

Preventing, Detecting, and Litigating Trade Secret Theft in the Age of AI

Companies are increasingly aware that trade secrets can be lost through careless AI use—such as when employees upload proprietary information to consumer AI platforms and thereby expose it to third parties. But many remain unaware of a new and more troubling risk—AI as a tool for *deliberate* theft. Gone are the days when a thief must carefully collate files before transferring them to a thumb drive. Poorly managed AI allows trade secrets to be summarized in seconds and exported in a form that bears no resemblance to the original, often without a conventional paper trail. This alert addresses both prevention and cure for this new trade secret conundrum.

I. The Rising Threat of Insider AI Theft

AI has collapsed the traditional three-step process of trade secret theft—locate, obtain, exfiltrate—into a simple conversation. An employee can ask an unguarded enterprise AI to summarize pricing strategy, describe the development status of an unreleased product, or explain a proprietary process. No document is opened, no file is moved, and the synthesized answer may itself constitute a new form of protected information. Trade secret theft has long occurred through attempts to disguise the stolen data.¹ Now, proprietary information that has been converted through AI may bear little resemblance to its sources—further complicating detection as AI-assisted exfiltration becomes increasingly sophisticated.²

1 See, e.g., Former Google Engineer Found Guilty of Economic Espionage and Theft of Confidential AI Technology, U.S. Dep't of Justice, <https://www.justice.gov/opa/pr/former-google-engineer-found-guilty-economic-espionage-and-theft-confidential-ai-technology> (conviction on trade secret theft by engineer who copy-pasted source code into PDF files that evaded Google's loss prevention systems).

2 See, e.g., 'Exploit Every Vulnerability': Rogue AI Agents Published Passwords and Overrode Anti-Virus Software, The Guardian (Mar. 12, 2026), <https://www.theguardian.com/technology/ng-interactive/2026/mar/12/lab-test-mounting-concern-over-rogue-ai-agents-artificial-intelligence> (describing AI agents that cooperated to smuggle information out of supposedly secure systems).

A determined employee need not attempt to obtain a trade secret in a single session. They can query an enterprise AI about different components across multiple sessions—each query seemingly innocuous in isolation—and later collate the results into a functional reproduction of the trade secret. This pattern defeats the bulk-download alerts that conventional data-loss-prevention tools are designed to catch. The relevant forensic question is not whether a single large extraction occurred, but whether a pattern of isolated queries adds up to one.

Enterprise AI platforms maintain significant context about an employee’s work through memory and project features. Many platforms now offer memory export functions that allow users to transfer their conversation history and context to other accounts.³ A departing employee can achieve wholesale knowledge transfer—in summarized, or even disguised form—with a few simple prompts.

Over-permissioned enterprise AI tools can automate the location and collation phase of a theft. If an organization has not enforced least-privilege access controls across its AI-accessible data sources, its enterprise AI may synthesize results from sources that an errant employee might never have located manually—effectively acting as an expert accomplice to the theft.

II. Risk Considerations: AI Access and Controls

The risks of AI-enabled theft may lead some organizations to evaluate controls that go beyond standard acceptable-use policies and NDAs. Various approaches have emerged, each involving tradeoffs that will vary by organization.

Some organizations may want to consider role-based access controls that limit what queries a given employee can make through an enterprise AI platform. Such controls can both constrain the scope of insider threats and generate audit trails—but they require ongoing maintenance as roles and systems evolve and may not be feasible for all platforms. A separate question concerns the underlying data permissions that an enterprise AI inherits. Where a company’s data repositories have accumulated broad access rights over years of organic growth, an enterprise AI can surface sensitive data that an employee could never have located manually. Organizations may want to consider auditing and restricting the underlying permissions before enabling broad AI retrieval capabilities; a process that requires tradeoffs between a particular AI deployment and the organization’s underlying data architecture.

Some AI platforms offer administrator controls that limit or disable memory export features. Limiting memory portability can reduce the risk of knowledge transfer at departure, though the availability and effectiveness of these controls depends on the platform and may require specific provisions in third-party vendor contracts. Organizations evaluating this area may find that their ability to implement controls depends substantially on what is negotiated at the time of vendor contract.

Some organizations have updated employment agreements and NDAs to address AI-specific conduct—including personal account use, memory export, and disclosure of AI tools used during employment. AI-specific provisions may serve as deterrents and, where a dispute arises, as evidence bearing on willfulness under the Defend Trade Secrets Act (“DTSA”). The DTSA definition

3. See, e.g., *Leaving ChatGPT for Claude? Here’s the Trick to Taking Your AI Memory With You*, PCMag, <https://www.pcmag.com/explainers/leaving-chatgpt-for-claude-heres-the-trick-to-taking-your-ai-memory-with>; *Make the switch: Bring your AI memories and chat history to Gemini*, Google, <https://blog.google/innovation-and-ai/products/gemini-app/switch-to-gemini-app/> (describing AI agents that cooperated to smuggle information out of supposedly secure systems).

of “employee” extends to independent contractors and consultants,⁴ and equivalent considerations apply to those agreements. Agreements covering trade secrets will likely want to include the whistleblower immunity notice required by DTSA § 1833(b)⁵ as omitting this notice forfeits the ability to seek exemplary damages and attorney’s fees in a subsequent DTSA action—a potentially significant penalty in a high-stakes case where conduct can be egregious but damages difficult to quantify.

