts and Articles | Insights | Ballard Spahr). OpenAI had argued it shouldn't be liable because the user prompting ChatGPT supposedly knew the output was false, and because OpenAI's terms of use warn that ChatGPT may "hallucinate" inaccuracies (Judge Denies Motion to Dismiss AI Defamation Suit | Alerts and Articles | Insights | Ballard Spahr). The judge rejected those arguments at the pleading stage, signaling that an AI developer can potentially be treated as the **publisher of its AI's statements** for defamation purposes (at least where the user did not supply the false information). Notably, OpenAI may test a defense under Section 230 of the Communications Decency Act, which immunizes platforms from liability for *user-generated* content (Judge Denies Motion to Dismiss AI Defamation Suit | Alerts and Articles | Insights | Ballard Spahr). However, it is unsettled whether Section 230 applies when the "content" is produced by the platform's own algorithm rather than a human user (Judge Denies Motion to Dismiss AI Defamation Suit | Alerts and Articles | Insights | Ballard Spahr). The Walters case – likely the first of many involving AI output that harms reputation – could help define whether AI firms face publisher liability or enjoy immunity. Separately, other individuals have reported being defamed by AI-generated falsehoods (for example, a professor falsely named in an imaginary harassment case, a politician inaccurately described as convicted of bribery, etc.), although those incidents have not yet resulted in U.S. lawsuits (ChatGPT falsely accuses law prof of sexual harassment; is libel suit ...) (Can AI be sued for defamation? - Columbia Journalism Review). The risk of **misinformation litigation** is real: if a generative AI chatbot delivers false and damaging statements about a person or company, the injured party may pursue legal remedies. Companies integrating AI into publishing or communication tools should take note – they might be held accountable for defamatory outputs just as traditional publishers are, absent clear legal protections.

Data Privacy and Cybersecurity

Generative AI also presents novel **data privacy and breach** risks. One concern is that AI systems may ingest or expose personal information without authorization. In mid-2023, a class-action lawsuit (*Cousart v. OpenAI*) was filed in California accusing OpenAI and its partner Microsoft of scraping millions of individuals' personal data from the internet (including private information from social media and websites) to train ChatGPT (OpenAI, Microsoft defeat US consumer-privacy lawsuit for now | Reuters). The plaintiffs alleged violations of privacy rights and property rights in their data. In May 2024, U.S. District Judge Vince Chhabria dismissed that sweeping 204-page complaint, criticizing it as excessively verbose, filled with policy arguments, and lacking focus on specific legal harms (OpenAI, Microsoft defeat US consumer-privacy lawsuit for now | Reuters) (OpenAI, Microsoft defeat US consumer-privacy lawsuit for now | Reuters). The judge described the pleading as "containing swaths of unnecessary and distracting allegations" and noted that a court is not a "town hall meeting" for airing general grievances about AI (OpenAI, Microsoft defeat US consumer-privacy lawsuit for now | Reuters) (OpenAI, Microsoft defeat US consumer-privacy lawsuit for now | Reuters). However, the dismissal was **without prejudice**, giving plaintiffs an opportunity to re-file a trimmed complaint (OpenAI, Microsoft defeat US consumer-privacy lawsuit for now | Reuters). This outcome suggests that while courts are open to privacy claims against AI companies, they will require well-pleaded facts tying the technology to concrete injuries under existing laws (such as the California Consumer Privacy Act or intrusion upon seclusion). The *OpenAI* privacy case also demonstrates the difficulty of organizing a class around broad harms from AI data practices; as of this writing, no amended complaint or new class action has succeeded on such claims.

