GUNDERSON DETTMER

Client Insight: AI Tools and Confidentiality: Practical Guidance for Companies for Protection Under Attorney-Client Privilege and Work-Product Doctrine

Insights
September 23, 2026

Federal and state courts have begun to address whether prompts, outputs and uploaded materials created using AI tools are protected from disclosure. Courts are applying existing attorney-client privilege and work-product doctrine rules to AI tools, with outcomes turning on who used the AI tool, why the materials were created, whether legal counsel directed the work and whether the AI workflow was used with the intent to preserve confidentiality.

For companies, the practical message is thoughtful usage. A conversation with a public chatbot on its own is not protected merely because it concerns a legal issue. Although AI-assisted preparation in anticipation of litigation may qualify as work product in some circumstances, the recent rulings are narrow, fact-specific and largely determined at the trial level. AI chat histories may also be subpoenaed, obtained through a warrant, preserved in discovery or used as evidence of intent.

In this article, you will find:

Key Takeaways

- **There is no automatic protection for AI chat history, even if communicated to an attorney.** An AI chatbot is not a lawyer. Sending an AI-generated document to counsel does not retroactively protect the underlying chatbot exchange as privileged or work product.
- **Protection under the work-product doctrine is not guaranteed even when involving the preparation of litigation materials.** Several courts ruled to protect AI-assisted litigation work product, but most of these cases involved self-represented litigants or a state rule that expressly protects materials prepared by a party. Notably, these rulings concern protection under the work-product doctrine, not attorney-client privilege, and do not establish that employees of a represented company may freely use consumer AI tools for sensitive legal matters.
- **Enterprise AI terms and account settings may help support confidentiality, but they do not determine privilege or work-product protection on their own.** Use of a free or personal account may undermine the confidentiality required for attorney-client privilege. Enterprise controls restricting model training, retention, disclosure and human review may support an intent for confidentiality, but they do not create attorney-client privilege or work-product protection on their own.
- **AI prompts and outputs can become evidence against a company.** Company AI chats have already been quoted in court as direct evidence of knowledge, motive and intent. They may also be discoverable, subject to preservation obligations and obtained through subpoenas or warrants. Companies should address these risks through internal and external AI policies, including appropriate disclaimers and guardrails.

- **Mitigate these risks with deliberate AI governance.** Companies should route sensitive AI use through counsel and approved enterprise or legal-first tools, keep privileged, litigation and other high-risk materials out of unapproved tools and meeting assistants, and adopt a written AI Usage Policy with training that reaches founders, executives and board members. For more details, see “Practical Guidance for Companies” below.

Practical Guidance for Companies

Distinguish attorney-client privilege from work-product protection. ***Attorney-client privilege*** generally protects confidential lawyer-client communications made for legal advice, while the ***work-product doctrine*** protects qualifying materials prepared in anticipation of litigation. Depending on the governing rule, work product may cover materials prepared by a party or its representative even when the AI user is not an attorney. Since these doctrines have different requirements and waiver standards, companies should identify which protection they are seeking before selecting an AI workflow.

- **Use approved enterprise or legal-first AI environments.** Do not rely on a free consumer account or an employee’s personal account. Review the enterprise terms and settings for model training, retention, deletion, human review, sub-processors, access controls and notice of government or third-party demands. These controls may support the confidentiality required for attorney-client privilege, but do not establish privilege or work-product protection. For work-product protection, the purpose of the materials and their connection to anticipated litigation also matter.

- **Keep sensitive information out of unapproved tools**. Employees should not upload attorney-client communications, materials prepared in anticipation of litigation, board materials, investigation notes, trade secrets, source code, personnel records, customer data, diligence materials or confidential discovery without express approval for the specific workflow. Privileged communications and potential work product should be evaluated separately because the protections attach and may be waived under different standards.

- **Involve legal counsel before using AI on a legal matter.** The legal team can help define the task, confirm whether AI use is appropriate and identify the permitted tool, data and workflow. Counsel can also assess whether the task involves confidential communication for legal advice, litigation preparation or

ordinary business activity, and whether attorney-client privilege, work-product protection or neither is likely to apply.

- **Control AI meeting assistants.** Disable or tightly manage transcription and AI note-taking tools during communications with counsel, board meetings and internal investigations. These tools create additional records and may introduce third-party confidentiality issues that affect attorney-client privilege. Materials created for litigation may raise separate work-product questions.

