

A framework for governing employees' existing use of generative AI

By Tara Lawler, Esq., and Matthew Hamilton, Esq., Morgan Lewis

OCTOBER 7, 2026

Generative AI (GenAI) entered the workplace through the side door. Long before many companies adopted formal policies, employees began using public tools to summarize meetings, draft emails, analyze spreadsheets, translate documents, write code, brainstorm marketing copy, prepare interview questions, and find answers faster than traditional systems allow.

For business leaders, the immediate question is no longer whether employees should be allowed to use GenAI, but rather how to govern use that is already occurring.

The scale of informal adoption is significant. The 2026 Akamai Enterprise AI Usage Risk Report (<https://bit.ly/4hK3eob>) warns that 47% of all AI interactions at work take place through personal identities rather than corporate-managed accounts, while the PagerDuty Shadow IT survey (<https://bit.ly/4yuRrRQ>) found that 66% of office professionals have used AI tools at work despite not being permitted under company policies.

This combination of widespread use and individual selection of tools creates the classic “shadow IT” problem, updated for the GenAI era: business activity is taking place through systems that IT, legal, privacy, security, and compliance teams may not see, control, or understand.

The instinctive response may be to ban unapproved GenAI tools until a comprehensive policy is finalized. For many organizations that will be both unrealistic and counterproductive.

Employees are using AI to address legitimate workplace challenges: too much information, too many routine communications, too many administrative tasks, and growing pressure to produce more with fewer resources. A prohibition that does not provide usable alternatives is likely to drive use underground, making the risk harder to manage.

A better starting point is a practical assumption: employees are already using GenAI. From there, companies can build governance that channels that behavior and requires use of only approved tools, appropriate use cases, and meaningful oversight.

The risks extend beyond confidentiality

When an employee uploads a customer contract, source code, board deck, trade secret, personnel file, or litigation document

into an unapproved AI tool, that reduces the company's ability to control its data, including how long it is retained and whether it will be used to train the GenAI model. The level of risk will depend on the tool, subscription type, settings, and contractual terms.

While enterprise AI products may offer materially different protections from free consumer accounts, that does not eliminate the need for vendor, security, privacy and contractual review.

Cybersecurity risk is also evolving. IBM's 2026 Cost of a Data Breach Report (<https://bit.ly/4yUvHPL>) described an “AI oversight gap,” reporting that 68% of organizations breached lacked AI governance policies to manage AI or prevent the proliferation of shadow AI, and that security incidents involving use of shadow AI more than doubled to 43% this year from 20% the year before.

GenAI tools may also produce inaccurate, biased, incomplete, or fabricated outputs. They may generate misleading communications, or create records that are difficult to explain later in an audit, investigation, or lawsuit.

NIST's GenAI profile (<https://bit.ly/3TG08d1>) identifies risks involving data privacy, harmful bias, information integrity, information security, and intellectual property. It recommends that organizations manage these risks throughout the AI lifecycle.

Start with visibility

Companies should conduct a focused AI use assessment, not a punitive investigation. Business units should be asked what tools they use, what tasks they use them for, what data they input, the outputs on which they rely, and where AI may already be embedded in vendor products.

Monitoring of network traffic, browser extensions, software procurement records, and other indicators of unauthorized or unmanaged AI use should be conducted consistent with applicable employee monitoring, privacy, labor, and other company policy requirements.

Classify use cases by risk

A practical policy should distinguish among low-risk productivity support, moderate-risk business drafting,

and high-risk uses that require approval, testing, documentation, or human sign-off.

Risk classifications should be based not only on the task being performed, but also the sensitivity of the data involved, the intended audience, the consequences of an error, and the extent to which the output will influence a significant decision.

Provide tools employees will actually use

Shadow AI problems arise where an employee is using the unofficial tool due to it being faster, easier, or more capable than the sanctioned alternative. If a company requires employees to use an internal or approved AI platform, that platform should be accessible, reliable, and integrated into ordinary workflows.