Apart from consumer lawsuits, companies using generative AI must be cautious about inadvertent **data leakage**. Employees may input sensitive personal or confidential data into AI tools, which could be stored or even used to further train the AI, potentially violating privacy obligations. For instance, if healthcare staff used ChatGPT with patient information, it might trigger HIPAA privacy violations. Likewise, proprietary data could be exposed – a well-publicized example occurred when an employee pasted confidential source code into an AI chatbot, only to realize it might be retained on external servers. Such scenarios blur the line between an internal data breach and an external cyber incident. A generative AI platform itself could also suffer a security failure: indeed, in March 2023, OpenAI disclosed a bug that briefly allowed some users to see excerpts of other users' chat history and payment info, a lapse that could be characterized as a data breach. While no lawsuit ensued from that incident, it highlighted that **AI services are not immune to typical cybersecurity issues**. Another emerging risk is **"poisoning" or corrupting AI models**. Attackers might manipulate the training data or prompt inputs to induce malicious outputs or to extract sensitive info from the model (a form of data exfiltration). If an AI deployed by a company is compromised and leaks personal data or allows unauthorized access to systems, the company could face liability for failing to secure it. The FTC has specifically warned that companies must consider whether AI tools "are prone to adversarial inputs or attacks that put personal data at risk" (AI and the Risk of Consumer Harm | Federal Trade Commission). In sum, generative AI can create new vectors for privacy breaches – either by the **improper use of personal data in training**, or by **introducing vulnerabilities** that hackers exploit. Companies should treat AI systems as part of their attack surface and governance scope, implementing safeguards to prevent and respond to data leakage. When incidents do occur, they are likely to be treated by courts and regulators under existing breach notification and privacy laws, even if the technology involved is cutting-edge.

Professional Malpractice and Errors

Generative AI is increasingly used to assist human professionals – from lawyers and doctors to software developers – which raises the question of who bears responsibility when the AI's **errors** cause harm. A cautionary tale widely cited in the legal community occurred in 2023, when a law firm filed a brief written with the help of ChatGPT that cited **fictitious case law** (Generative AI's Impact on Insurance Coverage: An Interview with ChatGPT-4 and Coverage Counsel on What Policyholders Should Be Doing Now — Policyholder Pulse — July 5, 2023). The AI had convincingly fabricated judicial decisions to support the attorney's argument. The mistake was only uncovered when opposing counsel and the judge could not find the cited cases, resulting in the embarrassed attorneys being sanctioned for violating their duty of candor (Generative AI's Impact on Insurance Coverage: An Interview with ChatGPT-4 and Coverage Counsel on What Policyholders Should Be Doing Now — Policyholder Pulse — July 5, 2023). While that incident (*Mata v. Avianca*) did not involve a client lawsuit, it underscores the risk of **professional negligence** if practitioners rely on AI outputs without verification. In a different field, software developers have leveraged AI coding assistants (like GitHub's Copilot or AWS's CodeWhisperer) to generate code. If the AI-suggested code contains bugs or security flaws, it could lead to product failures or breaches. Consider a scenario where an engineer uses generative AI to write a piece of software for a client, and a hidden error later causes a critical system outage or a data leak – the client might sue for malpractice or product liability. In fact, experts have noted that AI-written code may introduce vulnerabilities that **wrongdoers can exploit to hack a company's network** (Generative AI's Impact on Insurance Coverage: An Interview with ChatGPT-4 and Coverage Counsel on What Policyholders Should Be Doing Now —

Policyholder Pulse — July 5, 2023). Thus, **hallucinations and mistakes by AI** can translate into real-world damages: incorrect financial analysis, wrong medical advice, or faulty engineering designs, to name a few. Thus far, we have not seen a reported U.S. court decision squarely holding a professional or company liable for following flawed AI output. It is likely, however, that traditional standards of malpractice and negligence will apply. The human professional or the company using the AI remains responsible for exercising reasonable care. Using an AI tool won't excuse a doctor's misdiagnosis or an architect's design defect if they ought to have caught the error. Conversely, if a firm explicitly delegates tasks to an AI (for instance, an investment advisor letting AI allocate client assets), clients might argue the firm should be held vicariously liable for the AI's actions as if it were an employee or subcontractor. We are in uncharted waters, but companies should assume **"the buck stops with the human."** They should institute internal policies for AI use – such as requiring human review of AI-generated work products – both to reduce the risk of harm and to strengthen their defense that they met the standard of care (Generative AI's Impact on Insurance Coverage: An Interview with ChatGPT-4 and Coverage Counsel on What Policyholders Should Be Doing Now — Policyholder Pulse — July 5, 2023) (Generative AI's Impact on Insurance Coverage: An Interview with ChatGPT-4 and Coverage Counsel on What Policyholders Should Be Doing Now — Policyholder Pulse — July 5, 2023).

