AI PRODUCT LIABILITY: THE NEXT WAVE OF LITIGATION

Date: 27 March 2026

US Litigation and Dispute Resolution Alert

By: Jin J. To

Artificial intelligence (AI) litigation is beginning to consolidate around a familiar body of doctrine: product liability. Early cases are testing whether consumer-facing AI applications are treated as products (not services) and whether alleged harms are framed as design defects, inadequate warnings, or foreseeable misuse. That shift is also being reinforced by lawmakers—most notably the European Union's revised directive on liability for defective products (the PLD)[1] and a growing set of US state enactments (including California). Together, these developments suggest product liability will be a primary lens for the next wave of AI litigation.

This is a notable turn because many of the first headline AI disputes were framed through various adjacent doctrines—consumer protection, privacy, defamation, and intellectual property. Product liability is different: It is built to evaluate mass-distributed technologies through the lenses of defect, warnings, and foreseeability, with liability that can extend across a chain of entities involved in making a product available. As AI functionality becomes embedded into everyday consumer and enterprise workflows, plaintiffs have stronger incentives to describe the AI-enabled experience as a product and to litigate it the way courts already litigate other complex technologies.

CASE SNAPSHOTS: HOW PLAINTIFFS ARE PLEADING AI PRODUCT CLAIMS

A recurring threshold issue in these disputes is how courts should characterize generative AI outputs. Defendants often argue that chatbot responses are expressive content, seeking to reframe claims as attempts to impose liability for speech rather than for product design. Plaintiffs, by contrast, increasingly draft complaints to target the architecture of the deployed system—guardrails, defaults, escalation pathways, and marketing—so the case looks like a product-defect dispute instead of a content dispute.

Garcia v. Character Technologies, Inc.[2] is an early bellwether for how plaintiffs are attempting to fit chatbot-related injuries into a traditional products framework. The plaintiffs alleged a 14-year-old user formed an intense emotional relationship with a Character.AI chatbot and died by suicide. The complaint ties the alleged harm to product design and interaction patterns, and the court treated the mass-marketed chatbot app as a “product” under a strict-liability pleading. The ruling also permitted theories aimed at an upstream technology provider and the manufacturer to proceed at the pleading stage, reflecting how product-liability concepts can extend beyond the branded application to alleged component or enabling actors.

Raine v. OpenAI[3] and coordinated OpenAI matters (California) illustrate how plaintiffs are reframing “bad outputs” as allegations about AI architecture. In 2025, the parents of 16-year-old Adam Raine filed suit alleging that ChatGPT fostered emotional dependency, contributed to self-harm by providing instructions on hanging, and that

©2005-2026 K&L Gates LLP. All Rights Reserved.

the product lacked adequate safeguards. The pleadings emphasize guardrails, crisis-intervention behavior, and whether monitoring signals should have triggered different product behavior. The coordination of multiple actions also signals a familiar products-litigation dynamic: Plaintiffs may seek to develop pattern-of-conduct evidence around design choices, testing timelines, and warning strategies.

Nippon Life v. Open AI[4] underscores that AI product-liability-adjacent theories may not be limited to end-user personal injury. An insurer sued OpenAI in federal court in Illinois seeking to recover costs from an AI-assisted, meritless legal filings (including at least one citation to a nonexistent case). The case highlights institutional economic-harm theories and the potential for third-party plaintiffs, and it illustrates how terms and disclosures may be litigated as notice and risk-recognition timelines, not solely as defenses.

Nevada v. MediaLab AI, Inc.[5] demonstrates how some US states are trailblazers for the new wave of AI product-liability litigations. In 2025, the Nevada attorney general filed a lawsuit against a tech holding company and its social messaging app for alleged harms caused to Nevada's youth. The Nevada attorney general claims that the app is defective and "unreasonably dangerous" to the youth because there are no safety features to protect minors from "being contacted by predators." The lawsuit illustrates Nevada's aggressive direction to hold AI platforms responsible for user harm while shaping public policy.

Across these matters, a common strategy is emerging: Treat the "AI system" not as an abstract model, but as the deployed product experience—its interface, defaults, guardrails, and marketing claims. That framing is designed to sidestep threshold fights over whether a particular output is protected expression and instead litigate the system's design choices as the alleged defect. It also tees up claims against multiple entities involved in deployment, from branded application providers to alleged component or enabling actors.

WHY PRODUCT LIABILITY FITS AI DEPLOYMENTS

Product-liability doctrine is designed for technologies that reach users at scale through repeatable experiences—precisely how many AI applications are now distributed. As courts decide whether specific AI applications are "products" or "services," plaintiffs are increasingly pleading traditional product liability theories: design defect (guardrails, interaction design, and lack of safety features), failure to warn (limitations and foreseeable misuse), and negligence (reasonable testing/monitoring in context).

