Shadow AI is already here: Take control, reduce risk, and unleash innovation

August 2025

Addressing a pervasive problem

The use of generative artificial intelligence (GenAI) by employees across industries is exploding. According to a recent KPMG study, up to 58 percent of employees are using AI productivity tools on a daily basis.[1] There are some compelling reasons for this surge—workers love that the technology eliminates mundane tasks, it can save a lot of time, it's easy to use, and it amplifies productivity. A clear example is the rise of "vibe coding," where employees use GenAI to generate code from natural language prompts—accelerating prototyping but creating risks when done outside approved environments. But this surge isn't just about the tools themselves; it reflects something deeper. Workers are turning to personal or external AI tools they perceive as faster or more capable than the ones their company provides—or hasn't provided—using them to move quickly, solve problems independently, and bypass slow or outdated enterprise systems. In fact, almost half of employees admit to uploading sensitive company information to unauthorized platforms.[2]

And now this dynamic must address "shadow AI": the unsanctioned or unauthorized use of AI tools without the explicit approval or oversight of the IT department or a central AI governance team. Unlike traditional shadow IT, this isn't about rogue apps or hardware—it's about employees using AI tools and platforms that feel frictionless, flexible, and fast, even if they fall outside the organization's governance perimeter. Employees often resort to shadow AI because it is faster, easier, and less restrictive than a company's official tools. This trend is driven by the accessibility to AI tools, competitive pressures, and the familiarity with GenAI among AI-savvy employees.

Many executives would be shocked to know just how pervasive the problem is within their own organizations. The results from a recent global survey conducted by KPMG and the University of Melbourne titled, "Trust, Attitudes and Use of Artificial Intelligence," uncovered some disturbing trends (see stats on the right).

A KPMG survey uncovered disturbing AI workplace trends:

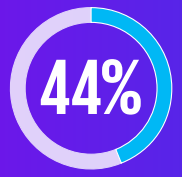

of employees have used AI in ways that contravene policies and guidelines, indicating a significant prevalence of shadow AI in organizations.

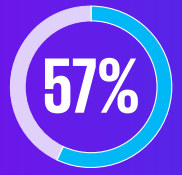

of employees have made mistakes due to AI, and 58 percent have relied on AI output without evaluating its accuracy.

of employees report that their organization has a policy guiding the use of GenAI, highlighting a huge gap in guardrails.

[1] "KPMG AI Quarterly Pulse Survey: What executives are saying now," KPMG LLP, 2025

[2] "Nearly half of workers using AI at work admit to doing so inappropriately," Fast Company, April 30, 2025

Shadow AI is outpacing enterprise controls

Shadow AI is actually a symptom of deeper friction between employee needs and AI capabilities provided by an enterprise. It often arises when official tools are outdated, lack essential features, or aren't well integrated into daily workflows. Rigid provisioning and overreaching security protocols further widen this gap, driving employees toward faster, frictionless, more intuitive alternatives.

Shadow AI thrives on behavioral drivers like convenience, curiosity, and the appeal of consumer-grade simplicity. Many GenAI tools are free, intuitive, and familiar—mirroring the user experience of popular consumer apps. Employees, particularly digital natives, are eager to experiment, reduce friction, and stay ahead of the curve. This "bring your consumer life to work" mindset—combined with a fear of missing out and shifting workplace expectations—accelerates adoption across the enterprise.

Organizationally, slow-moving AI strategies, ambiguous policies, and a lack of psychological safety can create uncertainty around what's permissible. When outcome is prioritized over compliance, and AI enthusiasm outpaces governance, shadow AI often emerges to fill the void. The perception of harmlessness and a lack of visible consequences can create a halo effect, normalizing risky behavior and masking serious threats like data leakage, compliance failures, and tooling chaos.