Bias, Discrimination, and Other Torts

Generative AI's outputs can sometimes reflect biased or offensive content, which in turn can create legal exposure. As noted, the EEOC has already taken action against an employer for **biased AI hiring practices**, and more broadly warned employers that using AI in employment decisions must comply with anti-discrimination laws (Tutoring firm settles US agency's first bias lawsuit involving AI software | Reuters) (Tutoring firm settles US agency's first bias lawsuit involving AI software | Reuters). While the iTutorGroup case involved a relatively straightforward misuse (explicitly programming age cut-offs), more subtle bias in AI-generated content could spark litigation in the future. For example, if a generative AI chatbot used by customers consistently gave poorer service or offensive responses to individuals of a certain race or gender, it could lead to claims of discrimination or harassment. In the employment context, imagine an AI HR assistant that generates biased performance evaluations or a resume screening tool that, unbeknownst to the employer, disproportionately filters out minority candidates due to biased training data. These scenarios could result in **hostile work environment claims or disparate impact lawsuits**. Indeed, concerns have been raised that generative AI could produce material that creates a **harassing or offensive workplace**, giving rise to claims by employees (Generative AI's Impact on Insurance Coverage: An Interview with ChatGPT-4 and Coverage Counsel on What Policyholders Should Be Doing Now — Policyholder Pulse — July 5, 2023). Companies deploying AI in interactions with the public also face potential **consumer protection and emotional distress claims** if the AI says or does something harmful – for instance, an AI financial advisor that steers users into unsuitable investments (leading to loss) could face negligence or fraud claims, or an AI companion bot that encourages self-harm might even trigger novel tort claims. We are only beginning to see such issues, but they highlight the need for rigorous testing and content moderation of AI systems. The legal system has long grappled with algorithmic bias in other contexts (credit scoring, housing ads, etc.), and generative AI will amplify those challenges by creating new content on the fly. Prudent companies should proactively address and audit for biases in AI outputs. On the flip side, **plaintiffs' lawyers and**

regulators are actively looking for egregious examples to test in court, so we can expect more litigation if high-profile incidents occur. As always, clear documentation of efforts to prevent AI-driven bias can be a key part of a legal defense.

Securities and "AI-Washing" Litigation

One of the newest liability fronts involves **shareholder lawsuits** against companies for misrepresenting or overstating their use of AI. In 2023, investors filed suits against firms like **Innodata Inc.** and **Telus International** after those companies made bold claims about incorporating AI into their business, only to have setbacks or disclosures that contradicted the AI hype (Protecting Your Business: AI Washing and D&O Insurance). Fast-forward to 2025, and this trend has accelerated. In March 2025, two more securities class actions were filed in California alleging that executives engaged in "AI-washing" – painting an overly rosy picture of their AI capabilities. In *Nunez v. Skyworks Solutions*, a semiconductor company was sued for allegedly overstating its "position and ability to capitalize on AI" in the smartphone market, which the complaint says led investors to buy stock at inflated prices (Protecting Your Business: AI Washing and D&O Insurance) (Protecting Your Business: AI Washing and D&O Insurance). The very next day, *Quiero v. AppLovin Corp.* was filed, accusing a mobile technology firm of **misleading investors** by touting its use of "cutting-edge AI" to drive its advertising business, when in reality the AI claims were exaggerated (Protecting Your Business: AI Washing and D&O Insurance) (Protecting Your Business: AI Washing and D&O Insurance). These cases underscore that public companies face not only technical and operational AI risks, but also **market disclosure risks**. If management overhypes AI initiatives or fails to disclose AI-related problems (like bias issues, regulatory inquiries, or lack of AI integration), they can be hit with shareholder suits for securities fraud or breach of fiduciary duty. From a legal standpoint, these claims will turn on the usual securities litigation questions – were any false statements made, were they material, and did executives act with scienter (intent or reckless disregard)? AI is simply the subject matter of the misstatements. However, what makes them noteworthy is how quickly AI has become a focus of investor expectations. Directors and officers should be aware that plaintiffs' attorneys (and the SEC) are listening to earnings calls and press releases for **buzzwords like "AI-driven"**, and they will not hesitate to sue if reality falls short of the talk. In-house counsel should counsel leadership to avoid speculative or conclusory assertions about AI and ensure any AI-related disclosures are accurate and not misleading. These suits also have **insurance implications** discussed below, as D&O policies will be the first line of defense for AI-related securities claims.

