

CLARION AI LAW

CLIENT MEMORANDUM • AI GOVERNANCE & LITIGATION SERIES

Generative AI, Privilege, and Discovery: What Every Organization and Its Counsel Must Know

How U.S. courts, judicial conferences, and ethics authorities now treat AI prompts, conversations, and outputs — and what to do about it

TO: Clients and Friends of Clarion AI Law
FROM: Bennett B. Borden, Clarion AI Law
DATE: July 31, 2026
RE: Effect of generative AI on attorney-client privilege and work-product protection; discoverability of AI prompts and outputs; and governing court, judicial, and ethics guidance

How to read this memorandum. Part I is a stand-alone executive summary written for business leaders and boards — it gives you the bottom line and the actions to take. Parts II–VI provide the detailed legal analysis for counsel, organized by each of the three questions we researched. Part VII sets out recommended best practices. The Appendices contain case-by-case digests, a snapshot of court AI-disclosure rules, and a source index.

Contents

- I.** Executive Summary
- II.** Why This Matters Now
- III.** The Legal Framework in Brief
- IV.** Question 1 — Does Using Generative AI Affect Privilege or Work-Product Protection?
- V.** Question 2 — When Are AI Prompts, Conversations, and Outputs Discoverable?
- VI.** Question 3 — Court, Judicial, and Ethics Guidance on Lawyers' Use of AI
- VII.** Recommended Best Practices
- VIII.** How Clarion AI Law Can Help
- A.** Appendix A — Case Digests
- B.** Appendix B — Snapshot of Representative Court AI-Disclosure Requirements
- C.** Appendix C — Key Authorities and Source Index

I. Executive Summary

Generative AI is now woven into everyday work. Employees and clients routinely use tools such as ChatGPT, Microsoft Copilot, Claude, and Google Gemini to draft correspondence, summarize documents, analyze data, and think through problems — including problems that are, or may become, the subject of litigation. Every one of those interactions creates a record: a prompt, an output, and an activity log. In a remarkably short window in early 2026, federal courts began deciding, for the first time, whether those records are protected by the attorney-client privilege and the work-product doctrine, when they are discoverable by an adversary, and how lawyers may ethically use these tools at all.

We commissioned focused research on three related questions. This memorandum consolidates the results. The short answers follow; the detailed analysis is in Parts IV through VI.

Question 1 — Does using generative AI affect attorney-client privilege or work-product protection?

For the most part, yes — and not in your favor if you are careless. Communications with a *public* AI platform are almost certainly *not* protected by the attorney-client privilege: the AI is not your lawyer, the platform’s privacy terms defeat any expectation of confidentiality, and the tool disclaims giving legal advice. Work-product protection, by contrast, *can* survive AI use — courts have held that an AI tool is a “tool, not a person,” so using it is not a disclosure to an adversary — but only when the material was prepared in anticipation of litigation. Critically, when a *represented* client uses AI on their own, without counsel’s direction, they risk losing *both* protections. The law is developing at the district-court level and the leading decisions are already in tension; no appellate court has resolved the core questions.

Question 2 — When are AI prompts, conversations, and outputs discoverable in litigation?

They are treated like any other electronically stored information. Prompts, chats, and outputs are ESI subject to the ordinary rules of relevance, proportionality, privilege, and work product — courts have refused to create any AI-specific exemption. Prompts crafted by counsel to interrogate an AI generally qualify as *opinion* work product (the most protected category). Merely using an AI tool does not waive protection. But a party that puts its AI investigation *at issue* — for example, by relying on AI outputs to prove a claim — can be forced to produce the full investigation, including unfavorable results. Courts are also beginning to write AI restrictions directly into protective orders.

Question 3 — What guidance have courts and ethics authorities issued on lawyers’ use of AI?

No blanket ban, but your existing duties fully apply. The ABA and state bars treat AI as permissible so long as lawyers satisfy their duties of competence, confidentiality, communication with clients, candor to the court, supervision, and reasonable fees. On top of that, a fast-growing patchwork of court standing orders and local rules requires attorneys to disclose or certify AI use and to verify AI output before filing. Courts have already imposed sanctions — including public reprimand, disqualification, and bar referrals — on lawyers who filed AI-“hallucinated” citations without checking them.

