ALM | LAW.COM

New York Law Journal

TECHNOLOGY TODAY

FEDERAL E-DISCOVERY

Court Extends Protective Order's AI Restrictions to All Discovery Materials

By H. Christopher Boehning and Daniel J. Toal
April 14, 2026

Protective orders have long served as key mechanisms for restricting the dissemination of certain information disclosed by parties in a proceeding. Traditionally, these orders were designed to safeguard confidential material, often by imposing restrictions on how such information may be accessed, shared, or utilized. However, the changing landscape of litigation technology has introduced new challenges—particularly with the integration of artificial intelligence (AI) tools into discovery workflows. While AI can offer substantial efficiencies, its use may also present heightened concerns if not used judiciously. For example, some widely available AI tools are "public" systems that retain the data they process. That data may become part of the system's training data and ultimately could be presented to other users. Uploading discovery materials into these public AI platforms can make it virtually impossible to claw back or delete information, increasing the risk of both unintended disclosure and loss of control over privileged, confidential, or sensitive information.



H. Christopher Boehning and Daniel J. Toal

Modern protective orders increasingly include provisions applicable to all information produced, particularly by establishing minimum standards for safeguarding the security and privacy of document productions and other disclosures—including restrictions on the use of public AI tools.

H. CHRISTOPHER BOEHNING *and* DANIEL J. TOAL *are litigation partners at Paul, Weiss, Rifkind, Wharton & Garrison. Ross M. Gotler, deputy chair and counsel, e-discovery, and Lidia M. Kekis, e-discovery attorney, assisted in the preparation of this article.*

In a recent matter, the parties had put in place a protective order limiting the use of public AI tools, but only with respect to confidential materials. However, in its ruling, the court extended this restriction to all discovery materials, over a party's objection, illustrating the adaptable nature and evolving function of protective orders. Taking into account concerns regarding public AI use, the court determined that confining AI usage to secure, private systems was necessary to sufficiently protect even non-confidential information. This decision exemplifies how protective orders can respond to emerging technological risks, capturing the benefits of AI while preserving comprehensive safeguards for all discovery materials.

Modern protective orders increasingly include provisions applicable to all information produced, particularly by establishing minimum standards for safeguarding the security and privacy of document productions and other disclosures—including restrictions on the use of public AI tools.

Motion to Amend Protective Order

In *Jeffries v. Harcros Chemicals,* 2026 WL 820218 (D. Kan. Mar. 25, 2026), local residents filed putative class actions against the owners and operators of a nearby chemical facility, alleging toxic airborne emissions caused serious health injuries. The court initially entered a protective order that restricted the use of AI tools for confidential information. In particular, the protective order limited use of AI tools to those that are secure and private, restricted any tool training to this suit, and required deletion of all confidential information at the conclusion of the case, including information used for tool training, with access limited to authorized persons. In essence, the protective order's restrictions applied only to confidential information and therefore limited use to secure, private AI tools. As such, the protective order permitted the parties to use public AI tools for nonconfidential information.

The defendants moved to amend the protective order to apply the AI restrictions "more broadly to all 'discovery materials'—i.e., all documents and information produced in discovery." The defendants thus sought to prevent the upload of non-confidential materials "into public or 'open loop' generative artificial intelligence tools," which the existing order would have allowed.

Parties' 'Good Cause' Arguments

To amend a protective order, Federal Rule of Civil Procedure 26(c)(1) requires a party to show good cause. The defendants presented several arguments in an effort to demonstrate good cause, arguing first that AI use in litigation is a "paradigm shift" that could jeopardize the security and integrity of information exchanged in discovery. Second, since public AI tools retain and continually learn from ingested data, it can be "practically impossible to claw back" privileged material, or delete it at the conclusion of the case, thereby creating risks of privilege waiver or exposure of personal information. Third, ingesting discovery materials "wholesale" into public AI tools could violate U.S. privacy laws and the EU General Data Protection Regulation's (GDPR) "strict disclosure rules," impacting two defendants here. Finally, and notably, the amendment was needed "to

protect against exposure of critical infrastructure and data breach, given that Defendants are part of the chemical sector designated as Critical Infrastructure and vital to national security."