Implications for Insurance Coverage

The multifaceted risks of generative AI cut across several lines of insurance. Policyholders – and their insurers – must analyze how traditional coverage applies to AI-related incidents, and whether new endorsements or policies are needed to fill gaps. Below we examine how standard **cyber, professional (E&O), and D&O** policies may respond or exclude these risks, as well as considerations for general liability and other coverages. In all cases, the specific policy language and the facts of the claim will be critical, but emerging patterns can be observed.

(image) *Generative AI adoption is driving insurers to adapt coverage. Some insurers have begun offering endorsements to address AI-related perils such as data poisoning and IP infringement, recognizing that traditional cyber policies often excluded these novel risks (Insurers Explore New AI Coverage Options, Potentially Filling Coverage Gaps for Policyholders Developing Generative AI | The Policyholder Perspective) (Insurers Explore New AI Coverage Options, Potentially Filling Coverage Gaps for Policyholders Developing Generative AI | The Policyholder Perspective).*

Cyber and Media Liability Coverage

Stand-alone cyber insurance policies have become a common risk management tool for data breaches and network security incidents. Many cyber policies also include **media liability** coverage for harms like defamation, copyright/trademark infringement, and privacy violations (especially when arising online). These coverages are directly implicated by generative AI. For instance, if an AI system deployed by an insured causes a data breach or privacy loss – say an employee's use of an AI chatbot leads to exposure of personal data – that could trigger the cyber policy's privacy breach insuring agreement. Likewise, if a company is sued for content liability (e.g. defamation or IP infringement) based on AI-generated material it published, a cyber policy's media liability section might respond, or a traditional media liability policy could. (Generative AI's Impact on Insurance Coverage: An Interview with ChatGPT-4 and Coverage Counsel on What Policyholders Should Be Doing Now — Policyholder Pulse — July 5, 2023) (Generative AI's Impact on Insurance Coverage: An Interview with ChatGPT-4 and Coverage Counsel on What Policyholders Should Be Doing Now — Policyholder Pulse — July 5, 2023) It's important to check the scope of such coverage. Many **Commercial General Liability (CGL)** policies, for example, cover "personal and advertising injury" which can include offenses like defamation or copyright infringement *in advertising* (Generative AI's Impact on Insurance Coverage: An Interview with ChatGPT-4 and Coverage Counsel on What Policyholders Should Be Doing Now — Policyholder Pulse — July 5, 2023). However, CGL coverage may be limited – it typically would not cover an IP infringement claim outside of the advertising context, and many modern CGL policies have IP exclusions or require the offense to relate to the insured's advertisement of goods and services. Generative AI claims might not fit neatly; for example, if an AI is used internally and inadvertently generates infringing content that is not part of an advertisement, CGL might not apply. Cyber policies can fill this gap, as they often cover a broader range of media liability arising from online content or technology activities (Generative AI's Impact on Insurance Coverage: An Interview with ChatGPT-4 and Coverage Counsel on What Policyholders Should Be Doing Now — Policyholder Pulse — July 5, 2023).