Bottom line for your organization

- **Treat public AI tools as non-confidential.** Assume anything typed into a consumer AI platform could end up in an adversary’s hands and is likely unprotected.
- **Route litigation-related AI use through counsel.** Work-product protection turns on counsel’s direction and on preparation in anticipation of litigation — not on the client’s good intentions.

- **Use closed, enterprise tools with the right contract terms.** No training on your inputs, no third-party disclosure, and deletion on request — the emerging baseline courts now demand before confidential information may touch an AI system.
- **Verify everything.** Every AI-generated citation, quotation, and factual assertion must be independently checked. “Read and reason” — do not simply confirm a case is real.
- **Govern it.** Adopt an AI-use policy, train your people, cover AI in litigation holds, and address how your clients and employees use AI on their own.

II. Why This Matters Now

Three developments make this an urgent governance issue rather than an academic one.

First, adoption is universal and largely invisible. AI use is no longer confined to an innovation team. It is happening in every department, often through free consumer tools and personal accounts, and often without any record that it occurred. Each interaction generates prompts, outputs, and logs that may contain relevant — and discoverable — information about a matter.

Second, the governing law crystallized almost overnight. The entire body of directly on-point federal authority on AI and privilege was produced within roughly a six-week window in early 2026: *United States v. Heppner* (S.D.N.Y., Feb. 17, 2026), *Warner v. Gilbarco, Inc.* (E.D. Mich., Feb. 10, 2026), and *Morgan v. V2X, Inc.* (D. Colo., Mar. 30, 2026). The pace of decisions on discovery, protective orders, disclosure, and sanctions has only accelerated since.

Third, the exposure runs in two directions. An organization can *lose* protection over its own analysis by using the wrong tool the wrong way; and it can face *sanctions and credibility damage* if AI-generated errors reach a court. Both risks are readily managed with modest, well-understood safeguards — but only if they are put in place before a dispute arises.

III. The Legal Framework in Brief

The new AI decisions do not invent new doctrine. They apply four familiar bodies of law. Understanding them is essential to understanding the results.

A. AI materials are electronically stored information (ESI)

Federal Rule of Civil Procedure 34 authorizes discovery of documents and ESI “stored in any medium.” Prompts, AI conversations, outputs, and logs are stored electronically and therefore qualify as ESI. Under Rule 26(b)(1), they are discoverable if they are relevant to a claim or defense and proportional to the needs of the case, and if they are not privileged or protected. As one court put it, AI’s novelty does not exempt it from longstanding legal principles — the same rules that govern other documents govern AI materials.

B. Attorney-client privilege

Under Federal Rule of Evidence 501 and federal common law, the privilege protects communications (1) between a client and an attorney, (2) intended to be and in fact kept confidential, (3) for the purpose of obtaining or providing legal advice. All three elements must be present. Federal Rule of Evidence 502 governs the scope of waiver — including the protection against inadvertent waiver where the holder took reasonable steps to prevent and to rectify disclosure, and the availability of a Rule 502(d) order to guard against waiver in a proceeding.

C. The work-product doctrine

Rooted in *Hickman v. Taylor* (1947) and codified at Rule 26(b)(3), the doctrine protects materials prepared in anticipation of litigation by or for a party or its representative. Courts distinguish *fact* work product (discoverable on a showing of substantial need and undue hardship) from *opinion* work product — an attorney’s mental impressions, conclusions, and legal theories — which is protected far more stringently and is “almost always” exempt from discovery.

D. The pivotal distinction: how privilege waiver differs from work-product waiver

This contrast drives nearly every AI result below. The attorney-client privilege is fragile: virtually *any* voluntary disclosure to a third party can waive it. Work-product protection is far more durable: it is waived only by disclosure *to an adversary*, or in a manner that substantially increases the likelihood an adversary will obtain the material. That is why the same act — typing into a public AI tool — can forfeit privilege while leaving work-product protection intact.

The through-line: Three variables decide almost every case — (1) *who* used the tool and whether counsel directed it; (2) *what the platform's terms say* about collecting, training on, and disclosing inputs; and (3) *whether the party put its AI use at issue* in the case.

IV. Question 1 — Does Using Generative AI Affect Privilege or Work-Product Protection?

Three federal district-court decisions — all from early 2026 — form the entire directly on-point body of law. They are not in flat conflict, but they emphasize different things, and the differences are where the risk lives.

