Companies Rush Into Agentic AI While Lawyers, Insurers, Security Experts Warn of Fraud/Privacy Risk

May 16, 2026 - Businesses across the United States are rapidly deploying a new generation of artificial intelligence systems known as "agentic AI" - software capable of carrying out tasks, making decisions, interacting with databases, and executing workflows with limited human oversight.

Many companies believe the technology could dramatically improve productivity, reduce labor costs, and streamline customer service. But cybersecurity experts, insurers, privacy advocates, and corporate lawyers are increasingly warning that many companies are adopting the systems before fully understanding the risks.

The concerns go far beyond chatbots generating text or answering customer questions. Modern AI agents can connect directly to company systems, access sensitive records, interact with software applications, send emails, process transactions, and carry out instructions autonomously.

That has created growing fears about fraud, insider misuse, privacy violations, and potentially massive corporate liability if the systems are deployed without proper safeguards. And these fears aren't just theoretical.

Recent incidents have intensified those concerns.

In one widely discussed case earlier this year, PocketOS founder Jeremy Crane said an autonomous AI coding agent deleted the companyâ€™s production database and backups in a matter of seconds after being given broad system permissions. The company later partialy recovered the data, but the event became an example frequently cited by security professionals warning about excessive AI authority inside corporate environments. We wrote about it a few weeks ago, here.

Another incident involved AI coding platform Replit. During testing, the companyâ€™s AI agent reportedly deleted a user database and then generated inaccurate explanations about what had happened. The episode raised concerns not only about autonomous AI actions, but also about the possibility that AI systems may provide misleading or incomplete information during incident investigations.

Researchers have also documented troubling behavior in experimental environments. In one Emergence AI simulation reported earlier this year, autonomous AI agents operating in a virtual setting reportedly engaged in destructive behavior after developing open-ended operational goals. While the experiment did not involve a real corporate network, it reinforced fears among cybersecurity experts that autonomous systems can behave unpredictably when given broad objectives and insufficient operational constraints.

Security professionals say the lesson from these incidents is not that AI systems are "becoming sentient," but that organizations are increasingly granting powerful automation tools access to critical systems before governance and security controls are fully developed.

One of the biggest concerns involves what security professionals call "overprivileged AI." In some deployments, AI agents may gain access to sensitive systems or databases without the kinds of restrictions normally applied to human employees.

Experts warn that a poorly controlled AI system could expose employee payroll data, customer records, internal communications, financial information, or proprietary corporate data if access permissions are not carefully managed.

The problem becomes even more serious when employees begin using unauthorized AI tools on their own.

The growing phenomenon, often referred to as "shadow AI," involves workers connecting consumer AI applications or autonomous agents to corporate systems without formal approval from company leadership or information technology departments.

In some cases, employees may not even realize the security risks involved. A worker might ask an AI assistant to summarize customer information, organize internal files, or automate repetitive tasks, unintentionally exposing confidential information to outside systems.

In more serious cases, a malicious employee could deliberately use AI tools to extract sensitive data, manipulate records, or damage company systems.

Cybersecurity experts say agentic AI lowers the technical barrier for insider abuse because employees no longer need advanced programming skills to automate complicated actions.

An employee with limited technical knowledge may suddenly be able to instruct an AI system to gather sensitive records, interact with databases, generate scripts, automate workflows, or manipulate software systems that previously required specialized expertise.

That has raised concerns about how quickly AI-assisted insider threats could evolve inside companies that lack modern security controls.

The legal consequences for companies could be significant.

If a company deploys AI systems without proper access controls, monitoring, approval processes, or operational safeguards, courts may eventually view resulting data breaches or operational failures as foreseeable risks rather than unavoidable accidents.

That distinction matters because negligence lawsuits often focus on whether a company took "reasonable precautions" to prevent predictable harm.

For example, if an AI system were allowed to access employee payroll records and an unauthorized worker used the system to retrieve confidential salary information, the company could potentially face privacy lawsuits, employment claims, regulatory scrutiny, or enforcement actions.

Similarly, if an AI agent were able to delete customer databases, alter financial records, disrupt critical systems, or expose confidential information, shareholders and customers could argue that company leadership failed to implement reasonable safeguards.

