

Akerman Intelligence Alert

AI, Privilege & Work Product: The Current Legal Landscape and Practical Guidance

March 16, 2026

By [Darryl R. Graham](#)

The rapid adoption of artificial intelligence (AI) in legal practice has created new challenges for the application of attorney-client privilege and work product protection. Two recent federal court decisions — *United States v. Heppner* and *Warner v. Gilbarco, Inc.* — offer early substantive guidance on how courts may treat information processed by AI tools in both criminal and civil contexts. However, these cases also highlight the unsettled nature of the law and the need for attorneys and clients to proceed with caution and clear documentation when using AI in connection with legal matters.

This practice update analyzes the key holdings and reasoning in *Heppner* and *Warner*, identifies the open questions that remain, and provides actionable recommendations for clients and counsel seeking to minimize risk and preserve privilege in the evolving AI landscape.

I. The Recent AI Privilege Cases

A. *Heppner*: Privilege and Work Product Denied for Client-Directed AI Use

In *Heppner*,^[1] as we discussed previously ([Court Rules That Information Disclosed by Layperson to AI Tools Is Not Protected by Attorney-Client or Work-Product Privileges - Akerman LLP](#)), the defendant

Related People

[Darryl R. Graham](#)

Related Work

[Litigation](#)

Related Offices

[New York](#)

Akerman Intelligence

Akerman Intelligence is Akerman's platform for AI strategy, research, and thought leadership. Through original analysis, client collaboration, and strategic partnerships, we help organizations navigate the legal, operational, and governance challenges of artificial intelligence.

independently used the consumer version of the Claude AI tool to run queries related to a criminal investigation, without any direction or oversight from counsel. After the FBI seized his devices and found 31 Claude-created documents, defense counsel asserted privilege. The court rejected this claim, holding that:

- Claude is not an attorney, cannot provide legal advice, and cannot form an attorney-client relationship.
- The consumer AI's terms of service did not provide a reasonable expectation of privacy or confidentiality.
- The defendant's self-directed use of AI was not protected by attorney-client privilege or work-product doctrine.
- Disclosure of privileged information to the AI tool (deemed a "third party") constituted a waiver of privilege.

The court emphasized that privilege and work-product protection are unlikely to apply where a client uses a consumer AI tool independently, especially in the absence of confidentiality safeguards and attorney involvement. The court further warned that sharing privileged communications with an AI tool is likely to waive privilege, and that human counsel should be directly involved in all AI interactions, ideally using enterprise-tiered products with robust confidentiality protections.

B. *Warner*: Work-Product Protection for Pro Se Litigant's AI Use

In *Warner*,^[2] a *pro se* plaintiff relied heavily on ChatGPT to prepare and litigate her employment discrimination case. Defendants sought discovery of all documents and information related to her use of AI tools, arguing that her use of ChatGPT constituted disclosure to a third party and thus waived work-

product protection. The court disagreed, holding that:

- The *pro se* plaintiff's AI-assisted materials were protected as opinion work product under Fed. R. Civ. P. 26(b)(3)(A).
- AI tools were considered “tools, not persons,” so using ChatGPT did not constitute disclosure to a third party or waiver.
- The court denied the defendants' motion to compel, finding the request irrelevant and a “fishing expedition.”

Notably, however, *Warner* appears to limit its holding to affording work-product protection to the mental impressions and opinion work product of counsel (extended here to a *pro se* litigant). On that point, *Warner* and *Heppner* seem aligned, as they focus on the thought processes of *counsel*, not the client.

II. Where the Cases Agree and Diverge

Both *Heppner* and *Warner* recognize that the core purpose of privilege and work product protection is to safeguard the mental impressions and legal strategies of counsel (including *pro se* litigants). However, the cases diverge sharply on the question of waiver and the status of AI tools:

- *Heppner* treats AI as a “third party,” so disclosure to an AI tool waives privilege.
- *Warner* treats AI as a “tool,” so use of ChatGPT does not waive work-product protection.

This split underscores the current legal uncertainty and the present need for conservative, well-documented practices until courts or rulemakers provide further clarity.

III. Open Questions and Practical Risk Factors

Despite these early rulings, several critical questions remain unresolved:

1. Is AI a “third party” or a “tool” for waiver purposes?

If AI is deemed a third party, any disclosure of privileged information could waive privilege. If AI is a tool, there is no waiver. Until this issue is resolved, the prudent approach is to assume AI is a third party for privilege analysis.

2. Do enterprise-tier AI tools provide sufficient confidentiality?

Heppner found that consumer AI tools lacked adequate confidentiality, as their terms appear to allow for data retention, training, and disclosure to third parties. Enterprise-grade tools with robust confidentiality terms *may* strengthen privilege arguments, but this remains untested in the courts.

3. Can AI tools be treated as *Kovel*-type “agents”?

Heppner provides some support for the argument that, if used at counsel’s direction and with sufficient safeguards, AI tools could be considered agents under the *Kovel* doctrine. However, this approach is not yet settled law and carries some risk. The Second Circuit’s analysis in *Ackert* may also inform this analysis — there, the court rejected *Kovel* privilege where the third party was “sought out” for information neither the client nor the lawyer had, rather than to “translate or interpret” information the client provided. *See United States v. Ackert*, 169 F.3d 136, 139 (2d Cir. 1999).