A. *United States v. Heppner* (S.D.N.Y. 2026) — the landmark privilege decision

Facts. A securities-fraud defendant, *without any direction from his lawyers*, used the public Claude platform to prepare roughly thirty-one documents outlining potential defense strategies after receiving a grand-jury subpoena. He later shared them with counsel and claimed privilege and work-product protection. The government moved to establish that the documents were unprotected.

Holding. Judge Jed Rakoff, addressing a question of first impression nationwide, held the documents were protected by *neither* doctrine. The communications failed at least two — and, the court indicated, likely all three — elements of privilege: Claude is not an attorney; the communications were not confidential because Anthropic’s privacy policy states that it collects users’ inputs and outputs, uses them to train the model, and reserves the right to disclose them to third parties “including governmental regulatory authorities”; and the defendant did not use Claude to obtain legal advice (indeed, Claude disclaims giving it). The court invoked black-letter law that non-privileged communications are not “alchemically” transformed into privileged ones merely by later sharing them with counsel.

Work product. Even assuming the documents were prepared in anticipation of litigation, they were not prepared “by or at the behest of counsel.” The defendant acted on his own volition and was not counsel’s agent; the documents did not reflect counsel’s strategy at the time they were created. That was dispositive.

Reach. Binding only in the Southern District of New York, but as the first published decision on the question it carries significant persuasive weight.

B. *Warner v. Gilbarco, Inc.* (E.D. Mich. 2026) — work-product protection survives AI use

Facts. A *pro se* former employee used ChatGPT and similar tools in preparing her discrimination case. The defendants moved to compel “all documents and information” about her AI use, arguing the use waived work-product protection.

Holding. Magistrate Judge Anthony Patti denied the motion as to the AI materials. The materials were prepared in anticipation of litigation and reflected the plaintiff’s internal drafting, analysis, and thought process. On waiver, the court held that ChatGPT and other generative AI programs are “tools, not persons, even if they may have administrators somewhere in the background,” and that work-product waiver requires disclosure to an adversary. Accepting the defendants’ theory “would nullify work-product protection in nearly every modern drafting environment, a result no court has endorsed.” The motion also improperly sought the plaintiff’s mental impressions rather than any existing evidence.

Reach. Binding only in the Eastern District of Michigan, but already cited approvingly elsewhere.

C. *Morgan v. V2X, Inc.* (D. Colo. 2026) — extending *Warner*, and adding a disclosure requirement

Holding. Magistrate Judge Dominguez Braswell followed *Warner*, holding that Rule 26(b)(3) can protect a *pro se* litigant’s mental impressions and litigation-preparation materials even when AI is used, because the rule protects materials prepared by *any party*, and because AI use does not automatically place information in an

adversary's hands absent legal process. The court analogized to Fourth Amendment cases holding that routing information through a third-party system does not forfeit all privacy.

Distinguishing Heppner. The court gave two grounds: *Heppner* was a criminal matter, while *Morgan* was civil and governed by the Federal Rules; and in *Heppner* there was a gap between the party and his attorney, whereas a *pro se* litigant is simultaneously party and advocate, eliminating that gap.

But note two important qualifications. The court ordered the plaintiff to disclose the *name* of the AI tool used with confidential information (that fact alone reveals no mental impressions and lets the opponent assess whether confidential information was compromised). And it amended the protective order to bar inputting confidential information into mainstream AI — including standard ChatGPT, Claude, and Gemini — unless the provider is contractually barred from storing or training on inputs and from disclosing them to third parties, and affords deletion on request.

D. Appellate authority and the emerging split

No federal circuit has resolved the core privilege or work-product questions. The only circuit-level engagement, *Malkeet Lnu v. Blanche* (9th Cir. 2026), concerned sanctions for AI hallucinations, not privilege. *Heppner*, *Warner*, and *Morgan* are all district-court decisions binding only in their own districts.

A workable synthesis emerges. On privilege, there is strong early consensus that communications with public AI platforms are not privileged — the decisive variable is the platform's terms of service. On work product, the consensus is that using AI as a drafting or analytical tool does not by itself waive protection. The live fault line is the threshold question of whether materials a *represented* client prepares *without* counsel's direction can be work product at all: *Heppner* says no; *Warner* and *Morgan* did not reach it because their litigants were *pro se*. The practical reading is that *Heppner*'s work-product holding is limited to represented parties who use AI independently of counsel.