> **"Shadow AI isn't a fringe issue—it's a signal that employees are moving faster than the systems designed to support them. Without trusted oversight and a coordinated architectural strategy, even a single shortcut can expose the organization to serious risk. But with the right guardrails in place, shadow AI can become a powerful force for innovation, agility, and long-term competitive advantage. The time to act is now—with clarity, trust, and bold forward-looking leadership."**
>
> —Swami Chandrasekaran, Principal, US and Global AI and Data Labs leader at KPMG

When it comes to shadow AI, organizations must address a range of significant challenges including security vulnerabilities, data leakage, exfiltration, licensing and copyright risks, tooling sprawl, erosion of trust, and potential reputational damage. The use of unauthorized AI tools and systems can expose sensitive information such as client-confidential data, intellectual property, financial records, personal identifiers, and legal documents. National security concerns can also come into play: when employees use unvetted AI platforms, there's no visibility into where that data may end up—including jurisdictions with known interests in cyber malfeasance.

In January 2025, a breach at the AI chatbot platform DeepSeek served as a stark reminder of the risks posed by unsanctioned AI use. More than 1 million records—including chat logs, API keys, and backend system details—were exposed due to misconfigured infrastructure. While the breach was not caused by employee uploads, it highlighted how easily proprietary data can be compromised when AI tools operate outside formal governance. The incident triggered regulatory scrutiny and reputational fallout, underscoring the need for centralized oversight and secure AI environments.[3]

This incident underscores the fact that shadow AI is not a theoretical risk—it is a present and growing threat. It means that organizations must act decisively to implement governance, visibility, and safe experimentation environments to ensure the next breach is not their own.

To better understand the situation, imagine a team racing to finalize a quarterly financial report. Under pressure, they upload sensitive, unreleased data into an external, unapproved AI tool

[3] "DeepSeek Cyber Attack: Timeline, Impact, and Lessons Learned," CM-Alliance.com, March 25, 2025

because it's faster and easier. It seems harmless—but that shortcut could trigger serious fallout.

The data shared with unauthorized AI tools may be stored or reused by the provider to further train or tune their models—leading to uncontrolled data exfiltration. If the AI generates content based on licensed or copyrighted material, it can introduce intellectual property and compliance risks. When this behavior is multiplied across departments using a patchwork of tools, it results in tooling chaos and AI sprawl—leaving organizations with little visibility and no centralized governance.

These types of incidents can erode trust internally and, if exposed, cause significant reputational damage. What began as a productivity boost can quickly escalate into a brand and regulatory crisis.

And while "vibe coding" can bring concepts to life quickly, especially for nondevelopers, when done through unsanctioned tools, it can introduce serious risks, including undocumented and unreviewed code, client confidential data being used, hidden vulnerabilities, and untraceable logic. In the context of shadow AI, this can amplify AI sprawl, bad software design, data leakage, and compliance gaps. Organizations should support vibe coding within secure, approved environments to enable innovation without sacrificing accountability.

While shadow AI contains inherent risks, organizations should push themselves to create safe environments that allow for fast, creative experimentation. Doing so can create real opportunities, including:

Accelerated experimentation: Employees can rapidly test new ideas and workflows without waiting on approvals, surfacing high-impact use cases organically.

Increased engagement and empowerment: When employees feel trusted to explore new tools, it fuels motivation, creativity, and retention—especially among digital natives.

Real-time insight into system gaps: The tools people choose reveal friction points in existing AI systems—highlighting where modernization and upgrades are required.

Tool discovery for future investment: Shadow AI surfaces emerging tools that can inform enterprise-scale adoption and innovation strategy.

Organic upskilling: Hands-on use of a variety of GenAI tools builds prompting, goal automation, and critical thinking skills—boosting AI fluency on the job.

Grassroots innovation: Employees often discover novel use cases that top-down strategies might overlook, helping organizations stay ahead of emerging trends.

Cross-functional collaboration: Shadow AI often emerges in problem-solving teams that span business and technical roles, fostering innovation across silos.

Pressure-testing enterprise AI strategy: This can reveal where official AI systems, tools, policies, or platforms are falling short—providing real-world feedback to refine enterprise AI roadmaps.

By recognizing and channeling these benefits into sanctioned programs, organizations can turn shadow AI from a liability into a strategic advantage. This is also important, because as innovation continues to thrive with the use of shadow AI, organizations must remain vigilant about potential enterprise IP leakage.