A second recurring theme is supply-chain liability. Pleadings and early rulings suggest plaintiffs will test theories that reach beyond the model developer to the enterprise that brands and deploys the system, as well as upstream providers that allegedly enabled or substantially participated in the final product's integration. In parallel, legislation like California's AB 316 (addressing "autonomy" defenses) reflects a policy trend toward keeping causation disputes fact-bound rather than allowing "AI did it" to operate as a categorical shield.

REGULATION IS STEERING TOWARD PRODUCT-LIABILITY CONCEPTS

Policy developments are increasingly using the language of products doctrine—what a product is, who is in the distribution chain, and how responsibility is allocated when software causes harm. For AI, that matters because these frameworks influence pleading strategies and can supply persuasive authority for defect, foreseeability, and standard-of-care arguments even where claims remain common-law tort.

©2005-2026 K&L Gates LLP. All Rights Reserved.

Several developments illustrate the shift. The PLD treats software—including AI systems—as "products," extends strict-liability concepts across the distribution chain, and captures parties that substantially modify AI systems; member states must transpose the directive by December 2026. In the United States, the AI LEAD Act (a US Senate proposal sponsored by Illinois Sen. Richard Durbin) reflects a similar policy interest in product-liability framing for certain AI systems, regardless of whether the proposal advances in its current form. At the state level, targeted enactments such as California's AB 316 (addressing "autonomy" defenses) and SB 243 (companion chatbots) may be cited by plaintiffs to argue foreseeability and to frame what safety features are reasonable in particular deployment contexts.

For multinational products, the EU framework can influence more than European litigation. The PLD's concepts—software as a product, coverage of substantial modifications, and supply-chain responsibility—are likely to appear in US complaints and expert reports as persuasive reference points, particularly where companies market a single AI-enabled product across jurisdictions. Likewise, detailed state statutes can function as "standard-setting" signals in tort cases: Even when they do not apply directly, plaintiffs may argue they reflect what risks were foreseeable and what safeguards were reasonable for a given category of AI deployment.

WHAT THIS MEANS FOR THE NEXT WAVE

Looking ahead, several themes are likely to define how this next wave develops. Courts will continue testing the product-versus-service line, a characterization that can determine whether strict-liability theories are available and how warnings and design are evaluated. Pleadings are also increasingly litigating AI "architecture"—guardrails, escalation design, and user experience choices that invite reliance—rather than focusing on isolated outputs. At the same time, liability theories are moving up and down the AI supply chain as plaintiffs explore component-part and substantial-participation theories that can reach upstream and downstream actors. Finally, regulation is becoming a shared liability vocabulary: The PLD and targeted state statutes are likely to appear in complaints and expert reports as reference points for defect and foreseeability, while testing artifacts, monitoring signals, and change histories remain central in discovery and can shape both causation narratives and settlement leverage.

For companies looking to reduce exposure in this environment, two practical disciplines consistently matter in product cases: defining the product and substantiating the design story. Mapping the deployed system—model and version, prompts, tool connections, retrieval sources, and safety settings—helps avoid ambiguity about what the product was at a given point in time, particularly where behavior changes with updates. In parallel, contemporaneous documentation of testing, risk identification, and safety tradeoffs often become the evidentiary backbone of defect and foreseeability arguments; it is the record that allows a defendant to explain not just what was built, but why the design choices were reasonable when made.

CONCLUSION

The early AI cases now in litigation—alongside developments like the PLD and California's targeted statutes—signal a broader trend: Established product-liability doctrine is migrating into AI contexts. Over the next several years, courts will supply threshold answers on product-versus-service characterization, the viability of design-defect framing for AI architecture, and how autonomy and causation arguments are handled. That combination of litigation and legislation makes product liability a likely focal point for the next wave of AI disputes.

©2005-2026 K&L Gates LLP. All Rights Reserved.

FOOTNOTES

[1] Directive (EU) 2024/2853.

[2] Case No. 6:2024-cv-01903, (M.D. Fla., Oct. 23, 2024).

[3] Case No. GCG-25-628528 (Cal. Sup. Ct., Aug. 26, 2025).

[4] Case No. 1:26-cv-02448 (N.D. Ill., Mar. 4, 2026).

[5] Case No. A-25-92-5766-C (Nev. Dist. Ct. Clark County, 2025).

KEY CONTACTS



JIN J. TO
ASSOCIATE

CHICAGO
+1.312.807.4224
JIN.TO@KLGATES.COM

This publication/newsletter is for informational purposes and does not contain or convey legal advice. The information herein should not be used or relied upon in regard to any particular facts or circumstances without first consulting a lawyer. Any views expressed herein are those of the author(s) and not necessarily those of the law firm's clients.

©2005-2026 K&L Gates LLP. All Rights Reserved.