Issue	Emerging position	Key authority
Privilege over public-AI communications	Almost certainly not privileged; terms of service are decisive	Heppner
Work product from AI used as a tool	Not waived merely by using the tool	Warner; Morgan
Represented client's solo AI use, no counsel direction	Likely no work-product protection	Heppner
Pro se litigant's AI use	Work product can apply; tool name may be disclosable	Warner; Morgan
Appellate resolution	None yet; district-level split	Malkeet Lnu (sanctions only)

V. Question 2 — When Are AI Prompts, Conversations, and Outputs Discoverable?

Courts have declined to create a special AI-discovery regime. Discoverability turns on the ordinary framework: relevance and proportionality, privilege, and work product. A federal magistrate judge in the *In re OpenAI* copyright MDL compelled production of AI logs — user prompts and model outputs — under Rule 26(b)(1), confirming there is no AI-specific exemption from ESI discovery.

A. Relevance and proportionality still gate everything

A party cannot obtain an adversary's AI materials just because AI is “in the case.” In *New York Times Co. v. Microsoft Corp.* (S.D.N.Y. 2024), the court refused to compel the Times's use of third-party AI tools and its internal AI positions, because that discovery was not relevant to OpenAI's fair-use defense (which scrutinized OpenAI's conduct, not the Times's). *Warner* likewise found a sweeping demand for a party's AI use disproportionate. The materials must bear on the actual claims and defenses.

B. Counsel-crafted prompts are opinion work product

In *Tremblay v. OpenAI, Inc.* (N.D. Cal. 2024), the court held that ChatGPT prompts crafted by counsel to investigate infringement were *opinion* work product — they reflected counsel's mental impressions about how to interrogate the model — and reversed a magistrate's order classifying them as mere fact work product. *Concord Music Group, Inc. v. Anthropic PBC* (N.D. Cal. 2025) reaffirmed that prompts and related settings used in counsel's AI investigation are opinion work product.

C. Merely using AI does not waive protection

Consistent with *Warner* and *Morgan*, submitting information to an AI tool is not a disclosure to an adversary and does not, by itself, waive work-product protection. The tool-not-person principle carries directly into the discovery analysis.

D. But “at-issue” waiver is the real trap

A party that places its AI investigation at issue can be compelled to produce all of it — including unfavorable results. In *Concord Music*, the publishers announced they would call witnesses to testify that their prompts were the kind “anyone could do” and produced infringing outputs. That put the *post-suit* investigation at issue, and the court ordered production of *all* prompts and outputs — negative as well as positive — used by investigators and client-witnesses, including attorney-drafted prompts handed to those witnesses. It withheld only counsel's own prompts that were never given to a testifying witness, and it found no waiver as to the *pre-suit* investigation. *Tremblay* applied the same limiting principle: disclosing certain prompts and outputs in the complaint waived protection *only* as to those materials, not the entire testing process. The lesson: waiver is tailored to what you actually put in play — so control what you put in play.

E. Privilege in the discovery posture

The privilege analysis from *Heppner* (Part IV) applies with equal force here: a represented party's independent use of a public platform to analyze its situation risks losing both privilege and work-product protection, and non-privileged AI materials do not become privileged by being forwarded to counsel.

F. Protective orders are being rewritten for AI

Courts increasingly build AI restrictions into protective orders so that confidential discovery material is not fed into open models. *Morgan v. V2X* (D. Colo. 2026) barred inputting confidential information into any AI platform unless the provider is contractually prohibited from storing or training on inputs and from disclosing them, with deletion on request. *Jeffries v. Harcross Chemicals Inc.* (D. Kan. 2026) limited AI use in discovery to closed models, reasoning that open models risk data-privacy breaches and incentivize under-production or over-redaction. Expect AI-specific protective-order clauses to become standard in complex litigation.

G. Disclosure obligations and sanctions

Separate from discovery, a growing set of standing orders requires disclosure or certification of AI use in filings, grounded in Rule 11's certification duty. The Ninth Circuit's *Malkeet Lnu v. Blanche* held that a competent, diligent attorney must do more than confirm that AI-provided citations are real — the lawyer “must read and reason.” Courts have imposed real sanctions: public reprimand, disqualification, and bar referral in *Johnson v. Dunn* (N.D. Ala. 2025); strong deterrent sanctions in *ByoPlanet Int'l, LLC v. Johansson* (S.D. Fla. 2025). And *White v. Walmart, Inc.* (S.D. Ind. 2026) held that exclusive reliance on AI-generated discovery responses does not satisfy an attorney's duty to meaningfully meet and confer.