Strategies for managing and harnessing shadow AI

To transform shadow AI from a liability into a dynamic catalyst for creativity and innovation, organizations should strive to move beyond merely reactive controls and embrace a proactive, structured strategy. This means not only mitigating risks but also enabling safe experimentation, guided by clear oversight and intentional design.

01. Establish a central AI transformation organization:

Create a dedicated transformation organization or office to govern AI strategy, architecture, tools, trust, and standards across the enterprise. This group should coordinate policy, trust, infrastructure, and innovation initiatives while enabling safe and effective experimentation.

02. Create an AI Technology Review SteerCo:

Form a cross-functional governance body to evaluate, approve, and monitor AI platforms, tools, and systems. This board should ensure alignment with enterprise standards, security protocols, legal, and trust guidelines.

03. Develop clear AI governance policies:

Effective governance strategies are necessary to address shadow AI risks. Organizations must implement policies and processes to prevent data loss, ensure compliance, and align AI use with organizational policies. This includes establishing stringent guidelines for AI applications and educating employees on the importance of using only corporate-approved AI tools.

04. Stand up an AI labs function:

Provide a sandboxed environment where teams can safely explore new AI platforms and tools, as well as test ideas and validate use cases before scaling. Labs reduce the need for unsanctioned experimentation while accelerating innovation.

05. Implement guardrails for responsible AI use:

Deploy internal AI experiments and pilots, secure development environments, and automated audits to minimize risk. Balance oversight with usability to encourage adoption of approved tools.

06. Promote employee education and awareness:

Train teams on AI risk, trusted AI principles, security in the age of AI leading practices, and data handling. Include scenario-based examples that clarify gray areas and make responsible use intuitive.

07. Leverage AI system inventory tools:

Use discovery platforms to identify and catalog AI systems used within the organization. Also scan and discover shadow AI tools use, map existing AI assets, and close visibility gaps. Pair this with identity and access controls to manage exposure.

08. Enable "choose your own AI" within approved boundaries:

Offer a curated menu of vetted GenAI tools tailored to different roles and use cases (e.g., copilots for developers, content tools for marketing, agentic platforms for builders). While employees can receive their choice, IT retains oversight.

09. Promote a culture of transparency and innovation:

Make it safe for employees to share how they're using GenAI. Reward responsible innovation and engage teams early in piloting new tools or workflows.