- **Adopt and enforce an AI Usage Policy.** Policies should distinguish public, enterprise and legal-approved tools, identify prohibited data and high-risk uses, and require escalation to legal, privacy and security teams. Policies and training should explain the different requirements for attorney-client privilege and work-product protection and make clear that prompts and outputs may be company records subject to preservation and discovery. Training should include founders, executives and board members, not only technical personnel. Your Gunderson Dettmer team can assist with drafting an AI Usage Policy and preparing internal AI best practice trainings.

Recent Court Decisions

1. **Materials independently created using a consumer AI tool and later shared with counsel may not be protected: *United States v. Heppner* (February 17, 2026).**

   In *Heppner*, a represented criminal defendant used the free, public version of Anthropic's Claude to develop potential defenses and legal arguments, then later shared the materials with counsel. On the government's motion for a ruling that the materials were unprotected, Judge Rakoff of the Southern District of New York held that:

   - **Attorney-client privilege** did not apply because (1) the defendant, and not counsel, communicated with Claude, (2) the consumer AI terms undermined any reasonable expectation of confidentiality and (3) the defendant used Claude on his own initiative rather than at counsel's direction.

   - **Work-product protection** did not apply because, even assuming the documents were prepared in anticipation of litigation, they were not prepared by or at the behest of counsel and did not reflect defense counsel's strategy at the time they were created.

In reaching that conclusion, the court emphasized that the work-product doctrine "at its core" protects the mental processes of attorneys, providing "a privileged area within which counsel can analyze and prepare his client's case." That framing is the reason this decision matters: the protection tracks counsel's involvement, not the sensitivity of the subject.

The court left open the possibility that a different AI workflow could be analyzed as an agent of counsel if counsel directed the use and the traditional requirements were satisfied. Notably, that observation is not a safe harbor – an enterprise account alone would not establish protection and the intent for confidentiality of an AI work product must be apparent. The decision confirms that any legal AI workflow should be structured at the outset with the applicable attorney-client privilege or work-product protection requirements in mind.

2. **AI-assisted litigation preparation may qualify for protection under the work-product doctrine: *Warner, Morgan, Tate* and *Assini*.**

   Four recent trial-level civil decisions resulted in at least partial protection of litigation materials prepared with the assistance of AI under the work-product doctrine, rather than attorney-client privilege. Notably, the first of these decisions was issued one week before *Heppner*, and the courts that followed had to address the apparent tension directly.

   - ***Warner v. Gilbarco* (February 10, 2026):** The Eastern District of Michigan magistrate judge denied discovery into a self-represented plaintiff's use of ChatGPT, reasoning that the work-product doctrine protects qualifying litigation materials prepared by a party (including the plaintiff's own mental impressions) and that work product is not ordinarily waived merely by disclosure to a non-adversary. The court characterized generative AI programs as "tools, not persons." The court also held, as a threshold matter, that the AI-related discovery sought was not discoverable in the first instance.

   - ***Morgan v. V2X* (March 30, 2026):** Following *Warner* and distinguishing *Heppner* on the ground that a pro se litigant acts as both party and advocate, the District of Colorado held that the work-product protection extends to a self-represented plaintiff's AI-assisted litigation preparation, and that inputting information into a public AI tool does not automatically waive that protection. However, protection did not extend to the identity of the tool itself, which was disclosed under court order. Further, in a protective order, the court barred inputting confidential information into any AI platform unless the provider is

contractually prohibited from (1) storing *or* using inputs to train or improve its model and (2) disclosing inputs to any third party except where essential to delivering the service, in which case that third party must be bound by obligations no less protective than the order itself. The provider must also contractually afford the party the ability to delete all confidential information on request.

- ***Tate Group Automotive v. Legacy Automotive Capital* (June 3, 2026):** A Texas business court held that most ChatGPT conversations created by a non-lawyer company principal were protected under Texas's work product rule, which expressly covers material prepared, or mental impressions developed, in anticipation of litigation by or for a party. The court ordered a small number of pages produced as non-work product and separately ordered the plaintiff to identify which discovery materials had been shared with ChatGPT to test compliance with the existing protective order. *Tate* is notable for extending protection to a non-lawyer company principal, but it remains a narrow trial-level ruling specifically under Texas procedural law.

- ***Assini v. Hayward* (June 4, 2026):** A New York trial court quashed broad non-party subpoenas to OpenAI seeking a self-represented defendant's litigation-related prompts, uploads and outputs, holding the materials qualified as "litigation preparation material" and that the subpoenaing parties had not shown the substantial need and undue hardship required to overcome New York's conditional protection. The court followed *Morgan* and declined to follow *Heppner*, but cautioned that AI use cannot go "unfettered" and directed compliance with New York's new rules regulating the use of AI by attorneys and parties in New York state courts (*see* Part 161, effective June 1, 2026).

