

Insights

PROTECTING YOUR CONFIDENTIAL INFORMATION IN LITIGATION: WHY PROTECTIVE ORDERS MUST ADDRESS AI

Oct 13, 2026

There is a real and growing risk that your confidential documents, including trade secrets, financial records, and other sensitive business information, could be shared with an AI platform during the course of litigation. If that happens, your confidential documents may be retained by the platform provider, used to train or fine-tune its AI model(s), and made available to others in ways that are, as a practical matter, impossible to undo. Recent court decisions confirm that traditional protective orders are not equipped to address this risk, and that parties need to adapt to the proliferation of AI tool use.

Federal courts across the country are converging on this issue. Between 2025 and early 2026, at least five federal district courts issued or analyzed protective order provisions specifically restricting the use of generative AI platforms with protected discovery materials. Two of the most significant and detailed decisions—*Morgan v. V2X, Inc.* (D. Colo. Mar. 30, 2026) and *Jeffries v. Harcros Chemicals, Inc.* (D. Kan. Mar. 25, 2026)—directly confronted how AI tools should be regulated in the discovery process, and other courts (including in Minnesota, the Western District of Washington, and the District of Nevada) have imposed their own AI-specific restrictions. Both *Morgan* and *Jeffries* recognized that standard protective orders fail to account for the unique risks posed by generative AI platforms, and each imposed or endorsed specific AI-related safeguards.

The Gap in Traditional Protective Orders

In litigation, both sides exchange documents through a process called "discovery." Courts routinely enter protective orders, agreements governing how confidential information may be used and disclosed, to prevent the opposing party from sharing sensitive documents publicly or misusing them outside the litigation.

Most protective orders were drafted before AI tools became widely available. As a result, they typically do not address what happens when a lawyer or litigation vendor uploads confidential documents into a generative AI platform for review or analysis. This gap creates significant and largely unmitigated risk.

Lawyers and litigation support professionals increasingly use AI tools to review, categorize, and summarize large volumes of discovery materials. When a party uses a consumer-grade AI tool for this purpose, the platform may retain the submitted data, use it to train or refine its model, and operate under terms of service that permit broad data use and sharing with third parties. Once confidential information has been ingested by such a platform, it is extremely difficult to retrieve or delete as a practical matter, even by court order. Although emerging research into techniques for selective deletion or suppression of data from trained AI models is ongoing, these methods are not yet widely available or recognized as reliable remedies.

How Courts Are Drawing the Line

Not all AI tools carry the same level of risk, and courts have not converged on a single test. Several distinct approaches have emerged:

Blanket prohibition: Some courts have simply barred any upload of confidential discovery material to generative AI platforms outright, without distinguishing between free consumer tools and paid enterprise tools. While straightforward to administer, this approach is increasingly difficult to reconcile with the reality and pace of AI adoption across the legal sector, and does not account for AI tools that already incorporate robust contractual protections against data retention and leakage.

Contractual safeguards approach: Other courts, most notably in *Morgan v. V2X, Inc.*, permit AI use only where the provider is contractually prohibited from storing or using inputs to train its model and from disclosing inputs to third parties (except where essential to service delivery, and then only if the third party is bound by equally protective terms), and where the provider contractually guarantees the ability to delete the information on request. Parties relying on this approach must retain written documentation of these contractual protections. As the *Morgan* court acknowledged, this standard effectively excludes most mainstream, free or low-cost consumer AI tools, though it bears noting that privacy settings and data-use practices vary across these platforms; even user interactions such as providing feedback (e.g., "thumbs up" or "thumbs down" responses) may affect whether and how the platform provider retains and uses submitted data, underscoring the need for careful, tool-specific evaluation.

Still other courts have taken different approaches. In *United States v. Jackson* and *United States v. Navarro Hernandez* (W.D. Wash.), the courts applied a data-rights test: protected material may not be loaded into any AI tool where a non-bound party retains rights to the data. In *Jeffries v. Harcross Chemicals, Inc.*, the court adopted the most comprehensive framework to date, requiring advance notice and an opportunity to object, use in a secure environment, restrictions on model training, and deletion of all confidential information at the conclusion of the case; the court later expanded these restrictions to all discovery materials and barred their use with any "open loop" or public generative AI tool entirely. The common thread is that public, data-retaining consumer AI tools pose risks that closed, contractually-restricted enterprise tools do not, even though courts differ on exactly where to draw that line.

The Business Risks of Inaction

If AI-specific protections are not incorporated into your protective orders from the outset, you face several serious risks:

- **Permanent loss of control.** Once confidential information is incorporated into an AI model's training data, retrieving or deleting it is practically impossible.
- **Competitive harm.** Trade secrets, proprietary pricing strategies, or confidential product information could become accessible to third parties through a public AI platform.
- **Regulatory exposure.** If personal data is disclosed through an AI platform, potential liability may arise under federal, state, and international data protection laws, including the GDPR and state privacy statutes.
- **No effective judicial remedy.** Although courts and researchers are actively grappling with the question of how to compel an AI platform to unlearn or purge information that has already been absorbed into its model, no reliable or widely accepted mechanism currently exists to do so.

Recommended Steps

Here is what we recommend you do now, working with your litigation team:

- **Raise AI protections at the outset of litigation.** Address AI-related risks in the initial Rule 26(f) conference, or state court equivalent, and in any proposed protective order. Do not leave these issues to informal agreement or wait until a dispute arises. Once confidential documents have been uploaded to an open AI platform, the harm cannot be undone.
- **Require that any AI tool used in your case be "closed."** Any AI platform used to process, review, or analyze your confidential documents should be governed by a contract that expressly prohibits data retention beyond the scope of the engagement, model training on your data, and third-party sharing or access.
- **Secure contractual deletion rights.** Ensure that the terms governing any AI tool used in your litigation provide for the complete deletion of all confidential materials upon written request or at the conclusion of the matter.
- **Insist on disclosure and certification obligations.** Require any party or counsel using AI tools in connection with your confidential documents to identify the specific tool, certify in writing that it meets the required contractual safeguards, and maintain that documentation for the duration of the litigation.

- **Audit existing protective orders.** If you have active litigation with a protective order already in place, review whether it addresses AI. Do not assume that silence protects you: courts have observed that a broad existing prohibition on disclosing confidential information to any third party may already cover submissions to a data-retaining AI platform, so noncompliant AI use could violate an order even without AI-specific language. An amendment to the protective order may be warranted even where no dispute has yet arisen.

Related capabilities

- Litigation
- Retail & Consumer Products
- Data Privacy & Security

Meet the team



Ashley C. Bateman

Associate, Seattle

ashley.bateman@bclplaw.com

[+1 206 600 6637](tel:+12066006637)



Goli Mahdavi

Global AI Lead, San Francisco

goli.mahdavi@bclplaw.com

[+1 415 675 3448](tel:+14156753448)



John W. Amberg

Senior Counsel, Los Angeles

john.amberg@bclplaw.com

[+1 310 576 2280](tel:+13105762280)

This material is not comprehensive, is for informational purposes only, and is not legal advice. Your use or receipt of this material does not create an attorney-client relationship between us. If you require legal advice, you should consult an attorney regarding your particular circumstances. The choice of a lawyer is an important decision and should not be based solely upon advertisements. This material may be “Attorney Advertising” under the ethics and professional rules of certain jurisdictions. For advertising purposes, St. Louis, Missouri, is designated BCLP’s principal office and Kathrine Dixon (kathrine.dixon@bclplaw.com) as the responsible attorney.