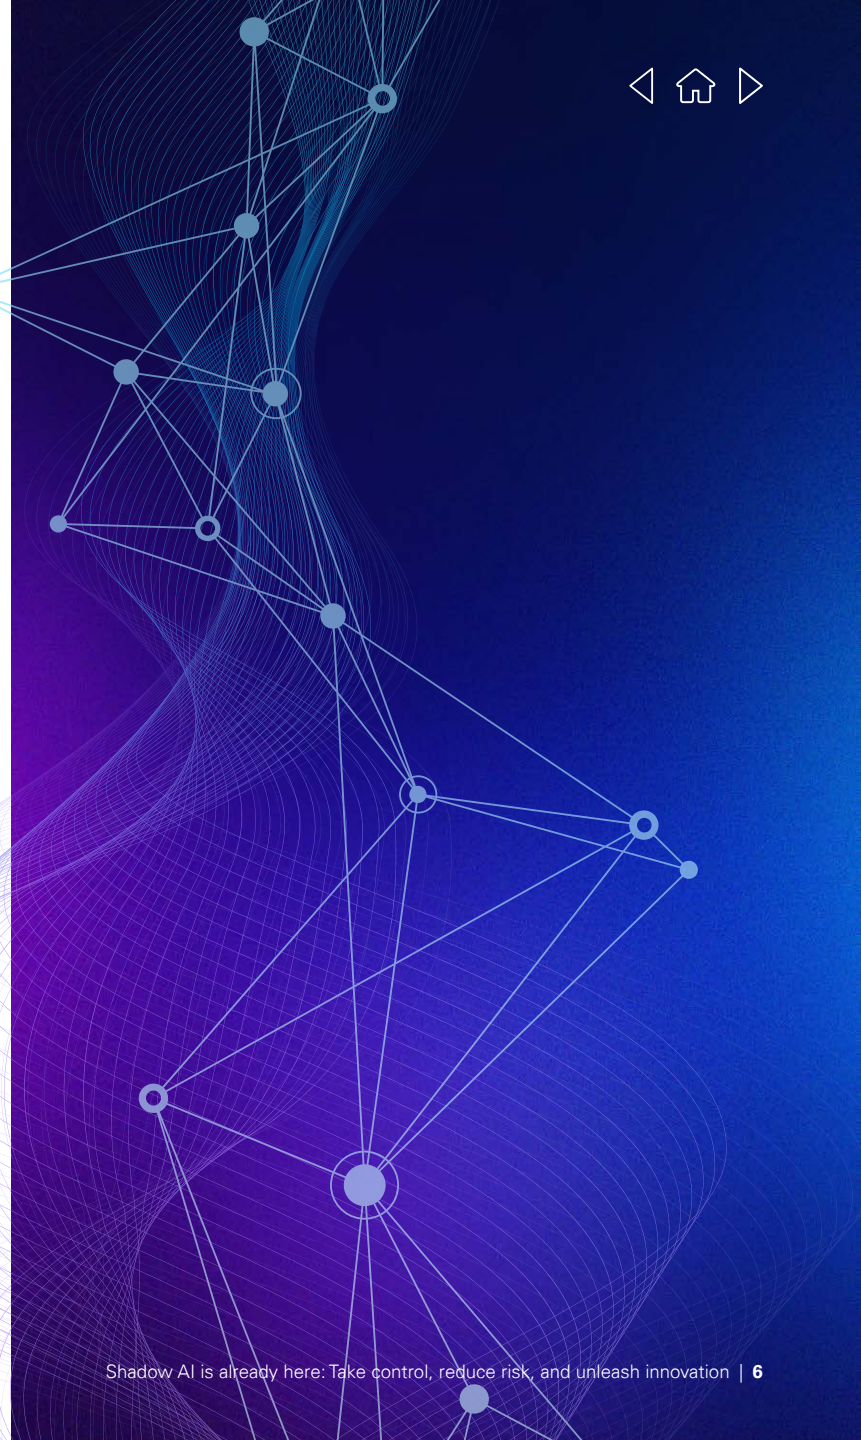

From liability to a potential engine of innovation

Shadow AI is no longer a fringe issue—it's a mainstream reality inside most organizations. While the risks are real—ranging from security breaches to compliance failures—it also represents a powerful signal: Your workforce is ready, willing, and already experimenting with AI to solve real problems. Organizations that respond with rigid control will stifle innovation. Those that respond with clear strategy, empowered oversight, and curated choice will unlock AI's full potential—safely and at scale. By establishing a central AI function, codifying governance, embedding secure guardrails, and embracing a "choose your own AI" approach, organizations can shift from policing AI to enabling it responsibly. The opportunity is clear: turn shadow AI from a liability into an engine of distributed innovation. The time to act is now—before shadow AI becomes unmanageable, or worse, irreversible.