That said, **traditional cyber policies have not universally kept pace with generative AI risks**. In fact, insurers initially were cautious – many cyber policies **excluded losses related to the development of AI models**, given the unquantified, potentially catastrophic nature of those exposures (Insurers Explore New AI Coverage Options, Potentially Filling Coverage Gaps for Policyholders Developing Generative AI | The Policyholder Perspective). For instance, if a tech company was building its own generative AI and got sued by a thousand copyright owners, a standard cyber policy might invoke exclusions for IP liability or for liability arising from providing a software product. This is analogous to how some cyber policies exclude product liability or professional services, pushing those into other lines like tech E&O. As the demand for coverage grows, though, insurers are starting to respond. In early 2024, cyber insurer **Coalition** introduced a policy endorsement specifically to cover certain **security breaches stemming from the use of generative AI** (Insurers Explore New AI Coverage Options, Potentially Filling Coverage Gaps for Policyholders Developing Generative AI | The Policyholder Perspective). And in October 2024, AXA XL launched one of the first tailored AI insurance solutions: an endorsement to its cyber policy that **expands coverage to address risks like data poisoning attacks, IP infringement in AI outputs, and even regulatory fines under laws such as the EU's AI Act** (Insurers Explore New AI Coverage Options, Potentially Filling Coverage Gaps for Policyholders Developing Generative AI | The Policyholder Perspective). This indicates a recognition that policyholders incorporating AI need protection beyond the standard cyber wording. We can expect other major carriers to follow with AI-focused enhancements, whether as add-ons or in next-generation cyber forms. In the meantime, in-house counsel should **review current cyber policies for potential gaps**: Are claims arising from AI-generated content covered or excluded? Is there coverage for unintentional copyright/trademark infringement by digital content? Are regulatory investigations (for example, an FTC inquiry into AI use) covered under network security or privacy liability sections? Also, consider **sub-limits** – some cyber policies might sub-limit certain coverages like regulatory fines or media liability. It may be prudent to negotiate higher limits or remove exclusions if AI-related exposure is significant for the insured's operations. Additionally, companies relying on third-party AI vendors should pay attention to contract terms and any **indemnities** (or lack thereof) from those providers, as that can affect how insurance would respond in a claim scenario.

Errors & Omissions (Professional Liability) Coverage

Professional liability or Errors & Omissions insurance covers financial losses to third parties caused by the insured's negligence or errors in the performance of professional services. For many organizations, if they incorporate generative AI into their services or advice, any AI-caused mistake could lead to an E&O claim. A key question is whether work output from an AI tool is considered part of the insured's "professional services" under the policy (Generative AI's Impact on Insurance Coverage: An Interview with ChatGPT-4 and Coverage Counsel on What Policyholders Should Be Doing Now — Policyholder Pulse — July 5, 2023) (Generative AI's Impact on Insurance Coverage: An Interview with ChatGPT-4 and Coverage Counsel on What Policyholders Should Be Doing Now — Policyholder Pulse — July 5, 2023). Most likely it is – for example, if a consulting firm uses AI to draft a report for a client, it's still delivering a consulting service. But insurers and policyholders should clarify this. Companies should **confirm that work product generated with AI is not excluded from E&O coverage** and indeed falls within the covered services definition (Generative AI's Impact on Insurance Coverage: An Interview with ChatGPT-4 and Coverage Counsel on What Policyholders Should Be Doing Now — Policyholder Pulse — July 5, 2023) (Generative AI's Impact on Insurance Coverage: An Interview with ChatGPT-4 and Coverage Counsel on What Policyholders Should Be Doing Now — Policyholder Pulse — July 5, 2023). If an AI-related failure occurs (say a software developer delivers faulty code written by an AI, or a marketing agency produces an ad with AI that inadvertently libels someone), the E&O policy should respond as it would to any error in the insured's work.