Preservation note. Because prompts, chats, and outputs are ESI, the duty to preserve under Rule 37(e) can attach once litigation is reasonably anticipated. Litigation-hold notices and legal-hold tooling should expressly capture AI tools and accounts — including personal and browser-based tools employees may use.

VI. Question 3 — Court, Judicial, and Ethics Guidance on Lawyers' Use of AI

U.S. guidance does not prohibit lawyers from using generative AI. Instead, it treats AI as subject to existing professional duties, while a growing number of courts layer on their own disclosure, certification, or (in a minority) prohibition requirements.

A. Ethics and bar-association guidance

ABA Formal Opinion 512 (July 2024) is the central national reference. It permits AI use provided the lawyer honors six duties: *competence* (understand how the tool works — its training, privacy settings, limitations, and hallucination risk); *confidentiality*; *communication* with the client; *candor* to the tribunal and meritorious-claims duties; *supervision* under Rules 5.1 and 5.3; and *reasonable fees*. State and local authorities — including Pennsylvania and Philadelphia, Texas, California, Florida, New Jersey, New York State, and New York City — largely follow this model, and some have amended their rules or comments to address AI expressly.

Confidentiality is the sharpest point. The ethical duty of confidentiality is *broader* than the attorney-client privilege and the work-product doctrine, so a disclosure can be an ethics violation *even if* no evidentiary protection has yet been lost. Lawyers are cautioned not to input sensitive or identifying information into a tool until they have confirmed adequate safeguards and a legitimate need.

B. Court rules, standing orders, and certification requirements

There is no uniform national court rule. Approaches range along a spectrum: silence; reminders that existing rules (such as Rule 11) already govern AI-assisted filings; affirmative disclosure or certification requirements; and, in a minority of courts, outright prohibitions. Certification requirements typically ask counsel to confirm either that no AI was used or that every AI-generated citation, quotation, factual assertion, and legal proposition was independently verified before filing; some also require identifying the tool used or confirming no confidential information was disclosed. Illinois offers a notable contrast — its Supreme Court AI policy emphasizes safe use and confidentiality but states that pleadings need not disclose AI use.

Because these rules vary by district and even by individual judge, counsel must check the applicable local rules, standing orders, and the assigned judge's individual rules before every filing. Failure to comply can result in stricken filings, sanctions, and parallel ethics exposure for disobeying a tribunal's rules. Appendix B provides a snapshot of representative requirements.

C. How the guidance ties back to privilege and work product

The ethics guidance and the case law point the same direction. Entering protected information into an open or poorly controlled tool can expose it to the vendor or other users, supporting a privilege-waiver argument and creating data-security risk; and placing AI prompts or outputs at issue can create at-issue waiver. The recommended safeguards mirror the case law: use closed tools with contractual protection, negotiate vendor confidentiality and deletion terms, consider a Federal Rule of Evidence 502(d) order, and add AI-specific confidentiality language to protective orders and discovery agreements.

VII. Recommended Best Practices

The following consolidates the practices the courts and ethics authorities converge on. They divide into three audiences.

A. For the organization (governance and risk)

- **Adopt a written AI-use policy** that identifies approved tools, prohibits inputting sensitive, confidential, or privileged information into unapproved (open) tools, and assigns ownership for updates as vendor terms evolve.
- **Provide mandatory training** on capabilities, limitations, hallucination risk, and confidentiality — at onboarding and at least annually.
- **Default to approved closed/enterprise systems** for any work touching client or matter information.
- **Address your people's and clients' own use of AI**, which can independently compromise confidentiality, privilege, and work product. Counsel clients — early — not to analyze a dispute in a public AI tool.
- **Cover AI in litigation holds.** Ensure preservation obligations reach AI prompts, chats, outputs, and logs across enterprise and personal tools.