These decisions suggest that AI-assisted litigation materials can sometimes be shielded from the opposing party under the work-product doctrine. However, these determinations have arisen primarily in cases where a party represented themselves or where a state rule expressly protects material prepared by or for a party. In those cases, courts treated the person as both client and advocate, so their AI prompts and outputs qualified as material prepared for litigation, and entering information into a public AI tool did not by itself waive the protection. Note, however, that the scope is limited: courts have still ordered disclosure of collateral facts, including which AI tool the party used and the confidential materials fed into such AI tool.

These rulings do not extend attorney-client privilege to AI conversations and do not protect ordinary business use of AI (e.g., market research, compliance analysis, employment decisions, board strategy and independent legal research) simply because the subject matter is legally sensitive. As a practical matter, parties should assume material entered into general-purpose AI tools may be discoverable absent attorney direction, and should confirm that provider terms bar training on inputs and afford a deletion right.

3. **AI chat histories may be exposed through discovery and used as evidence: *Fortis Advisors v. Krafton* (March 16, 2026) and *United States v. Kim* (June 22, 2026)**

   These cases do not resolve when attorney-client privilege or work-product protection applies, but they illustrate the consequences when no protection applies or before objections are resolved.

   - ***Fortis Advisors v. Krafton* (March 16, 2026):** The Delaware Court of Chancery found that an acquirer breached an equity purchase agreement by terminating the target's founders and CEO on a pretext to avoid an earnout of up to $250 million, relying on the acquirer's CEO's ChatGPT prompts asking how to avoid the payment as direct evidence that the stated grounds for termination were manufactured. Notably, the timing was as damaging as the content: an executive who turns to an AI tool before counsel on terminations, investor disputes, earnout obligations or regulatory exposure risks creates a record of the company's intent and state of mind.

   - ***United States v. Kim* (June 22, 2026):** The Southern District of New York permitted execution of a Stored Communications Act warrant for OpenAI records, including stored prompts and responses and account-related data, while leaving open arguments concerning privilege, suppression, admissibility and limits on the government's use of particular records for later judgment. The decision does not hold that every AI chat is unprotected, but it shows that data stored by AI providers may be produced in court before such issues are resolved.

The cases are a reminder that, absent applicable privilege, work-product protection or another limitation, relevant AI chats may be obtained from a user or provider through civil discovery, subpoena or warrant and used as evidence.

Conclusion

The law remains unsettled, and the current decisions are narrow trial-level rulings. Several of the more protective outcomes arise from unusual facts or procedural rules that may not apply to a represented company and, thus far, no appellate decision has decisively resolved when generative AI prompts and outputs qualify for attorney-client privilege or work-product protection.

Companies should therefore manage AI prompts and outputs as potentially discoverable records while structuring any legal use deliberately with counsel. That analysis should separately consider whether the workflow is intended to preserve confidential lawyer-client communications or to create materials in anticipation of litigation. Appropriate governance cannot guarantee attorney-client privilege or work-product protection, but it can materially improve the company's position and reduce the risk that a sensitive AI exchange becomes an avoidable exhibit.

If you have questions about any of the developments discussed above, please contact your Gunderson Dettmer attorney.

Related Services

AI & Machine Learning

Data Privacy

North America

Strategic Transactions & Licensing

Featured Insights

CLIENT NEWS

Tidemark Leads Adaptive's $30 Million Series B

PUBLIC VENTURES

Glass Lewis Invites Comment on Proposed 'Four Perspectives' Proxy Voting Research Model

CLIENT NEWS

Matter Venture Partners Closes $450 Million Fund II

FIRM NEWS

Gunderson Dettmer Commemorates 2026 Hispanic Heritage Month

CLIENT NEWS

Andreessen Horowitz Leads Lightfield’s Series A

CLIENT NEWS

Leyden Labs Acquires RIGImmune

CLIENT NEWS

Savvy Wealth Raises $100 Million Series C

CLIENT NEWS

Bluecore Energy Announces $50 Million Seed Financing

CLIENT NEWS

BrainChild Bio Closes $116 Million Series A

CLIENT NEWS

Chile-Based Magnar Announces $8 Million Series A

CLIENT NEWS

Point72 Ventures and Spark Capital Co-Lead Stoke Space’s $1 Billion Series E

PUBLIC VENTURES

SEC Files Subpoena Enforcement Action Against ISS