MELAND | BUDWICK

Non-Disclosure Agreement Drafting Must Account for AI's Risks

https://melandbudwick.com/2026/06/17/non-disclosure-agreement-drafting-must-account-for-ais-risks/

By Eric Ostroff *| June 17, 2026*

Your company has signed countless non-disclosure agreements. So has every company you do business with. They govern virtually every deal and vendor relationship that touches proprietary information. Your employees feed that protected information into artificial intelligence tools every day, and so do your counterparties. They aren't trying to steal anything; they're trying to do their jobs. But they may be violating the NDA in the process.

A standard NDA restricts the receiving party from disclosing confidential information to anyone other than employees and authorized representatives with a need to know. When an employee pastes a counterparty's confidential information into a consumer AI tool, the information leaves the receiving party's environment and is processed on a third party's servers. The provider isn't an employee or authorized representative.

Worse, terms of service for consumer-tier large language models often reserve the right to use inputs to train future models, which means the confidential information is potentially absorbed into a third-party product. Although these tools typically allow for user configurations that opt-out of training, many employees haven't properly configured their account to address this issue.

Whether software counts as a "recipient" is a question most NDAs were never drafted to answer. But the disclosing party's core concern is the same either way: The information is no longer under the receiving party's control.

The return-or-destroy problem is worse. Most NDAs contain a return-or-destroy clause requiring the receiving party to give back or delete all confidential information at the end of the relationship, often requiring certification of compliance in writing. A receiving party whose employees have fed confidential information into

an AI tool can't honestly make that certification.

While the receiving party can delete the chat history, it can't reach into the provider's server logs, cached prompts, model fine-tuning data, vendor backups, or sub-processor systems. Once that data has been used for training purposes, it becomes very difficult, if not impossible, to certify compliance. Lack of visibility into how the AI companies use data for training purposes complicates the problem.

With enterprise platforms that contractually commit to zero retention or no training on inputs, the gap may be manageable, if the receiving party can document which platform was used and what its terms were. With consumer tools, the receiving party is signing a certificate it has no way to support.

The scale of AI use also should worry general counsel. Multiple recent surveys put the share of employees who have pasted confidential company data into a public AI tool somewhere north of half, and many organizations still have no governing policy. If even a fraction of that activity involves information received under an NDA, the volume of potentially compromised confidential information across the economy is enormous.

Contract drafting hasn't caught up. The disclosing party assumes its NDA is being honored. The receiving party's employees, including its lawyers, often don't realize that running someone else's confidential information through an AI tool might implicate the agreement at all.

The exposure runs in both directions. Every company is on both sides of these agreements, and every company has employees feeding the other side's confidential information into AI tools while doing day-to-day work.

The drafting fix isn't complicated. NDAs going forward should address AI directly. At minimum they should:

- State whether the receiving party may use AI tools to process confidential information at all;
- If so, distinguish between enterprise platforms with contractual zero-retention and no-training commitments, which may be acceptable, and consumer tools, which generally aren't;
- Require disclosure upon request of which AI tools were used with the

information;
- Extend the return-or-destroy obligation, on a reasonable-efforts basis, to the AI provider's environment—and qualify the destruction certification to acknowledge what can't be reached;
- Address agentic AI in which the tool receives the information and acts on systems containing it. This is an emerging category that magnifies every problem described above.

This is the same exercise NDA drafting has always involved: identifying the channels through which confidential information can leave the receiving party's control and closing them. AI is a channel that didn't exist when most NDAs currently in force were drafted. The language needs to catch up.

Until it does, companies are carrying a risk they haven't accounted for. The disclosing party's information may be less protected than it believes, and the receiving party may be in breach of an obligation without knowing it. Because NDAs are such a commonplace tool, companies need to resist relying on old, pre-AI forms and instead update their language to account for this new risk.

This article was originally published by Bloomberg Law on June 17, 2026.