B. For counsel (in-house and outside)

- **Direct and document AI use in litigation preparation.** Work-product protection turns on counsel's direction and preparation in anticipation of litigation — so create that record.
- **Choose tools by their contract terms.** Require, in writing: no retention or training on inputs; no third-party disclosure except as essential to deliver the service; and deletion of matter data on request. This is the baseline the *Morgan* and *Jeffries* protective orders now demand.
- **Verify every output.** Independently confirm all citations, quotations, record cites, and factual and legal assertions with traditional tools before use or filing — “read and reason,” do not merely check that a case exists.
- **Obtain informed client consent where required or prudent.** Per ABA Opinion 512, boilerplate engagement-letter language alone is not sufficient when consent is needed to input client information.
- **Check the court and the judge before filing.** Review local rules, standing orders, and individual judges' rules for disclosure, certification, or prohibition requirements, and comply precisely.
- **Engineer against waiver.** Negotiate AI-specific confidentiality provisions in protective orders and discovery agreements, and seek a Rule 502(d) order at the outset to blunt inadvertent-waiver risk.
- **Supervise AI like a nonlawyer assistant** under Rules 5.1 and 5.3 — including vetting vendors and reviewing AI-assisted work product.

C. On protecting privilege and work product specifically

- **Treat public-tool inputs as presumptively non-confidential and unprotected.** Assume *Heppner's* reasoning could apply.
- **Keep represented clients' solo AI use out of the file.** A represented client's independent AI analysis, undirected by counsel, is the fact pattern most likely to lose both protections.
- **Be deliberate about putting AI use at issue.** Relying on AI outputs to prove a point can compel production of the entire investigation, including negative results.

A note on this memorandum's own subject. Consistent with the guidance above, the authorities summarized here should be independently verified against primary sources before being cited in any filing or relied upon for a specific matter. This memorandum is general information, not legal advice, and does not create an attorney-client relationship. For advice on a particular situation, please contact us.

VIII. How Clarion AI Law Can Help

Clarion AI Law works at the intersection of law, technology, and AI governance. We help organizations put the safeguards in this memorandum into practice — drafting and operationalizing AI-use policies, selecting and papering enterprise AI tools with defensible confidentiality terms, building AI into litigation-hold and discovery protocols, negotiating AI-specific protective-order and Rule 502(d) provisions, and training legal and business teams to use these tools competently and defensibly. If you would like to assess your organization's exposure or stand up a practical AI-governance program, we would welcome the conversation.

Appendix A — Case Digests

The following digests summarize the decisions discussed above, drawn from the underlying research. Citations should be independently verified against primary sources before use.

Privilege and work product

United States v. Heppner 820 F. Supp. 3d 292 (S.D.N.Y. Feb. 17, 2026); 2026 WL 436479

First-impression ruling that a client's solo use of a public AI platform is neither privileged nor work product.

A represented criminal defendant used Claude, without counsel's direction, to prepare defense-strategy documents and later shared them with counsel. Held: no attorney-client privilege (the AI is not an attorney; the platform's privacy policy defeated confidentiality; the defendant did not seek legal advice from the AI) and no work-product protection (the materials were not prepared by or at the behest of counsel). Non-privileged materials do not become privileged when later shared with counsel.

Warner v. Gilbarco, Inc. 820 F. Supp. 3d 629 (E.D. Mich. Feb. 10, 2026); 2026 WL 373043

AI-assisted materials are work product; using an AI tool is not disclosure to an adversary.

A pro se plaintiff's AI-assisted litigation materials were protected work product prepared in anticipation of litigation. Generative AI programs are "tools, not persons," so using them is not the kind of third-party disclosure that waives protection; the contrary theory "would nullify work-product protection in nearly every modern drafting environment."

Morgan v. V2X, Inc. 2026 WL 864223 (D. Colo. Mar. 30, 2026)

Extends Warner to pro se litigants; requires disclosure of the AI tool's name; amends the protective order.

Rule 26(b)(3) can protect a pro se litigant's mental impressions even when AI is used, and AI use does not automatically waive work product. The court distinguished Heppner (criminal; represented party acting apart from counsel), required the plaintiff to disclose the name of the AI tool used with confidential information, and amended the protective order to bar inputting confidential information into mainstream AI absent contractual no-training, no-disclosure, and deletion protections.

Discoverability, waiver, and protective orders

Tremblay v. OpenAI, Inc. 2024 WL 3748003 (N.D. Cal. Aug. 8, 2024)

Counsel-crafted AI prompts are opinion work product; waiver is limited to what was disclosed.

Reversing a magistrate, the court held ChatGPT prompts crafted by counsel to investigate infringement were opinion work product reflecting counsel's mental impressions. Waiver extended only to the specific prompts, outputs, and settings disclosed in the complaint — not to negative test results or the broader testing process.

