



Your AI Chats May Be Discoverable: Preservation and Privilege Considerations for Generative AI Prompts and Outputs

October 29, 2026

As generative AI tools become part of everyday business and legal work, a pressing question for litigants is whether prompts, outputs, and conversation logs must be preserved and produced in litigation. Recent decisions suggest that AI interactions are generally treated like any other electronically stored information (ESI), and courts have begun to define when those interactions may be shielded as irrelevant or as privileged, and when that protection can be lost.

Key Takeaways

- Once a duty to preserve attaches, preserve relevant AI prompts, outputs and related account settings.
- Expressly cover generative AI tools, both enterprise and consumer, in litigation hold notices.
- Do not assume AI interactions are ephemeral; most platforms retain histories that a court will consider within a party's control.
- Route litigation-related AI use through counsel and approved enterprise tools to strengthen work product protection.
- Avoid entering confidential or privileged information into open consumer AI tools.
- Before relying on AI outputs in a pleading or filing, weigh the risk that doing so will waive protection over the related prompts.

AI Prompts and Outputs Are ESI

AI prompts, outputs and conversation logs are generally treated as Electronically Stored Information (ESI) under Federal Rule of Civil Procedure 34. They are subject to ordinary discovery and preservation rules, and a failure to preserve them is analyzed under Rule 37(e). Storing communications on an AI platform does not relieve a party of its discovery obligations. Rule 34

reaches ESI in a party's "possession, custody, or control," which includes data on third-party servers that the party can access. A user who can log in and retrieve or export a chat history therefore likely controls those conversations. Courts have applied similar reasoning to text messages held by a third-party service provider (see *Flagg v. City of Detroit*, 252 F.R.D. 346 (E.D. Mich. 2008)), and to stored emails (see *Al Noaimi v. Zaid*, No. 11-1156-EFM, 2012 WL 4758048 (D. Kan. Oct. 5, 2012)).

Courts also have been willing to impose broad AI-specific preservation obligations. In *In re OpenAI, Inc. Copyright Infringement Litigation*, Magistrate Judge Wang ordered OpenAI to "preserve and segregate all output log data that would otherwise be deleted on a going forward basis." The order applied whether the data would be deleted at a user's request or under "numerous privacy laws and regulations." Order, Case No. 1:23-cv-11195-SHS-OTW (Dkt. No. 551) (S.D.N.Y. May 13, 2025). The district judge upheld the order over OpenAI's objection that it would require keeping billions of conversations, suggesting that courts will treat AI transcripts no differently than emails or text messages.

"Ephemeral" Arguments Are Unlikely to Succeed

Companies should be cautious about treating AI interactions as ephemeral or not reasonably discoverable. In January 2024, the DOJ Antitrust Division and the FTC updated their standard preservation letters, second requests and compulsory process, including grand jury subpoenas, to address "collaboration tools and ephemeral messaging platforms." The DOJ said it expects counsel to "preserve and produce any and all responsive documents, including data from ephemeral messaging applications," and warned that "[f]ailure to produce such documents may result in obstruction of justice charges." The FTC added that preservation obligations apply to "new methods of collaboration and information sharing tools, even including tools that allow for messages to disappear." The release does not mention AI tools by name, but its reasoning seems consistent with AI tools that may be described as "ephemeral" in nature.

While certain courts' early ESI protocols treated truly transient data as generally not discoverable (such as RAM or history, cache, cookies, etc.), that carve-out covers data that is not kept in the ordinary course. See *Travelers Prop. Cas. Co. of Am. v. Cannon & Dunphy*, S.C., 997 F. Supp. 2d 937 (E.D. Wis. 2014) (attaching then-relevant standing order). Most enterprise and consumer AI platforms save prompt histories. Unless a party can show that its AI interactions genuinely cannot be retrieved, courts will likely apply the usual case-specific relevance and proportionality analysis.

When AI Interactions May Be Protected

Relevance. Relevance remains a meaningful limit. In *In re OpenAI, Inc., Copyright Infringement Litigation*, 800 F. Supp. 3d 602 (S.D.N.Y. 2025), the court found that requested AI chat logs were not relevant, so their inadvertent loss did not implicate Rule 37(e). But relevance cuts both ways. In

Tremblay v. OpenAI, Inc., No. 23-cv-03223-AMO (RMI) (N.D. Cal. June 24, 2024), the magistrate judge rejected the plaintiffs' argument that their negative prompt results were irrelevant, holding that even if not relevant to the plaintiffs' claims, "they are plainly relevant to Defendants' defenses."

Attorney work product. Prompts drafted by in-house or outside counsel may be protected work product under Rule 26(b)(3)(B) if they reflect counsel's "mental impressions, conclusions, opinions, or legal theories" concerning litigation. However, this can be waived if the party relies on the outputs in the litigation. In *Tremblay*, the magistrate judge initially compelled production of all prompts and outputs, viewing account settings and negative test results as "more in the nature of fact work product." The district judge reversed, finding that the prompts "were queries crafted by counsel and contain counsel's mental impressions and opinions about how to interrogate ChatGPT, in an effort to vindicate Plaintiffs' copyrights." *Tremblay v. OpenAI, Inc.*, 2024 WL 3748003 (N.D. Cal. Aug. 8, 2024). *Concord Music Group, Inc. v. Anthropic PBC*, No. 24-cv-03811-EKL (SVK), Dkt. 377 (N.D. Cal. May 23, 2025), reached a similar result, protecting undisclosed attorney-generated prompts.

Waiver. In both *Tremblay* and *Concord*, however, the court further held that protection had been waived because the parties placed the prompts and outputs at issue in the complaint and/or other court filings. A party that relies on AI outputs in a pleading should expect a court to order production of the related prompts.

Consumer AI tools. The analysis may change when a third-party consumer AI tool is used. In *United States v. Heppner*, 820 F. Supp. 3d 292 (S.D.N.Y. Feb. 17, 2026), the court held that a criminal defendant's search prompts were not protected by attorney-client or work product privilege even when asking about information provided by his attorney. The decision turned on three factors: (1) the inquiries were not communications between the defendant and his counsel; (2) counsel did not direct the communications; and (3) the prompts were not confidential because the platform disclaims confidentiality. The first two points are important but not alarming – as a matter of general practice, non-attorneys should be instructed that actions undertaken independent of counsel's direction are not privileged.

The third issue, though, is concerning for any attorney who uses AI for research because it suggests that communications with AI do not have the same protections as communications through other electronic means. That said, it is not clear that this analysis would hold up if the underlying decision were about an attorney's own legal research. In addition, the decision expressly notes that "Anthropic reserves the right to disclose such data to a host of 'third parties,' including 'governmental regulatory authorities,'" so it likely would not apply to confidential, internal systems.

Separately, the Delaware Court of Chancery entered a stipulated protective order providing that "[p]roviding Confidential Discovery Material to an open [Generative AI] tool is considered a disclosure

to a third party.” *Slam Corp. v. Lynk Glob., Inc.*, C.A. No. 2025-0693, 2025 WL 1842381 (Del. Ch. July 2, 2025). While the entry of that order is not a holding, it may signal judicial acceptance that inputting information into an open AI tool constitutes disclosure outside the litigation team, which is significant for the attorney-client privilege.

Nature of the task. Whether the type of task matters turns on a standard work-product analysis: does the prompt reflect the attorney’s mental impressions or case strategy? A prompt asking an AI tool to draft an argument, test a legal theory or frame an investigation is more likely to be protected opinion work product. A routine factual lookup, or a prompt that simply restates facts, is less likely to be protected. Business use by non-attorneys, with no anticipation of litigation, will likely receive no protection at all.

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.