IV. Current Privilege and Work-Product Analysis

Attorney-Client Privilege

Warner did not address attorney-client privilege, so *Heppner* remains the primary guide. The key points are:

- A client alone cannot establish or sustain an attorney-client relationship with an AI tool, regardless of confidentiality or privacy protections.
- Privilege *may* extend to a client's use of AI if: (i) the use is at the direction of counsel, (ii) robust confidentiality safeguards are in place, and (iii) the use is for the purpose of obtaining legal advice.
- All elements should be contemporaneously documented, including the attorney-client relationship, the purpose of AI use, and counsel's involvement.

Work-Product Protection

The work-product doctrine protects materials prepared in anticipation of litigation. Both *Heppner* and *Warner* focus on the mental impressions of counsel (or a *pro se* litigant in civil cases). However, neither case explicitly extends protection to non-attorney parties or representatives acting independently. Therefore, the safest approach currently is to ensure that any client use of AI for litigation preparation is at the direction of counsel and clearly documented.

Waiver

Heppner and *Warner* diverge on whether disclosure to an AI tool constitutes waiver. Until this issue is settled, treat all disclosures to consumer AI tools as if they are to a third party, which could waive privilege. To maintain protection, use only enterprise-tiered AI tools with robust confidentiality safeguards, and only at the direction of counsel.

V. The Risk Spectrum for AI Use in Legal Contexts

The following table summarizes the risk levels associated with various AI use scenarios:

Scenario	Risk Level	Assessment Basis
Consumer AI, client-directed, no counsel involvement	HIGH	Privileges likely unavailable and waived; consumer privacy policy destroys confidentiality; no counsel direction for work product
Consumer AI, privileged info disclosed to tool	HIGH	Likely waiver of underlying privilege; voluntary disclosure to “third party” with permissive data practices
Consumer AI, attorney-directed, no confidentiality terms	MODERATE-HIGH	Work product may apply; privilege doubtful; no confidentiality assurance undermines <i>Kovel</i> for privilege
Enterprise AI, client-directed, no counsel involvement	MODERATE-HIGH	Better confidentiality argument; work product still weak without counsel direction; no privilege per <i>Heppner</i>
Enterprise AI, attorney-directed, with confidentiality terms	LOW-MODERATE	Strongest available position; still untested on privilege/ <i>Kovel</i> ; work-product protection stronger

VI. Best Practices: The Lawyer-in-the-Loop Standard

Given the current uncertainty, the most defensible approach is to ensure that all AI use in legal matters is directed and overseen by counsel, and that robust confidentiality safeguards are in place. Both *Heppner* and *Warner*, read together, point toward a governing principle we refer to as the “lawyer-in-the-loop” standard: privilege and work-product protection are most defensible when attorneys are meaningfully involved in directing and overseeing AI use in the legal context. The following best practices are recommended:

- **Engagement letters and retainer agreements** should specifically address AI use, authorize counsel-directed AI tools, specify enterprise-grade platforms with confidentiality provisions, and obtain informed client consent where required.
- **Firm-wide AI use policies** should distinguish permissible attorney-directed AI use from impermissible client-independent use, and specify approved platforms.

- **Exclusive use of enterprise-tier AI** for all legal work, with the choice of platform documented.
- **Explicit client counseling** about the risks of independent AI use, including the risk of destroying privilege and work-product protection.
- **Litigation hold notices** should explicitly address whether employees are permitted to use AI tools to summarize documents, draft responses, or analyze materials covered by the hold. If AI use is permitted, define the scope, require logging, and treat prompts and outputs as potentially discoverable.
- **Review and update *Kovel* agreements** with retained professionals to address the AI tools those professionals use.
- **Protective order negotiations and Rule 26(f) conferences** should address the parties' AI use policies and whether existing orders need to be modified to address AI platform uploads.
- **All client use of AI relating to legal matters** should be at the direction of counsel and exclusively through an enterprise-tiered product with sufficient confidentiality safeguards. The client should note counsel's direction and involvement in each interaction with AI, and all outputs should be kept confidential and shared exclusively with counsel and/or at counsel's direction.

Conclusion

Heppner and *Warner* provide early but divergent guidance on the intersection of AI, privilege, and work product. Until the law is more firmly settled, the safest and most defensible approach is to treat AI as a powerful tool that must be used within the attorney-client relationship, not as a substitute for it. Clients who use AI autonomously on legal matters, without counsel's direction and oversight, risk potential waiver of privilege and work-product protection. By adopting thoughtful, well-documented practices and using enterprise-grade AI tools under

attorney supervision, clients and counsel can best position themselves to defend privilege claims in this rapidly evolving area.

We will continue to monitor developments and update our analysis as courts, bar associations, and rulemakers provide additional guidance.

[1] *United States v Heppner*, 25 CR. 503 (JSR), 2026 WL 436479 (SDNY Feb. 17, 2026) (oral ruling issued Feb. 10, 2026; written memorandum Feb. 17, 2026).

[2] *Warner v Gilbarco, Inc.*, 2:24-CV-12333, 2026 WL 373043 (ED Mich Feb. 10, 2026).