SIDLEY

ARTIFICIAL INTELLIGENCE UPDATE

Generative AI in Discovery: Protective Orders as an Emerging Point of Dispute

APRIL 6, 2026

As courts have begun addressing generative AI in the privilege and work product context, they are also confronting related disputes in the context of protective orders. Recent decisions *Morgan v. V2X, Inc*. and *Jeffries v. Harcros Chemicals, Inc.* show that disagreements about how protective orders should address the use of AI in discovery — issues previously handled through negotiation — now will be informed by guidance from the courts.

In February 2026, disputes concerning generative AI and privilege protections took center stage, with courts addressing matters of first impression concerning the application of the attorney-client privilege and work product doctrine to AI use in discovery. In *United States v. Heppner*, the court declined to extend privilege or work product protection to AI-generated materials created outside the direction of counsel and not for the purpose of obtaining legal advice. In *Warner v. Gilbarco*, the court held that a *pro se* litigant's AI-assisted materials were protected work product where they reflected the plaintiff's own mental impressions and that use of a public AI tool did not, by itself, constitute waiver.

Against that backdrop, *Morgan* and *Jeffries* address a related set of questions through a different lens: how courts should manage the use of AI in discovery through protective order provisions. Together, these decisions offer early insight into how courts are evaluating confidentiality risks associated with AI and the role protective orders may play in addressing them. They also raise important questions about *pro se* access to the courts and how increased use of AI in litigation may both expand access for unrepresented litigants while also introducing new challenges for, and burdens on, parties and courts.

I. Two Courts Recently Resolved Disputes Regarding the Limits of AI Use in Discovery Protective Orders

Morgan v. V2X, Inc.

In *Morgan*, the United States District Court for the District of Colorado approached the use of generative AI in discovery through the lens of confidentiality risks and how they may be mitigated through a protective order. Although the parties agreed that the existing confidentiality order should be amended to address the submission of confidential material to an AI tool, they disputed the appropriate scope of restrictions. The court focused on the risks associated with submitting discovery materials to AI tools, particularly where those tools may retain, use, or disseminate data, including for model training.

To address confidentiality concerns, the court required that any AI tool used to process confidential information be subject to contractual safeguards, including prohibitions on using inputs to train models, restrictions on onward disclosure, and the ability to delete data. As a practical matter, the court recognized that these requirements would limit the use of most widely available, consumer-facing AI tools for confidential discovery materials — at least as of the date of the decision. In expressly referring to AI as "one of the most powerful knowledge tools ever to become available to the masses," the court distinguished between AI tools operating "in a secure, closed-circuit environment" and less secure, "mainstream low-to-no-cost AI."

The court also addressed work product protections in this context. Consistent with *Warner*, it confirmed that Rule 26(b)(3) can protect a *pro se* party's mental impressions and litigation materials generated with the assistance of AI and that use of such tools does not automatically waive those protections.

The court's analysis also includes a broader discussion of how AI systems may be conceptualized within existing legal frameworks, including analogies to third-party service providers — an issue we explore further below.

Jeffries v. Harcros Chemicals, Inc.

In *Jeffries*, the United States District Court for the District of Kansas took a more expansive approach, granting a motion to amend a protective order to restrict the use of "open" or publicly accessible AI tools (as opposed to "closed" or "secure" AI tools) not only for confidential information but also for all discovery materials, including specifically those that are not confidential. The court, like the court in *Morgan*, drew a distinction between publicly accessible AI systems, which may retain and use submitted data to train models, and more secure or closed systems that operate in a more controlled environment. It concluded that the use of such publicly accessible AI tools presents unique risks, including the practical inability to claw back or delete data once it has been incorporated into a model.

The court also emphasized the impact on the discovery process itself. Allowing parties to

upload materials into such tools, even if they are not designated confidential, could incentivize more conservative behavior in discovery, including underproduction or excessive redaction, as parties seek to avoid exposing large volumes of data to systems outside their control. By restricting the use of these tools while permitting more secure alternatives, the court sought to facilitate broader production and reduce friction in discovery for producing parties.

II. Practical Takeaways

Much like the risk management takeaways outlined in our **prior article** on the *Heppner* and *Warner* decisions, the following guidance emerges from this next wave of recent case law addressing the use of generative AI in discovery.

Address AI use explicitly at the outset of discovery

Morgan and *Jeffries* make clear that protective orders serve as a key vehicle for managing AI-related risks in discovery. Parties should consider addressing these issues directly in protective order negotiations before discovery begins, particularly given the rapidly developing case law in this area, rather than leaving them to informal agreement or later dispute. Possible topics of inquiry include identifying the required confidentiality and security safeguards the AI tools must have in order to process discovery materials. Parties should be prepared to raise practical, real-world concerns regarding AI use proactively with both courts and counterparties. Early alignment may reduce motion practice and avoid disruption once discovery is underway.

Define permissible AI use through concrete safeguards that stand the test of time

Morgan and *Jeffries* reflect a move away from general prohibitions toward more specific, operational requirements governing the use of generative AI in discovery. Parties should consider whether proposed protective order language addresses issues such as data retention, use of inputs for model training, dissemination, and the ability to delete or segregate data. Framing restrictions in terms of these concrete safeguards — rather than by reference to particular tools or their current capabilities—can make provisions more durable and easier to apply as technologies evolve. This approach also underscores the broader importance of using closed AI systems, for confidentiality, privacy, and privilege reasons, rather than exposing sensitive data to open, commercial AI tools.