However, there could be grey areas. Some E&O policies have exclusions for certain types of acts – for instance, a financial advisor's E&O might exclude investment losses due to misrepresentation. If an AI chatbot the advisor deployed made an unauthorized guarantee about returns, an insurer might invoke such an exclusion. Another consideration is the **use of third-party AI platforms**: if an insured relies on an AI vendor and that vendor's tool fails, insurers might attempt to deny coverage by arguing the claim arose from the failure of a third-party product (this is uncommon, but something to watch in policy wording). Thus far we haven't seen new, AI-specific exclusions widely added to E&O policies, but **insurers are starting to ask questions in applications about AI usage** (Generative AI's Impact on Insurance Coverage: An Interview with ChatGPT-4 and Coverage Counsel on What Policyholders Should Be Doing Now — Policyholder Pulse — July 5, 2023). In-house counsel should be prepared to describe how their company controls AI-related risks, as underwriting scrutiny increases. On the flip side, insureds may request endorsements to **affirmatively cover AI** – some brokers report negotiations to explicitly include AI-driven services in the definition of professional services (Generative AI's Impact on Insurance Coverage: An Interview with ChatGPT-4 and Coverage Counsel on What Policyholders Should Be Doing Now — Policyholder Pulse — July 5, 2023). Given the "human in the loop" best practices, an insured can argue AI is just another tool, and any mistake is ultimately a professional mistake covered by E&O. Still, prudence dictates checking for any exclusions that could be interpreted to bar coverage (for example, some tech E&O policies exclude claims of intellectual property infringement – if an AI causes an IP claim against the insured, that could fall between the cracks if not addressed via media liability coverage as noted above). In summary, **E&O policies are a vital backstop for AI-related errors**, and companies using AI in delivering products or services should ensure their policies are up to date. This may involve working with brokers to adjust language at renewal, especially as **insurers could introduce new AI exclusions or sublimits** in response to growing losses. Staying ahead of that by negotiation is critical.

Directors & Officers (D&O) and Management Liability

Directors and Officers liability policies protect a company's executives and the entity against claims of wrongful acts in managing the company – including securities class actions, derivative suits, and regulatory investigations. As discussed, **"AI-washing" lawsuits** and other investor actions are emerging, which means D&O insurance will be a crucial line of defense (Protecting Your Business: AI Washing and D&O Insurance). Fortunately, a typical public company D&O policy is broad in scope, covering securities claims and breach of fiduciary duty allegations, subject to exclusions for fraud or personal profit (which usually apply only if there's a final adjudication of dishonest conduct). An AI-related securities lawsuit, such as those against Skyworks or AppLovin, should fall squarely within D&O coverage for "securities claims" – there is nothing fundamentally different about the nature of the claim; it's the subject matter (AI) that's new. **However, companies should review their D&O programs for any exclusions that could potentially limit coverage for AI-related matters.** For instance, some D&O policies for certain industries have exclusions for claims arising from professional services or technology errors (to avoid overlap with E&O). If a D&O policy had a broad "technology services" exclusion, one could imagine an insurer attempting to invoke it in a scenario where a lawsuit alleges the company misled customers about its AI product's functionality. Insureds will want to ensure no such exclusion bars coverage for the types of **AI misrepresentation claims** we've seen. Hunton Andrews Kurth, an insurance law firm, noted key considerations for maximizing D&O protection against AI risks, including: (1) **Policy review** to confirm AI-related losses aren't swept into any coverage exclusions (like a cyber exclusion in a D&O policy), (2) ensuring coverage for **regulatory investigations** (e.g. an SEC inquiry or state AG investigation into AI use), (3) **coordinating D&O with cyber/E&O** to avoid gaps or disputes over which policy applies, (4) exploring **AI-specific endorsements or policies** if available, and (5) maintaining robust Side A coverage for individual directors/officers as an extra safety net (Protecting Your Business: AI Washing and D&O Insurance). In practical terms, this means when renewing D&O coverage, companies should discuss AI exposures with their carrier – not to invite a restrictive endorsement, but to confirm that the insurer does not view AI issues as outside the intended coverage.

One area to watch is regulatory coverage under D&O. If the FTC or DOJ were to investigate a company for allegedly unfair or deceptive AI practices (say a probe into whether an AI violated consumer protection laws), the company might incur significant legal costs. D&O policies often cover "investigations" of insured persons and sometimes the entity, but the trigger language can be tricky (coverage might attach when there's a formal investigative order or a Wells notice, etc.). Companies heavily invested in AI may want to bolster entity investigation coverage by endorsement, to ensure early engagement with regulators is covered. Additionally, **Side C (entity) coverage** in public D&O will cover securities claims but not other claims against the entity – so a pure consumer class action (not securities) over AI might not be covered by D&O except as a derivative claim. For example, if consumers sued a company for fraud because an AI-powered product didn't work as advertised, that likely hits the CGL or cyber policy, not D&O (unless shareholders bring a parallel claim). Thus, **D&O is not a catch-all for all AI litigation**, but rather focused on governance-related claims. In the context of AI, that primarily means investor and shareholder claims, and possibly regulatory oversight. Given how hot AI is, it's conceivable that **shareholders could also sue for breach of fiduciary duty** if a board wholly fails to oversee AI risks (an oversight claim, akin to Caremark claims in derivative suits). D&O would respond to defend such claims. In short, D&O policies are as vital as ever in the AI era, and companies should **treat AI risks as part of their D&O coverage review**, just as they would emerging risks like cybersecurity or COVID-19 impacts in prior years.