Concord Music Group, Inc. v. Anthropic PBC 2025 WL 3677935 (N.D. Cal. Dec. 18, 2025)

At-issue waiver compels production of an entire AI investigation, including negative results.

Prompts and settings used in the publishers' investigation were opinion work product. But by announcing they would call witnesses about the post-suit investigation's results, the publishers placed it at issue and had to produce all prompts and outputs (positive and negative) used by investigators and client-witnesses — including attorney-drafted prompts given to those witnesses. Counsel's own prompts never given to a witness, and the pre-suit investigation, remained protected.

New York Times Co. v. Microsoft Corp. 757 F. Supp. 3d 594 (S.D.N.Y. Nov. 22, 2024)

A party's AI use is not discoverable merely because AI is in the case.

The court denied OpenAI's motion to compel the Times's use of third-party AI tools, its own AI products, and its internal AI positions, resting its denial on relevance — the fair-use defense requires scrutiny of OpenAI's conduct, not the Times's. Relevance is measured against the actual claims and defenses.

In re OpenAI, Inc., Copyright Infringement Litigation 2025 WL 1652110 (S.D.N.Y. May 30, 2025); MDL No. 25-MD-3143 (S.D.N.Y. Dec. 2, 2025)

No AI-specific exemption from ESI discovery; relevance still controls.

A magistrate judge compelled production of AI logs — user prompts and model outputs — under Rule 26(b)(1), confirming AI materials are ordinary ESI. In a related ruling the court declined to compel the Times’s AI-tool logs where not relevant, and found no Rule 37(e) spoliation implicated.

Jeffries v. Harcros Chemicals Inc. 2026 WL 820218 (D. Kan. Mar. 25, 2026)

Protective order limited AI use in discovery to closed models.

Allowing documents to be uploaded to open AI models would risk data-privacy breaches and incentivize under-production or over-redaction; the court limited discovery-related AI use to closed models.

White v. Walmart, Inc. 2026 WL 1002132 (S.D. Ind. Apr. 14, 2026)

Attorneys cannot outsource discovery judgment to AI.

Exclusive reliance on AI-generated discovery responses does not satisfy the duty to meaningfully meet and confer; attorneys must exercise independent judgment and oversight even when using AI.

Competence, disclosure, and sanctions

Malkeet Lnu v. Blanche 177 F.4th 1014 (9th Cir. June 3, 2026)

First circuit-level standard for attorney competence when using AI.

AI use is not inherently unethical, but an attorney must rigorously check its output for both fabrications (nonexistent cases) and inaccuracies (real cases cited for wrong propositions). A competent, diligent attorney must “read and reason,” not merely verify that citations exist. The court did not address privilege or work product.

Johnson v. Dunn 792 F. Supp. 3d 1241 (N.D. Ala. July 23, 2025)

Sanctions for filing AI-hallucinated citations.

The court imposed public reprimand, disqualification from the case, and referral to the state bar against attorneys who included AI-hallucinated citations in discovery motions without verification, finding recklessness tantamount to bad faith.

ByoPlanet Int’l, LLC v. Johansson 792 F. Supp. 3d 1341 (S.D. Fla. July 17, 2025)

Strong sanctions to deter blind reliance on AI.

Wholesale reliance on AI without further inquiry or diligence is conduct courts should deter; strong sanctions will discourage lawyers from blindly relying on AI for court submissions.

Willis v. U.S. Bank N.A. (Igloo Series Trust) 783 F. Supp. 3d 959 (N.D. Tex. May 15, 2025)

Illustrates court disclosure regimes.

Addresses the Northern District of Texas Local Civil Rule 7.2(f) (disclose AI use on a brief’s first page under a designated heading) and the Southern District of Texas General Order 2025-04 (cautioning against filing AI-drafted pleadings without checking accuracy).

Hunter v. TForce Freight Inc. 2026 WL 673285 (D. Ariz. Mar. 10, 2026)

Declaration-and-certification model.

Applies the District of Arizona’s policy requiring any party using AI in a filing to attach a declaration disclosing the use, certifying personal review and verification under Rule 11, and identifying the AI-generated portions.

Torres v. Freedom Mortgage Corp. 2026 WL 1615767 (M.D. Pa. June 4, 2026)

Certificate identifying AI portions and verifying citations.

