Federal Court Rules Client's AI-Generated Documents Not Privileged

February 25, 2026

In a case that has significant implications for those using AI tools for legal matters, on February 10, 2026, the United States District Court for the Southern District of New York held that documents a client created using a public artificial intelligence platform were not protected by either the attorney-client privilege or the work product doctrine. The decision represents one of the first federal rulings to directly address whether AI-generated materials are privileged and signals important considerations for parties incorporating public AI tools into their legal work.

What Happened?

The case, *United States v Heppner*, involves a criminal defendant indicted in October 2025 on securities fraud and related charges. Following his arrest, federal agents executed a search warrant and seized electronic devices from his residence. Defense counsel informed the government that the defendant used a commercial AI platform to generate approximately 31 documents containing queries and responses related to the government's investigation prior to his arrest, stored on the seized devices. Defense counsel asserted these “AI Documents” were privileged and unusable at trial.

At a February 10 conference, Judge Rakoff orally ruled from the bench in the government's favor, stating he saw "not remotely any basis for any claim of attorney-client privilege."

Why Did the Court Reject Attorney-Client Privilege?

The court found that communications with an AI tool are not communications with an attorney, agreeing with the government's analogy that using AI for legal research is no different than asking non-lawyer friends for input on a legal situation-which does not create privilege.

The court further noted that the AI tool’s terms of service expressly disclaim any attorney-client relationship and that its privacy policy permits collecting user prompts and disclosing them to governmental authorities-negating any reasonable expectation of confidentiality. Without such an expectation, privilege cannot exist.

Critically, the court rejected the argument that transmitting the AI Documents to attorneys created privilege. Well-settled law holds that pre-existing, non-privileged materials do not become privileged merely because a client later shares them with counsel.

Why Did the Court Reject Work Product Protection?

Defense counsel also asserted work product protection for the AI documents. This doctrine protects documents prepared in anticipation of litigation. The court noted that the defendant created the AI Documents "of his own volition" and not at counsel's direction. Because they did not represent Defense Counsel's strategy, they could not be attorney work product. The Court also concluded that Rule 16(b)(2) was inapplicable because the AI reports were seized at arrest pursuant to a search warrant rather than during discovery. Defense Counsel failed to challenge the warrant's validity, leaving unclear how Rule 16(b)(2) will apply to AI materials if this rationale is used again.

What Are the Risks?

This ruling highlights the risks inherent in Courts grappling with the rapidly expanding use of AI tools in legal matters. These issues are arising quickly, and the technology and how the public interacts with generative AI evolves daily.

Following this ruling, the public should recognize that queries to public AI tools may not be privileged, and sending AI outputs to counsel after the fact likely does not retroactively create privilege. It will take time for courts, and potentially legislatures, to fully resolve these evidentiary issues and provide clear guidance. Meanwhile, significant uncertainty remains on how these issues will be handled case-by-case.

Could This Ruling Face Reversal on Appeal?

This ruling may face significant challenges on appeal. The Court's reasoning that the AI platform's privacy policy created a *per se* waiver of confidentiality rests on grounds that could apply to virtually every digital communication platform attorneys and clients use daily. Many AI platform privacy policies are structurally identical to those of Gmail, Outlook, Slack, Microsoft Teams, and other common platforms. All reserve rights to collect user data and permit disclosure to governmental authorities under certain circumstances. None create an affirmative "expectation of privacy" in the legal sense.

If the standard is whether a platform's privacy policy disclaims confidentiality, the same logic would theoretically apply to all email systems. Under this reasoning, attorneys would need to review every email provider's privacy policy before sending privileged documents. If a client uses a free Gmail account, would that waive privilege over everything the attorney sends? Courts have historically answered no, because privilege analysis focuses on reasonable expectations of confidentiality in practice-not boilerplate terms drafted to maximize corporate flexibility.

Moreover, users reasonably expect AI chats to be private in practice. These platforms are typically protected by login credentials, often with multi-factor authentication. Sharing chat results requires affirmative user action, and conversations are siloed to the specific user. The functional experience is arguably more private than email, where messages traverse multiple servers and can be forwarded without restriction.

The Court did not appear to engage with these considerations in detail. It accepted the government's framing that the AI platform's privacy policy created a *per se* waiver, moved to the work product analysis, and ruled without fully addressing how this reasoning might distinguish AI platforms from the ubiquitous digital tools attorneys rely upon daily. This gap may provide grounds for appellate review.

For example, on the same day, in *Warner v. Gilbarco, Inc.,*, another court ruled that a *pro se* litigant’s AI outputs are protected by work product doctrine. That court reasoned that generative AI is not a person but a tool, so sharing impressions with it does not waive privilege. Notably, because the Michigan case involved a *pro se* litigant functioning as her own lawyer, the question of whether the AI chat was conducted at counsel's behest was not at issue. Nevertheless, the court's identification of AI as a tool-disclosure to which cannot break privilege-stands in opposition to the Southern District of New York’s decision.

What Steps Should You Take Now?

The main takeaway: individuals and companies aware of an investigation, regulatory inquiry, or potential litigation should be cautious about immediately logging into Claude, ChatGPT, or other AI platforms for guidance. Treat your AI prompts and outputs as potentially discoverable and exercise caution with public platforms. Best practice is to engage legal counsel and discuss the matter with them promptly to reduce the risk of your inquiries becoming public record. If you have questions about how this decision may impact your use of AI tools or litigation strategy, please contact one of the Saiber attorneys familiar with these issues: Jack Losinger, Esq., Michael Shortt, Esq., or Katherine Escanlar, Esq.