Katten

ADVISORY



Every Prompt Leaves a Record: How Businesses Can Manage Litigation Risk From Everyday AI Use

October 30, 2026

Use of artificial intelligence (AI) may have significant implications in litigation. AI prompts, outputs and conversation histories may be requested in discovery, and companies may need to preserve relevant AI records once litigation is reasonably anticipated.

Generative AI is now part of the workday at many companies. Employees may use it to draft emails, summarize contracts, develop strategy and research questions they once would have taken to a colleague or an attorney. These conversations are often saved and could later be reviewed by an opposing party, a regulator or a jury. Recent litigation underscores the importance of assessing how employees use AI and establishing safeguards before a dispute arises.

Key Takeaways

- AI chats can be requested and used as evidence in litigation and government investigations, much like email.
- AI chat histories may remain accessible. Companies should assess platform retention settings and preserve relevant records when litigation is reasonably anticipated.
- Information entered into public or consumer AI tools may be treated as shared with an outside party, which can undermine confidentiality and privilege.
- Employees should not assume that AI conversations about legal issues are privileged, particularly when conducted without direction from counsel.
- A clear AI use policy, approved enterprise tools, updated retention practices and employee training can help reduce these risks.

Can Your Employees' AI Chats Be Used in Litigation?

AI prompts, outputs and conversation histories can constitute electronically stored information (ESI) under the Federal Rules of Civil Procedure, which is the same category that covers email, text

messages and other digital business records. Relevant AI chats may therefore be subject to discovery.

Storage on an AI provider's servers does not, by itself, place those records beyond discovery. Discovery can reach information within a party's possession, custody or control, including records held by outside providers. An employee's ability to view or export a chat history may support a finding that the company controls it. Even if those records are not deemed to be within a party's possession, custody or control, they may still be obtained in the course of litigation from a third-party through a third-party subpoena.

Courts have also been willing to order preservation of AI chats on a very large scale. In *In re OpenAI, Inc. Copyright Infringement Litigation* (S.D.N.Y. 2025), the court ordered OpenAI to preserve output logs that would otherwise have been deleted, including at users' request. The order was upheld over objections that it would require keeping billions of conversations. These decisions suggest courts are unlikely to treat AI chats differently from other business communications.

AI records remain subject to relevance requirements. In another In re OpenAI decision (S.D.N.Y. 2025), the court declined to sanction the loss of AI chat logs because they were not relevant. In *Tremblay v. OpenAI, Inc.* (N.D. Cal. 2024), however, the court found a party's AI test results relevant to the opposing party's defenses even though they did not support the producing party's own claims. For businesses, the practical implication is that AI conversations about customers, transactions, product issues, workplace complaints or competitors that become relevant to litigation may become evidence, including when the results do not support the company's position. Employees should apply the same care to AI prompts that they apply to other business communications.

Do Companies Have to Preserve AI Chats?

Companies should assess relevant AI records as part of their preservation response when litigation or an investigation is reasonably anticipated, rather than assume those conversations are transient. In 2024, the Department of Justice (DOJ) Antitrust Division and the Federal Trade Commission (FTC) updated their standard preservation demands to cover collaboration tools and ephemeral messaging apps, and the DOJ warned that failing to produce such materials could result in obstruction-of-justice charges. Although that guidance does not identify AI tools by name, its reasoning may inform expectations for preserving AI records.

Courts have historically excused truly transient data, such as browser caches, from discovery because it is not kept in the ordinary course of business. *See Travelers Prop. Cas. Co. of Am. v. Cannon & Dunphy, S.C.* (E.D. Wis. 2014). Many enterprise and consumer AI platforms, by contrast, may save chat histories. Companies should understand whether those histories can be retrieved and

assess their preservation and production obligations when the records are relevant to reasonably anticipated litigation.

Litigation hold procedures should therefore address AI chats where relevant to the matter, including whether automatic deletion must be suspended. Failure to preserve AI records subject to a preservation obligation could expose a company to sanctions.

Is What You Share With AI Confidential?

Companies should assess an AI tool's confidentiality terms before employees enter sensitive business or legal information. In *United States v. Heppner* (S.D.N.Y. 2026), the court held that a criminal defendant's AI chats were not protected by attorney-client privilege or work product protection, even though he was asking about information his attorney had given him. The court pointed to three factors: the AI chats were not with his attorney, his attorney did not direct them, and they were not confidential because the AI platform reserved the right to share user data with third parties, including government regulators.

