

THURSDAY, JUNE 18, 2026

PERSPECTIVE | MCLE

Privilege survives AI only if lawyers do their job

Heppner fired the warning shot. *Morgan* drew the map. *Conservation Law* marked the boundaries: the emerging law of AI, privilege and work product.

By Arnold P. Peter

We are witnessing the birth of an entirely new body of law governing artificial intelligence, attorney-client privilege, confidentiality and attorney work product. As generative AI tools migrate from novelty to daily fixture in legal practice, courts are being asked to decide questions that the existing doctrine never anticipated. Two recent federal decisions, in particular, have begun to shape how courts will approach these emerging questions, and together they offer the first meaningful judicial guidance on a subject that will affect virtually every practicing lawyer.

Heppner: The warning shot

Most lawyers have now heard about *United States v. Heppner*, 820 F. Supp. 3d 292 (S.D.N.Y. 2026), a Southern District of New York criminal case that addressed a question of first impression: whether materials generated through interactions with a public AI platform are protected by the attorney-client privilege or work-product doctrine. District Judge Jed Rakoff largely answered that question in the negative under the facts presented.

The defendant in *Heppner* used Anthropic's Claude platform to analyze legal issues, develop potential defenses and prepare materials related to a federal criminal investigation. The court concluded that those communications were neither privileged nor protected work



This art was created with the assistance of Shutterstock AI tools

product because they had been voluntarily disclosed to a third-party AI provider and were not created at the direction of counsel. Critical to the court's reasoning was the fact that the defendant acted alone, outside the supervision of any attorney, and accepted terms of service that placed no meaningful restrictions on how his inputs might be retained, reviewed or used. Under those circumstances, Judge Rakoff found it difficult to locate any reasonable expectation of confidentiality that the law could protect.

At first glance, *Heppner* appeared to send a warning shot across the bow of lawyers and clients using generative AI. Many commentators

interpreted the decision as standing for the proposition that entering information into an AI system may destroy otherwise available protections.

Just weeks later, however, another federal court provided what may ultimately become the more influential analysis. That decision approached the same underlying issues from a markedly different angle, with consequences that may prove far-reaching.

Morgan: Reframing the inquiry around safeguards

In *Morgan v. V2X, Inc.*, No. 25-CV-01991-SKC-MDB, 2026 WL 864223 (D. Colo. Mar. 30, 2026), Magistrate Judge Maritza Dominguez Braswell, confronted many of the same con-

cerns discussed in *Heppner* but reached a more nuanced conclusion. Rather than viewing AI itself as the problem, the *Morgan* court focused on the circumstances under which AI is used and the safeguards surrounding its use.

In many respects, *Morgan* distinguished itself from, rather than disagreed with, *Heppner*. The two opinions addressed fundamentally different factual settings, which largely accounts for their divergent conclusions.

The *Heppner* court focused on a criminal defendant's use of a public-facing AI platform without attorney direction and under terms that arguably undermined any reasonable expectation of confidentiality. *Morgan*, by contrast, addressed AI use within the context of ongoing litigation, discovery obligations, protective orders and attorney supervision.

Most importantly, *Morgan* rejected any suggestion that the mere involvement of AI automatically destroys work-product protection. The court recognized that lawyers routinely use technologies operated by third parties, including email providers, cloud-storage vendors, e-discovery platforms, litigation-support systems and document-management tools. AI, the court suggested, should not necessarily be treated differently simply because it is new.

That distinction is significant because it shifts the focus away from the existence of AI and toward the management of risk. In doing so, it reorients the inquiry around the user's conduct rather than the underlying technology.

Rather than asking whether AI was used, *Morgan* asks whether the user maintained reasonable confidentiality protections. Did the platform train on user data? Were contractual safeguards in place? Could confidential information be deleted? Were third-party disclosures restricted? Was the AI being used in a manner consistent with existing protective orders and professional obligations? Each of these questions maps onto a concrete, verifiable feature of how a given tool is deployed, which means that the privilege analysis becomes a function of diligence rather than of the technology's mere presence. A lawyer who selects an enterprise platform with contractual data protections, disables model training and operates within the bounds of a protective order is in a materially different position from a defendant typing sensitive facts into a consumer chatbot.

The court's answer was that those questions matter far more than the simple fact that artificial intelligence was involved. The presence of adequate safeguards, not the involvement of AI, drives the analysis.

Morgan therefore reframed the debate. Instead of treating AI as a privilege-destroying technology, the court treated AI as another technological tool whose use must be governed by reasonable safeguards.

This distinction also highlights one of the potential limitations of broader readings of *Heppner*. The fact that AI is not a lawyer does not necessarily answer the privilege question. Courts have long recognized that privilege may survive the involvement of interpreters, consultants, forensic accountants, litigation-support vendors, jury consultants and other third parties whose participation facilitates legal representation. Future courts will likely confront whether certain AI systems should be analyzed similarly.