Other Relevant Coverages (General Liability, EPLI, etc.)

Beyond the major lines above, a few other policies may be triggered by generative AI issues. **General Liability (GL)**, as noted, covers certain personal/advertising injuries – this could come into play if, for example, a company is sued for libel because of something an employee posted that was drafted by an AI. If the post was in the course of advertising the company's goods, the GL insurer might defend under the advertising injury coverage (Generative AI's Impact on Insurance Coverage: An Interview with ChatGPT-4 and Coverage Counsel on What Policyholders Should Be Doing Now — Policyholder Pulse — July 5, 2023). However, many companies rely on cyber or media policies for a broader shield on such risks, as GL can be limited and some GL carriers now exclude broad intellectual property claims. **Employment Practices Liability (EPLI)** is another line to consider in the AI context. EPLI covers claims by employees (or sometimes third parties) alleging discrimination, harassment, or wrongful employment decisions. If an employee alleges that an AI tool used by the employer created a hostile work environment (for instance, a HR chatbot that responded with sexist or racist remarks, or an AI system that systematically gave lower performance scores to a protected group), an EPLI policy could potentially cover the claim (Generative AI's Impact on Insurance Coverage: An Interview with ChatGPT-4 and Coverage Counsel on What Policyholders Should Be Doing Now — Policyholder Pulse — July 5, 2023). The first AI bias case by the EEOC (iTutorGroup) was handled as a government enforcement matter, and many EPLI policies do cover defense costs (and sometimes settlements) for EEOC actions or similar proceedings. Companies deploying AI in employment should review their EPLI coverage and, if necessary, seek endorsements to clarify that **automated decision-making falls within the policy's scope** (Generative AI's Impact on Insurance Coverage: An Interview with ChatGPT-4 and Coverage Counsel on What Policyholders Should Be Doing Now — Policyholder Pulse — July 5, 2023). Some insurers might introduce exclusions for decisions made by algorithms (to push that risk to a tech E&O policy), so this warrants attention.

Finally, for companies that produce or sell AI systems themselves (as opposed to just using them internally), **product liability insurance** could conceivably be implicated if an AI product causes physical injury or property damage. For example, consider an AI-powered tool that malfunctions and causes damage to equipment or a person – that could trigger a products claim under a general liability policy's products-completed operations coverage. However, most current generative AI applications are software-based and cause economic losses more so than physical harm. The autonomous vehicle realm is an exception (AI driving systems can cause accidents), but that implicates auto liability and is beyond our scope here. Still, it is worth noting that as AI is embedded in more physical devices (drones, robots, medical devices), the **line between cyber and physical damage liability blurs**, and insurers will likely adjust policy language to delineate what is covered by a tech/cyber policy versus a GL/product policy.