III. Building the Trial Record: What to Do When Theft is Suspected

Some of the most probative evidence in an AI-based trade secret case may be held by an AI platform—not an employer’s systems—and may be subject to deletion. Enterprise AI platform audit logs, browser history, clipboard logs, network traffic to AI platform domains, and data-loss-prevention alerts should be reviewed before an employee knows they are under investigation. Where platform-side records are needed, a preservation demand or, if necessary, an emergency TRO should be sought before the employee can delete their account. These contemporaneous logs may be the most powerful evidence in a trade secret case: they are difficult to fabricate and, unlike witness testimony, not subject to the credibility disputes that can dominate trade secret trials.

AI-based misappropriation frequently does not manifest as a single discrete event. The relevant forensic question can instead be whether a pattern of queries—across sessions, subjects, and time—constitutes a functional reproduction of a trade secret. That analysis requires a forensic expert who can work across AI platform logs, enterprise access records, and the content of AI outputs. Where a company suspects indirect prompt injection—a form of theft that is likely to grow as AI use increases—a targeted audit of documents provided by the employee to the company’s AI systems before departure may also be warranted. Planted exfiltration instructions can cause a system to continue surfacing sensitive data after the employee’s access has been nominally revoked.

In AI-based trade secret cases, where misappropriated information may have been incorporated into the parameters of a model, the harm can be technically irreversible. Courts are beginning to grapple with the available remedies. At one end is “algorithmic disgorgement”⁶—destruction of any model into which proprietary information has been unlawfully incorporated—a remedy the FTC has required in analogous contexts⁷ and that private plaintiffs have sought.⁸ A technically lesser alternative is “machine unlearning,” which selectively diminishes the influence of specific data on a model’s weights without destroying the model. Machine unlearning is not always a reliable cure: research indicates that, with some models, traces of removed data can often still be detected through adversarial auditing and that imperfect unlearning may itself provide a roadmap for inferring what was removed.⁹

4. 18 U.S.C. § 1833(b)(4) (defining “employee” to include “any individual performing work as a contractor or consultant for an organization”).

5. 18 U.S.C. § 1833(b)(3)(C) (providing that employer that fails to include whistleblower notice required by § 1833(b)(1) “may not be awarded exemplary damages or attorney fees” in subsequent DTSA action).

6. Daniel Wilf-Townsend, *The Deletion Remedy*, 103 N.C. L. Rev. 1809 (2025), https://northcarolinalegallawreview.org/wp-content/uploads/sites/5/2025/09/4-Wilf-Townsend_FinalForPrint.pdf

7. In a data privacy context, the FTC has ordered model deletion in several recent enforcement actions. See generally Joshua A. Goland, *Algorithmic Disgorgement: Destruction of Artificial Intelligence Models as the FTC’s Newest Enforcement Tool for Bad Data*, 29 Rich. J.L. & Tech. 1 (2024), <https://scholarship.richmond.edu/jolt/vol29/iss2/1/>

8. In a copyright context, the New York Times sought destruction of any AI models trained on the misappropriated data. See *The New York Times Co. v. OpenAI, Inc.*, No. 1:23-cv-11195 (S.D.N.Y., filed Dec. 27, 2023), Complaint at Prayer for Relief, <https://www.courtlistener.com/docket/68117049/1/the-new-york-times-company-v-microsoft-corporation/>

The DTSA permits exemplary damages of up to twice actual damages for willful and malicious misappropriation. AI-based investigations often develop evidence bearing directly on willfulness—queries outside an employee’s functional role, platform access from a personal device to avoid corporate monitoring, or memory export shortly before resignation. The contemporaneous log record can tell that story in a way that is difficult for the defense to rebut. Organizing and preserving that evidence from the outset with the willfulness argument in mind affects both the damages exposure and the settlement posture of the case.

Takeaways For Companies Managing AI Risk: How AI-enabled trade secret risk manifests—and what responses are appropriate—will vary significantly depending on how AI tools have been deployed, what access controls exist, and how agreements with employees and vendors are currently structured. The access architecture, memory controls, and agreement considerations described in this alert may be worth evaluating against current practices. Whether and how to address any gaps will depend on each organization’s specific circumstances, including operational constraints and the sensitivity of the information involved.

Takeaways For Companies That Suspect AI-Enabled Theft Has Occurred: The evidentiary window in AI-based trade secret cases is often short. Evidence held by AI platforms may be subject to deletion, and engaging counsel before notifying the employee can materially affect the strength of any subsequent case. Traditional endpoint review focused on file downloads may not capture the full picture of AI-based exfiltration. We are glad to discuss what a typical investigation and case-building process involves, and how an investigation might be tailored to an organization’s particular circumstances.

⁹ See European Data Protection Supervisor, *Machine Unlearning*, TechSonar (2024), https://www.edps.europa.eu/data-protection/technology-monitoring/techsonar/machine-unlearning_en; Chen et al., *A Survey of Security and Privacy Issues of Machine Unlearning*, AI Magazine (Wiley, Jan. 2025), <https://onlinelibrary.wiley.com/doi/full/10.1002/aaai.12209>

If you have any questions about the issues addressed in this memorandum, or if you would like a copy of any of the materials mentioned in it, please do not hesitate to reach out to:



Stacylyn Doore

Partner

Boston

stacylyndoore@quinnemanuel.com



Ryan Landes

Partner

Los Angeles

ryanlandes@quinnemanuel.com



Patrick Curran

Partner

Boston

patrickcurran@quinnemanuel.com



Toby Futter

Of Counsel

New York

tobyfutter@quinnemanuel.com

To view more memoranda, please visit www.quinnemanuel.com/the-firm/publications/

To update information or unsubscribe, please email updates@quinnemanuel.com