The analogy to other third-party service providers is more than rhetorical. When a law firm transmits privileged communications through an email provider or stores work product on a commercial cloud server, it has technically disclosed that information to an outside vendor. Privilege survives because the disclosure is consistent with maintaining confidentiality and because the

vendor is engaged to facilitate the representation rather than to share the information with adversaries. *Morgan* suggests that, properly configured, AI tools can occupy the same category, while *Heppner* illustrates what happens when none of those protective conditions are present.

Morgan may ultimately become the more influential decision. *Heppner* is a criminal case addressing privilege under highly specific facts. *Morgan* is a civil discovery decision authored by a judge who serves on judicial committees examining artificial intelligence and its impact on litigation and court administration.

Unlike *Heppner*, which primarily resolved the dispute before it, *Morgan* expressly attempted to create a workable framework for future cases. Several commentators have already described *Morgan* as the first judicial opinion to provide a practical blueprint for AI-specific protective orders, confidentiality protocols and litigation governance.

Conservation Law: From confidentiality to transparency

A third decision, handed down only weeks after *Morgan*, illustrates how rapidly this body of law is developing and adds a dimension that neither *Heppner* nor *Morgan* squarely addressed. In *Conservation Law Foundation, Inc. v. Shell Oil Co.*, No. 3:21-CV-00933 (VDO), 2026 WL 764396 (D. Conn. May 18, 2026), the court considered not whether AI use waived privilege, but whether the AI prompts an expert used to analyze documents were themselves discoverable. The court held that they were. Because the expert had used those prompts to narrow and shape her analysis, the court reasoned, the prompts were part of her methodology and therefore discoverable under Federal Rule of Civil Procedure 26(b).

Conservation Law is significant precisely because it shifts the analysis from confidentiality to transparency. Where *Heppner* and *Morgan* ask whether AI use preserves or forfeits privilege, *Conservation Law* reminds practitioners that the inputs supplied to an AI system can become part of the discoverable record when they inform an expert's reasoning. The prompts an expert chooses, such as her assumptions,

instructions and analytical choices, may reveal how she reached her conclusions and litigants should assume that such prompts may have to be produced. Taken together with *Heppner* and *Morgan*, the decision underscores that the manner of AI use, not the mere fact of it, will determine its legal consequences.

Reconciling the decisions:

A coherent framework

If I had to predict where appellate courts ultimately land, I believe these decisions will survive because they are not truly inconsistent. Each addresses a distinct set of facts and, properly understood, the opinions complement rather than contradict one another. *Heppner* and *Morgan* map the boundaries of privilege and confidentiality, while *Conservation Law* addresses the discoverability of the inputs themselves, together forming a coherent picture rather than competing rules.

Heppner stands for the proposition that careless use of public AI platforms may jeopardize privilege and confidentiality protections. *Morgan* stands for the proposition that responsible and properly governed AI use does not automatically destroy them.

Viewed together, the cases suggest that future courts will focus less on whether AI was used and more on how it was used. *Conservation Law* adds a further dimension to that inquiry, cautioning that the manner of use may not only determine whether privilege is preserved but also whether the AI inputs themselves become part of the discoverable record. That combined framework is likely the one that will govern legal ethics, discovery practice, privilege law and litigation strategy for years to come.

Practical lessons for practitioners

For practitioners, the practical lessons are already clear. Before deploying AI on any matter involving confidential or privileged information, lawyers should scrutinize the platform's terms of service, confirm that user inputs are not used to train models, and prefer enterprise agreements that contractually restrict data retention and third-party access. Counsel should supervise AI use rather than leaving it to clients

acting on their own, document the safeguards in place, and ensure that any AI workflow aligns with applicable protective orders and ethical obligations. In light of *Conservation Law*, practitioners should also assume that the prompts and instructions supplied to an AI system, particularly those informing an expert's analysis, may themselves be discoverable, and should craft and preserve those inputs accordingly. Firms would also be well advised to adopt internal AI-use policies now, before a privilege or discovery dispute forces the issue into litigation.

The broader significance of *Heppner*, *Morgan*, and *Conservation Law* lies not in the specific outcomes but in the analytical framework they begin to establish. As more courts confront these questions, the decisions that endure will likely be those that treat AI as a powerful but governable tool rather than as a categorical threat to confidentiality, while recognizing that the inputs supplied to that tool may carry their own discovery consequences. Lawyers who internalize those distinctions, and who build disciplined safeguards into their use of these technologies, will be far better positioned to protect their clients than those who either avoid AI entirely or use it without care.

Arnold P. Peter is the founder and managing partner of Peter Law Group, a California-based law firm representing employees, executives, entrepreneurs, and businesses in transactions and complex litigation, employment, entertainment and business matters.

