MELAND | BUDWICK

Cos. Must Update Protocols To Protect Trade Secrets From AI

https://melandbudwick.com/2026/04/23/cos-must-update-protocols-to-protect-trade-secrets-from-ai/

By Eric Ostroff *| April 22, 2026*

In March, Meta confirmed that an artificial intelligence agent used by an employee inside Meta's internal infrastructure exposed sensitive user and company data to unauthorized employees for two hours.

The agent autonomously generated a solution that bypassed access controls no human would have deliberately ignored. No one acted maliciously. The agent simply solved a problem using every resource available to it, including data it should never have touched.

If that can happen with a company-sanctioned agent operating inside one of the most sophisticated security environments on earth, consider what happens when employees deploy their own.

A similar scenario is playing out right now across every industry, largely invisible to the companies absorbing the risk.

A sales director at a software company sets up an AI agent — something like OpenClaw, one of a growing number of open-source autonomous agent frameworks — and connects it to his work email, his calendar and the company customer relationship management software. He is not trying to steal anything; he is trying to keep up. The agent drafts follow-up emails, summarizes meeting notes, flags renewal risks and pulls competitor intel from the web. It is, by every measure, making the sales director more productive.

Eighteen months later, he resigns and joins a direct competitor. The company runs through its standard offboarding checklist: laptop returned, email access revoked, nondisclosure agreements acknowledged, etc. But nobody asks about the agent. In fact, the HR director doesn't even know what an AI agent is.

The problem is that the agent has been running on a consumer-tier large language model without enterprise data protections. Over 18 months, it processed thousands of internal communications, synthesized customer pricing histories, learned deal structures and built a working model of the company's most sensitive sales strategy.

The employee may have forgotten half of what the agent learned. The agent has not forgotten any of it, and the employee still has access to the agent.

Companies need to take steps to guard against this novel risk factor.

What Autonomous Agents Actually Do

An autonomous agent is not a chatbot. It does not wait for questions. The agent is given a goal and a set of tool integrations — email, calendar, Slack, CRM, internal databases, etc. — and it pursues that goal by taking actions, retrieving information and making decisions across those systems.

Frameworks like OpenClaw have made this capability accessible to nondevelopers. An employee with basic technical curiosity can connect one of these agents to their work systems in an afternoon.

The result is a fundamentally different kind of data exposure. The employee is not copying files. The agent is continuously aggregating, synthesizing and retaining information across every system it touches. By the time anyone realizes this is happening, the exposure is months deep.

Three Threat Windows Companies Are Not Thinking About

The Unintentional Exfiltration Problem

An autonomous agent can exfiltrate trade secrets with zero intent on anyone's part. The employee wanted to work faster. And the agent did what it was designed to do. Somewhere in that transaction, pricing models, customer information and product roadmaps moved from inside the company's walls to a consumer AI platform. This is analogous to the Meta situation, though there the breach was only internal. The nightmare scenario is when the agent exfiltrates trade secrets externally.

The Nonenterprise LLM Problem

Enterprise AI platforms, like Claude Enterprise, Microsoft Copilot Enterprise and

comparable products, carry contractual commitments about data handling: no training on user inputs, audit logs or data isolation.

Consumer-tier tools typically offer no such protections. When an employee routes trade secrets through a consumer AI agent, the company may have lost meaningful control of that information in ways that cannot be remediated later.

This creates a serious reasonable-measures problem. The Defend Trade Secrets Act and its state equivalents condition protection on the company taking reasonable steps to maintain secrecy. Allowing employees to process trade secrets through consumer AI tools, even unknowingly, is a credible basis to challenge whether those measures were reasonable.

The Departure and Postdeparture Problem

When an employee leaves, companies know how to handle devices and access credentials. Nobody knows how to handle the agent.

If the agent was running with persistent memory, which many configurations enable, it may have synthesized a working model of the company's most sensitive information. Unlike a downloaded file, there is no clear artifact to image. And unlike an email, there is no metadata showing where it was sent.

The agent is personal property, running on the employee's own accounts, holding something that is functionally equivalent to a comprehensive internal briefing document.

The Reasonable-Measures Framework Needs an Update

The reasonable-measures doctrine was built around physical and digital controls operated by humans. Those controls assumed a relatively static relationship between an employee and information. The employee accessed specific files, in specific systems, at specific times.

Autonomous agents break that assumption. They operate continuously, across multiple systems simultaneously, and they synthesize in ways that no individual access event would reveal. Companies that are serious about protecting their trade secrets need to address this on three fronts.

First, technical controls. The systems that house trade secrets should be configured to block unauthorized agent integrations. This is not a novel technical challenge. Platforms like Microsoft 365 and Google Workspace already include admin controls that restrict which third-party applications can connect to employee accounts. But those controls need to be configured appropriately.

Second, policy. Companies need written AI usage policies that specifically address autonomous agents and personal AI tools. The policy should identify which platforms are approved for use with company data, prohibit the connection of personal AI agents to internal systems, and make clear that approved tools means enterprise-tier deployments with appropriate data handling commitments. A general technology-use policy written before 2026 almost certainly does not cover this.

Third, and most often skipped: training. A policy that employees have not read and do not understand provides little protection, either as a practical matter or as evidence of reasonable measures in litigation. Active, documented employee training on AI tool policies serves two purposes: It reduces the risk of inadvertent misuse, and it strengthens the company's reasonable-measures posture if litigation follows.

Conclusion

AI agents present a novel threat to companies' trade secrets. As this technology continues to evolve rapidly, companies need to think proactively about how to address this threat head on. This will require an understanding of both the developing technology and its interaction with long-standing trade secret law.

Indeed, the Defend Trade Secrets Act and its state-law equivalents require a plaintiff to prove both that it took reasonable measures to protect the information and that the defendant acquired it through improper means. Autonomous agents complicate both elements.

On reasonable measures, a company that restricts system access, requires NDAs and trains employees on confidentiality will often satisfy the standard. But if that same company allows employees to connect personal AI agents to those systems without restriction, a defendant will argue the company failed to protect the very information it now claims as secret. Courts may scrutinize the absence of AI-specific

controls.

On improper means, an employee who connects a personal agent to company systems in violation of a written policy, routes proprietary data through a consumer platform, and retains it through persistent memory after departure fits comfortably within the statutory language. But if no policy existed and no technical controls were in place, the company's argument weakens.

The two elements reinforce each other: failing to address AI agents in a trade secret protection program undermines both at once. Companies need to be thinking about these issues now.

This article was originally published by Law360 on April 22, 2026.