III. Emerging Issues

These two recent decisions highlight several questions that are likely to shape how courts approach generative AI in discovery going forward. Although *Morgan* and *Jeffries* provide practical guidance for structuring protective orders, they also offer early insight into how courts are conceptualizing AI systems and the risks those systems present.

Conceptualizing AI systems within existing legal frameworks

A central issue emerging from these cases is how AI systems should be understood within existing legal doctrines — whether as tools used by a party, intermediaries that facilitate analysis, or third-party platforms that may affect confidentiality and waive privilege. In *Morgan*, drawing on a line of Fourth Amendment cases involving email providers and other third-party systems, the court acknowledged that almost all electronic interaction passes through third-party systems and posed a rhetorical question: "Does that mean that anyone with a Gmail account has forfeited all rights to confidentiality and privacy?" The court's analysis suggests that the answer is no, reinforcing that the use of third-party systems, without more, does not automatically defeat expectations of confidentiality. But the rationale of these two cases also strongly suggests that courts will not take a binary approach to confidentiality, but rather look at the actual context of technology, including not only legal or contractual limitations, but also the practical realities of how third-party systems may access and use enormous amounts of data.

The role of user behavior and AI system design

The court in *Morgan* addressed an issue that other courts have not explored as deeply to date: how modern AI systems differ from more traditional technologies in terms of inducing user behavior. It noted that generative AI is not like a search engine, which passively returns results, but instead consists of systems "designed and trained to engage," inviting "candid and significant disclosure of information, including sensitive information." These systems "simulate empathy, foster trust, and interact in a way that feels genuine and intimate," and the court emphasized that research suggests that users often share sensitive information without fully understanding how it may be retained or used.

These observations highlight how courts may incorporate assumptions about user behavior and system design into future analyses of confidentiality, waiver, and privilege. How courts continue to characterize AI systems along this spectrum and account for these behavioral dynamics likely will have significant implications for privilege, work product, and confidentiality doctrines. So in this respect, legal doctrine may be driven not only from technological evolutions, but also by parallel (but not identical) evolutions in how people actually use and think about that technology.

AI use by unrepresented litigants and its implications for discovery

Each case also raises questions about access to justice for unrepresented litigants and their use of AI tools. As these tools become more sophisticated, the volume of more polished complaints and other legal filings may increase, potentially improving access to the courts, particularly for *pro se* litigants. At the same time, these developments raise questions regarding how information used in those tools, particularly discovery materials, is handled.

Courts and parties will need to navigate these considerations carefully. Making it easier for *pro se* plaintiffs to draft complaints, for example, may well make it easier for those complaints to survive motions to dismiss. This may make it easier for courts (and defendants) to reckon with meritorious claims that previously were obscured; it also may make it easier for claims that lack merit to impose additional costs on courts (and defendants). By contrast, the limits imposed by confidentiality orders may lessen the ability of *pro se* plaintiffs to take full advantage of the power of AI tools, a subject with which the *Morgan* court wrestles.

The use of publicly accessible AI tools also have the potential to impose costs on parties whose information is disclosed in litigation, even if not central to the claims asserted. As *Jeffries* underscores, information submitted to such public AI systems may be "impossible as a practical matter" to retrieve or delete fully once incorporated into a model, which can expose large volumes of information in ways that are not easily remedied through traditional discovery mechanisms. And, given the "massive amounts of information in complex litigation," much of which is "largely irrelevant," easier public digestion or use of that information by AI raises additional, serious questions.

In this context, protective orders may serve as one mechanism to manage these risks by placing early guardrails on the use of AI in connection with discovery materials, particularly where those materials include sensitive or confidential information. Because these AI tools are so new and their downstream effects remain matters of first impression for many courts, counsel for producing parties must plainly articulate these risks — particularly the potential exposure of large volumes of sensitive information — as concrete and significant harms.

* * *

Taken together, these decisions suggest that the application of existing privilege, work product, and confidentiality doctrines to generative AI is increasingly being addressed and at times contested in the context of protective orders. As courts continue to grapple with how to conceptualize AI systems and account for how they are used in practice, those underlying assumptions will shape their analysis. As additional disputes reach the courts, the structure and specificity of protective order provisions, and the actual guardrails put in place with respect to the use of the AI tools themselves, will play an increasingly important role in defining the use of AI in discovery.

professional advisers. Sidley and Sidley Austin refer to Sidley Austin LLP and affiliated partnerships as explained at **www.sidley.com/disclaimer**.

Contacts

If you have any questions regarding this Sidley Update, please contact the Sidley lawyer with whom you usually work, or

PARTNER

David A. Gordon

dgordon@sidley.com

Chicago +1 312 853 7159

PARTNER

Takayuki Ono

tono@sidley.com

**Not a registered foreign lawyer in Japan.*

Chicago +1 312 853 7296
Tokyo +81 3 3218 5096

COUNSEL

Matt S. Jackson

matthew.jackson@sidley.com

Chicago +1 312 853 4101

COUNSEL

Daniel Lim

daniel.lim@sidley.com

Washington, D.C.
+1 202 736 8699

MANAGING ASSOCIATE

Stephen Beemsterboer

sbeemsterboer@sidley.com

Chicago +1 312 853 0750

ASSOCIATE

Kseniya K. Belysheva

kbelysheva@sidley.com

Los Angeles +1 213 896 6028