Conclusion

The rapid rise of generative AI has opened up exciting opportunities but also a Pandora's box of liability issues. U.S. courts and regulators are actively addressing these challenges: we have early case law on how copyright doctrines apply to AI, the first defamation and discrimination suits sparked by AI outputs, and an uptick in shareholder litigation over AI hype. For insurance carriers and insureds alike, the key takeaway is that **traditional insurance will be tested against novel fact patterns**, and in many instances it can respond – but careful scrutiny of coverage is required. In-house counsel at insurers should be tracking this evolving case law to anticipate how claims might be handled. Likewise, they may consider working with underwriters to update policy forms and endorsements to either clarify coverage or exclude unmanageable risks. We have begun to see the market react, with endorsements covering things like AI training-data IP liabilities and AI-specific exclusions being contemplated. Policyholders, for their part, should **proactively review their insurance portfolios for AI-related gaps** (Generative AI's Impact on Insurance Coverage: An Interview with ChatGPT-4 and Coverage Counsel on What Policyholders Should Be Doing Now — Policyholder Pulse — July 5, 2023) (Insurers Explore New AI Coverage Options, Potentially Filling Coverage Gaps for Policyholders Developing Generative AI | The Policyholder Perspective). This includes evaluating cyber, E&O, D&O, CGL, and other policies in tandem, since a complex AI incident could potentially trigger multiple lines (or fall between them if not coordinated). For example, a single AI fiasco might lead to a privacy regulatory investigation (cyber/D&O), a consumer class action (cyber/E&O), and a shareholder suit (D&O). Ensuring that at least one policy will cover each dimension – and that insurers cannot easily point fingers at each other in denial – is crucial.

Ultimately, managing generative AI risk is a multidisciplinary effort: robust **governance and oversight** of AI use (to prevent harm in the first place), diligent **compliance with emerging laws and regulations**, and thoughtful **insurance risk transfer**. In-house counsel at insurance carriers should be prepared to advise both their underwriting teams and their insureds on these issues. By staying abreast of legal developments and tailoring coverage accordingly, the insurance industry can rise to meet the challenges of the AI era – providing the certainty and protection needed as businesses navigate uncharted territory. The landscape will continue to evolve rapidly, but with informed vigilance and adaptive strategies, insurers and policyholders can mitigate the cyber liability risks arising from generative AI while harnessing its potential benefits.

Sources: Recent case dockets and filings; legal news outlets (Reuters, Law360, Bloomberg Law); insurance industry publications and law firm insights on AI (Pillsbury (Generative AI's Impact on Insurance Coverage: An Interview with ChatGPT-4 and Coverage Counsel on What Policyholders Should Be Doing Now — Policyholder Pulse — July 5, 2023) (Generative AI's Impact on Insurance Coverage: An Interview with ChatGPT-4 and Coverage Counsel on What Policyholders Should Be Doing Now — Policyholder Pulse — July 5, 2023), Reed Smith (Insurers Explore New AI Coverage Options, Potentially Filling Coverage Gaps for Policyholders Developing Generative AI | The Policyholder Perspective) (Insurers Explore New AI Coverage Options, Potentially Filling Coverage Gaps for Policyholders Developing Generative AI | The Policyholder Perspective), Hunton AK (Protecting Your Business: AI Washing and D&O Insurance) (Protecting Your Business: AI Washing and D&O Insurance), Skadden (Motion To Dismiss Ruling Provides Further Insight Into How Courts View AI Training Data Cases | Insights | Skadden, Arps, Slate, Meagher & Flom LLP) (New FTC Initiative Targets Deceptive AI Claims and a Generative AI Service | Insights | Skadden, Arps, Slate, Meagher & Flom LLP)); and court decisions and orders in *Walters v. OpenAI* (Judge Denies Motion to Dismiss AI Defamation Suit | Alerts and Articles | Insights | Ballard Spahr) (Judge Denies Motion to Dismiss AI Defamation Suit | Alerts and Articles | Insights | Ballard Spahr), *Andersen v. Stability AI* (Takeaways from the Andersen v. Stability AI Copyright Case | Copyright Alliance) (Takeaways from the Andersen v. Stability AI Copyright Case | Copyright Alliance), *Doe v. GitHub* (Motion To Dismiss Ruling Provides Further Insight Into How Courts View AI Training Data Cases | Insights | Skadden, Arps, Slate, Meagher & Flom LLP), *Cousart v. OpenAI* (OpenAI, Microsoft defeat US consumer-privacy lawsuit for now | Reuters), *EEOC v. iTutorGroup* (Tutoring firm settles US agency's first bias lawsuit involving AI software | Reuters) (Tutoring firm settles US agency's first bias lawsuit involving AI software | Reuters), among others. This report reflects developments through early 2025 and will require updates as new laws and precedents emerge in this fast-moving area.