Separately, in *Slam Corp. v. Lynk Global, Inc.* (Del. Ch. 2025), the Delaware Court of Chancery entered a stipulated protective order stating that providing confidential discovery material to an open generative AI tool constitutes disclosure to a third party. Although the order was not a ruling on the issue, it illustrates how restrictions on third-party disclosure can apply to public AI tools.

The implications depend on the circumstances and the tool's terms. Heppner relied heavily on the AI provider's right to disclose user data, so confidential, internal AI systems may present different considerations. Companies should assess their own systems and contractual protections rather than assume that all AI tools present the same confidentiality risks.

AI use policies should clearly restrict the entry of trade secrets, legal advice and confidential customer or transaction information into public or consumer AI tools. Such disclosures can put confidentiality protections and attorney-client privilege at risk.

When Might Privilege or Work Product Protection Apply?

Companies should involve counsel before employees use AI to analyze a legal matter. Whether attorney-client privilege or work product protection applies depends on the circumstances, including who uses the tool, for what purpose and under what confidentiality protections. A conversation does not become privileged simply because it concerns a legal issue.

Attorney involvement and the purpose of the communication can be important for work-product protection. In *Tremblay* and in *Concord Music Group, Inc. v. Anthropic PBC* (N.D. Cal. 2025), the courts protected prompts written by attorneys because they reflected the attorneys' strategy and

thinking. A prompt seeking to develop an argument or test a legal theory may present a stronger basis for protection than a routine factual inquiry. Ordinary business use by employees outside the context of anticipated litigation generally presents a weaker basis for claiming work-product protection.

Attorney-directed AI use does not eliminate waiver risk. In both *Tremblay* and *Concord*, the courts found that the parties waived protection by relying on AI outputs in their court filings. Companies considering whether to rely on AI-generated results in litigation should assess with counsel whether doing so could expose the underlying prompts to discovery.

AI use connected to a legal matter should therefore be directed by counsel using approved tools, with confidentiality and potential waiver considered before the results are used.

What Businesses Should Do Now

Managing these risks does not require companies to stop using AI. It requires visibility into employee use, implementation of appropriate controls over sensitive information, and the ability to preserve relevant records, including AI prompts, outputs and conversation histories, that are relevant to reasonably anticipated litigation. Companies should consider incorporating AI into the governance framework that applies to email and other business communications. Steps to consider include:

1. **Take inventory of AI use.** Identify which AI tools employees are using, including company-approved enterprise tools and consumer tools accessed on their own, and what they're using them for.
2. **Adopt or update an AI use policy.** Specify approved tools, prohibited uses and restrictions on entering sensitive information, such as trade secrets, customer data and legal advice, into public or consumer AI tools.
3. **Favor enterprise tools with strong confidentiality terms.** Review vendor terms on data retention, confidentiality and disclosure to third parties, and direct business use toward tools with appropriate protections for company data.
4. **Align retention policies with AI.** Understand how long each AI tool retains chat histories, align retention settings with the company's records policies and confirm that relevant AI chats can be preserved and collected when needed.
5. **Address relevant AI records in litigation holds.** Assess whether hold notices should expressly cover AI prompts, outputs and account settings for enterprise and consumer tools, and confirm that the information technology (IT) team can suspend automatic deletion where required.

6. **Route legal questions through counsel.** Remind employees that independently using AI to analyze legal issues does not, by itself, establish privilege, and that AI use connected to a dispute should be directed by in-house or outside counsel.
7. **Train employees.** Explain that work-related AI conversations may constitute business records and may be reviewed in litigation or an investigation.
8. **Build AI into broader governance.** Assign ownership for AI oversight across legal, compliance, IT and business teams, and revisit your AI use policy as courts and regulators continue to address these issues.

Companies that address these issues now will be better positioned to protect confidential information and to respond efficiently if litigation or an investigation arises.

CONTACTS

For more information, please contact your Katten attorney or any of the following Securities Litigation attorneys.



Bruce G. Vanyo
+1.212.940.8787
bruce@katten.com



Christina L. Costley
+1.310.788.4485
christina.costley@katten.com



Jamie N. Noonan
+1.312.902.5492
jamie.noonan@katten.com

Attorney advertising. Published as a source of information only. The material contained herein is not to be construed as legal advice or opinion.

©2026 Katten Muchin Rosenman LLP.

All rights reserved. Katten refers to Katten Muchin Rosenman LLP and the affiliated partnership as explained at katten.com/disclaimer.