A policy that tells employees what not to do but gives them no practical way to complete their work is unlikely to succeed. The responsible path must also be the workable path.

Establish clear red lines

Employees should understand that certain categories of information may not be entered into unapproved AI tools.

Depending on the organization, those categories may include:

- Trade secrets
- Material nonpublic information
- Attorney-client privileged or work-product-protected communications
- Regulated personal data
- Customer confidential information
- Sensitive employee records

Employees should also be instructed that AI should not serve as the sole basis for consequential decisions involving employment, credit, housing, healthcare, legal rights, safety, or disciplinary action.

The policy should be written in plain language and supported by examples drawn from the company's actual business. Employees are more likely to follow a policy when they can recognize how it applies to their daily work.

Treat training as a control

Training should be treated as a control, not a box-checking exercise. Employees must understand what prompts and outputs are, why copying sensitive data into an AI tool can create risk, the risk of hallucinations/confabulations, when disclosure may be needed, and when human review is mandatory.

Training should also be role specific: a software developer, human resources professional, marketing employee, and in-house lawyer may face very different risks when using the same tool.

Revisit vendor contracting

AI tools raise questions that may not be fully addressed in standard software agreements. Businesses should determine:

- Whether customer data may be used for model training
- Whether prompts and outputs are retained and if so for how long
- Where data is stored
- Whether subprocessors are used
- What security certifications and controls apply
- What audit rights are available
- How material model or product changes will be disclosed
- What indemnities or liability limitations apply if the tool causes legal, security, or operational harm

Organizations should also consider whether their existing vendor review process is capable of identifying AI functionality embedded within products that were not originally purchased or marketed as AI tools.

Define meaningful human review

A vague requirement to keep a "human in the loop" is not enough. The reviewer must have the expertise, time, information, and authority to challenge the AI output.

For high-impact uses, the review process should require verification against reliable sources, documentation of key assumptions, and a record of who approved the final decision or communication. The more consequential the decision, the more rigorous the review should be.

Address workplace culture

Employees may hide AI use if they fear discipline, replacement, or criticism for relying on new tools. Companies should communicate a candid and balanced message: responsible AI use is encouraged where it improves work but it must be transparent, secure, and consistent with the company's legal, regulatory, and contractual obligations.

Managers should model appropriate use and reinforce that accuracy, confidentiality, and accountability remain human responsibilities.

Shadow AI is not merely a compliance failure. It is also a signal of unmet demand. Employees are showing companies where work is repetitive, slow, or poorly supported by existing systems.

For companies, the goal should be neither blind adoption nor blanket prohibition. The goal should be managed use: approved tools, clear rules, trained employees, protected data, and accountable human judgment.

Generative AI is already at work. The task now is to bring it out of the shadows.

About the authors



Tara Lawler (L) is deputy leader of **Morgan Lewis'** data analytics and governance practice, focusing on discovery, information governance, data privacy, security, and artificial intelligence. She partners with in-house legal and technology teams to assess, develop, and implement best practices for discovery and data management. She can be reached at tara.lawler@morganlewis.com.

Matthew Hamilton, (R) of counsel at the firm, is a civil litigator and e-discovery counsel who represents clients in complex pharmaceutical and medical products liability, data breach, antitrust, and commercial litigation. He can be reached at

matthew.hamilton@morganlewis.com. The authors are based in Philadelphia. The authors would like to thank associate Leonard Impagliazzo for contributing to this article.

This article was first published on Reuters Legal News and Westlaw Today on October 7, 2026.

© 2026 Thomson Reuters. This publication was created to provide you with accurate and authoritative information concerning the subject matter covered, however it may not necessarily have been prepared by persons licensed to practice law in a particular jurisdiction. The publisher is not engaged in rendering legal or other professional advice, and this publication is not a substitute for the advice of an attorney. If you require legal or other expert advice, you should seek the services of a competent attorney or other professional. For subscription information, please visit legalsolutions.thomsonreuters.com.