Reminds parties of the Middle District of Pennsylvania’s requirement that AI-assisted documents include a certificate identifying AI-generated portions and certifying that citations were verified as accurate.

Williamson v. TransUnion LLC 2026 WL 1262027 (M.D. Fla. May 8, 2026)

Standing-order certification.

References the Middle District of Florida’s standing order requiring a certificate on the use of generative AI in court filings.

Appendix B — Snapshot of Representative Court AI-Disclosure Requirements

Requirements vary by district and by individual judge and change frequently. This snapshot is illustrative only; always confirm the current local rules, standing orders, and the assigned judge's individual rules before filing.

Court / Authority	Nature of requirement
N.D. Tex. — Local Civ. R. 7.2(f)	Brief prepared using generative AI must disclose that fact on the first page under a designated heading; identify AI-prepared portions if the judge directs.
S.D. Tex. — General Order 2025-04	Cautions attorneys and self-represented litigants against filing AI-drafted pleadings without checking accuracy.
D. Ariz. — general policy (per Hunter)	Attach a declaration disclosing AI use, certifying independent review and verification under Rule 11, and identifying AI-generated portions.
M.D. Pa. (per Torres)	Certificate identifying AI-generated portions and certifying that citations were verified as accurate.
M.D. Fla. (per Williamson)	Standing order requiring a certificate on the use of generative AI in filings.
E.D. Pa. — individual judges (e.g., Baylson)	Disclose any AI use in preparing a filing and certify that AI output was verified.
C.D. Cal. / S.D.N.Y. / E.D. Pa.	Representative districts with standing orders or certification requirements addressing AI use.
E.D. Tex.	Amended attorney-practice standards to add provisions on generative-AI use.
Illinois (Supreme Court AI policy)	Emphasizes safe use and confidentiality; states pleadings need not disclose AI use.
Bankruptcy courts (e.g., N.D. Tex., W.D. Okla.)	Range from prohibition to disclosure or certification; definitions of “AI” vary and some exclude established research tools.

Appendix C — Key Authorities and Source Index

This memorandum synthesizes three CoCounsel/Westlaw research reports and their underlying sources. The principal authorities are listed below.

Rules and statutes

- Fed. R. Civ. P. 26 — scope of discovery, relevance and proportionality (26(b)(1)); work-product doctrine (26(b)(3)).
- Fed. R. Civ. P. 34 — production of documents and electronically stored information.
- Fed. R. Civ. P. 11 — certification duty for filings; basis for AI-verification obligations.
- Fed. R. Civ. P. 37(e) — failure to preserve electronically stored information.
- Fed. R. Evid. 501 — federal common law of privilege.
- Fed. R. Evid. 502 — scope of and limits on waiver, including inadvertent disclosure and 502(d) orders.

Ethics and professional-responsibility guidance

- ABA Standing Committee on Ethics and Professional Responsibility, Formal Opinion 512 (July 29, 2024) — generative AI and the duties of competence, confidentiality, communication, candor, supervision (Rules 5.1/5.3), and reasonable fees.
- State and local guidance, including Pennsylvania and Philadelphia, Texas, California, Florida, New Jersey, New York State, and New York City; and the Illinois Supreme Court AI policy.

Secondary sources and commentary

- eDiscovery for Corporate Counsel §§ 15:19 (feeding prompts into LLMs), 15:20 (AI notetaking/transcription), 26:15 (confidentiality, privilege, and work product).
- Am. Jur. Proof of Facts 3d (Proof of Waiver of Attorney-Client Privilege; Proof of Basis for Work-Product Protection); A.L.R. Fed. 3d (Waiver of Federal Work-Product Protection).
- Practical Law practice notes: Generative AI in Litigation: Overview; Generative AI Ethics for Litigators; Generative AI: Preserving the Attorney-Client Privilege and Work Product Protection; Work Product Protection: Waiver (Federal); Certification Regarding Generative AI for Court Filings; and related standard clauses and policies.
- JD Supra commentary (2026): analyses by Zelle LLP; Akin Gump; Krieg DeVault; Kilpatrick; and Esquire Deposition Solutions on privilege, work product, protective orders, and litigation holds.

Prepared by Clarion AI Law. © 2026 Clarion AI Law. This memorandum is provided for general informational purposes, reflects research as of mid-2026, is not legal advice, and does not create an attorney-client relationship. The authorities cited are drawn from AI-assisted legal research and should be independently verified before being relied upon or cited.