The plaintiffs argued the defendants had not shown good cause, including calling the proposal an "umbrella" protective order "disfavored by courts." They also argued that the amendment would unduly increase litigation costs by preventing them from using public AI tools to analyze discovery materials that were not otherwise protected. In addition, they argued that the modification requested by the defendants would violate their First Amendment rights by restricting use and dissemination of non-confidential discovery. Finally, they argued the clawback and deletion concerns were speculative and unsupported by concrete examples.

Public AI Risks Provide 'Good Cause'

In reviewing the motion, the court began by swiftly rejecting the plaintiffs' arguments on the "umbrella" protective order characterization, undue burden imposition, and constitutionality question, finding all such arguments were unfounded or unsupported. In particular, the court found the plaintiffs' undue burden argument not only lacked support, such as quantifiable increased costs, but also was misplaced since the proposal permitted AI use for all parties and, thus, imposed "the same financial burden" to both sides.

The court's primary analysis focused on the identified risks from using non-confidential information in public AI tools.The court found unpersuasive the plaintiffs' argument that the defendants failed to show a "need for the relief sought." The court noted that the plaintiffs' criticisms did not undercut the defendants' showing of the AIspecific risks at issue.

The court found the defendants' arguments on the risks more persuasive. For example, the court emphasized the plaintiffs never squarely addressed a central point on risk: once data is uploaded to a public AI tool, it cannot realistically be clawed back or deleted because the tool continues to learn from it. The plaintiffs instead dismissed those concerns as "speculative," relying on a misplaced publicposting analogy. The court, however, rejected that analogy, pointing out that using public AI tools "presents the opportunity for a centralized repository that makes information available to the public at a scale that was not historically available and ignores the very real security risks of public AI tools, including the inability to effectively claw back information from the AI tool."

The court acknowledged that the plaintiffs identified potentially compelling efficiencies from the use of AI, such as "significantly accelerating document review, quickly summarizing documents, and streamlining privilege review and logging." Even so, the court noted that the proposed protective order amendment still permitted the parties to use AI tools; it just restricted such use to secure, private AI tools. The court reasoned that "wholesale submission of discovery materials to an open AI Tool for these e-discovery tasks could expose massive amounts of data" to various risks such as violating "U.S. data privacy laws and the stricter GDPR disclosure rules."

The court, on balance, found the defendants had sufficiently demonstrated good cause. The court highlighted the risks "in this particular case"—where discovery comes from a critical-

infrastructure industry and disclosure could implicate data-privacy laws—as sufficient to justify an amendment to the protective order. Unconvinced by the plaintiffs' arguments, the court noted that in big cases, parties often have to "produce massive amounts of information in complex litigation" that is "largely irrelevant or nonresponsive." In these situations, the court explained, permitting a party to upload large volumes of non-confidential discovery into a public AI tool could expose that material "to public consumption." This, in turn, might lead producing parties to respond by producing less or redacting more to reduce risk. The court contrasted that with the defendants' proposal, which would allow private—not public—AI tools, noting this approach would "facilitate discovery by incentivizing more fulsome document productions." For these reasons, the court found good cause and granted the defendants' motion to enter their proposed amended protective order.

Conclusion

Jeffries is an important precedent because it extends AI restrictions in a protective order to *all* discovery materials—not just those designated confidential. The court determined that private AI tools offer reasonable and necessary protections that public AI tools do not offer. The court still recognized the efficiencies AI tools provide for e-discovery tasks and permitted their use. It found, though, that the resulting risks to privilege, security, and privacy were good cause to restrict use of AI to secure, private tools. And, while the defendants' designation as critical infrastructure may have influenced the court's decision, the overarching message—that protective orders are flexible and can provide appropriate guardrails to mitigate AI-use risks when properly drafted—is transferable to all cases.

Jeffries offers several practical takeaways. First, protective orders remain an evolving tool for managing discovery risk, including by regulating how discovery is handled and processed as technology evolves. Second, it reflects courts' growing AI fluency, including a willingness to distinguish between public and private systems rather than treating "AI" as a single, undifferentiated category. Third, it underscores that public AI tools can raise heightened privilege, security, and privacy concerns even for non-confidential information, making it advisable to address tool selection and data-handling requirements early, before large-scale AI-assisted review and productions are underway. Finally, parties should address AI expressly when negotiating protective orders and draft those provisions to manage the legal and practical risks of AI use across *all* discovery materials—not just those designated confidential.