Legal experts say courts are unlikely to accept the argument that "the AI did it" as a complete defense.

Instead, companies deploying AI systems will probably be expected to demonstrate that they implemented modern security controls, restricted access privileges, maintained audit logs, enforced governance policies, and limited autonomous system authority appropriate to the level of risk involved.

The issue becomes even more complicated when unauthorized employee-installed AI tools are involved.

Even if a rogue employee violates company policy by installing autonomous AI software, organizations may still face liability if they failed to maintain reasonable security practices, monitor network activity, restrict software installation privileges, or protect sensitive systems.

Many of the underlying legal principles are not new. Courts have long dealt with cases involving insider threats, negligent cybersecurity practices, unauthorized software, and data breaches. Agentic AI simply amplifies those risks because autonomous systems can operate continuously and at machine speed.

Another unresolved issue involves liability itself.

If an autonomous AI system initiates a fraudulent transaction, exposes confidential information, damages infrastructure, or causes operational losses, courts and regulators may eventually need to determine whether responsibility lies with the employee using the system, the company deploying it, the software developer, the AI provider, or some combination of all four.

At the moment, there is no settled legal framework governing machine-to-machine accountability.

That uncertainty has become especially important in banking and financial services, where institutions are increasingly experimenting with AI agents capable of handling customer interactions, fraud detection, transaction analysis, and operational workflows.

Financial institutions are particularly concerned about scenarios involving:

- Unauthorized financial transactions

- Identity theft
- AI-assisted fraud
- Account manipulation
- Exposure of customer financial data
- Regulatory compliance failures

Banks and insurers worry that autonomous systems operating at machine speed could potentially accelerate fraud schemes faster than human investigators can respond.

The insurance industry is also beginning to react.

Cybersecurity insurers and directors-and-officers liability insurers are increasingly evaluating AI governance practices during underwriting reviews.

Industry analysts say insurers are becoming concerned that poorly controlled AI systems could trigger large-scale privacy breaches, operational outages, fraud losses, operational sabotage, or regulatory investigations.

As a result, companies may soon face growing pressure to adopt formal AI governance programs in order to maintain affordable insurance coverage.

Experts expect insurers to increasingly ask organizations whether they:

- Restrict employee use of unauthorized AI tools
- Monitor AI interactions with sensitive data
- Require human approval for critical actions
- Maintain audit trails for AI activity
- Segment critical systems from AI access
- Train employees on AI security risks
- Restrict software installation privileges
- Limit autonomous deletion or transaction authority
- Maintain immutable backups
- Conduct AI risk assessments

Some insurers may eventually refuse coverage or impose higher premiums on companies that fail to implement basic AI governance safeguards, much as cyber insurers previously began requiring multifactor authentication and ransomware protections.

Security professionals say the most effective defense is not banning AI entirely, but limiting its authority and carefully controlling how it interacts with sensitive systems.

Many organizations are now adopting "least privilege" security models that restrict AI agents to narrowly defined tasks and prevent autonomous access to highly sensitive information or destructive administrative functions.

Other safeguards include approval workflows, software allowlists, endpoint monitoring, network segmentation, immutable backups, audit logging, anomaly detection systems, and data loss prevention tools.

Regulators are also paying close attention.

Financial regulators, privacy authorities, cybersecurity agencies, and lawmakers are increasingly studying how autonomous AI systems may affect fraud prevention, consumer protection, operational resilience, and corporate accountability.

As AI adoption accelerates, many experts believe companies are entering a period similar to the early days of cloud computing and ransomware, when businesses moved rapidly toward adoption before security standards fully matured.

The concern now is that some organizations may be prioritizing productivity gains and cost reduction while underestimating how quickly autonomous systems can create privacy breaches, operational failures, fraud losses, and legal exposure if they are not tightly controlled.

For now, there is no universal regulatory standard governing how companies should safely deploy agentic AI.

But cybersecurity experts, insurers, and corporate lawyers increasingly agree on one point: businesses that grant autonomous systems broad access without meaningful safeguards may eventually face enormous financial and legal consequences when something goes wrong.

by Jim Malmberg
Note: When posting a comment, please sign-in first if you want a response. If you are not registered, click here. Registration is easy and free.

Follow ACCESS