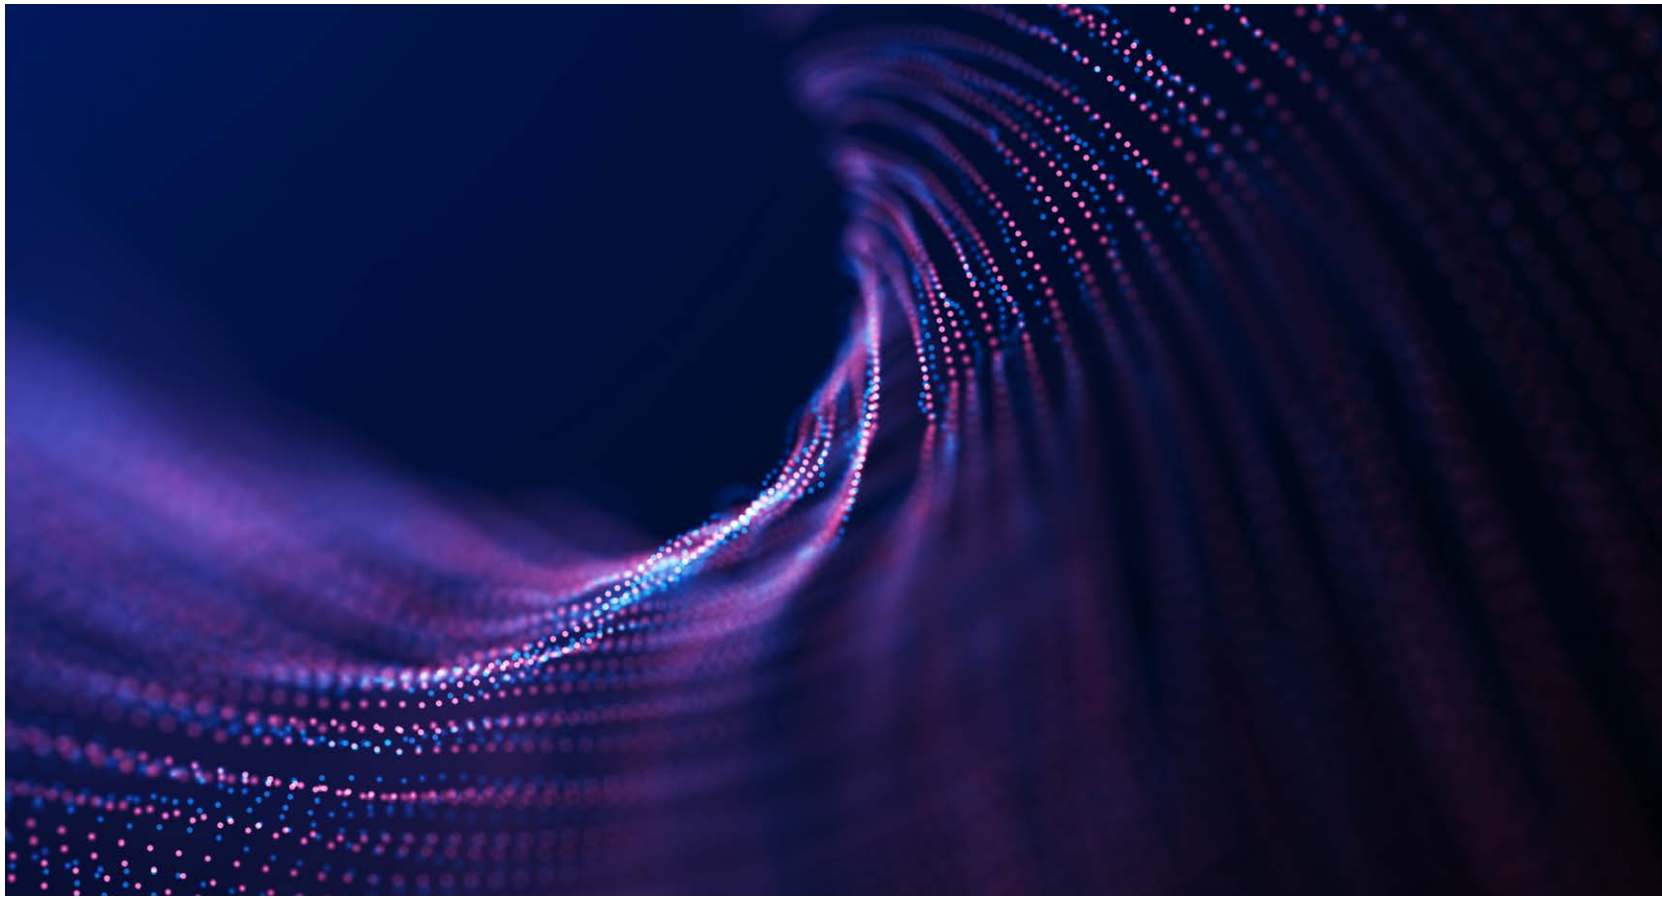

How KPMG can help

KPMG AI Trust Services helps organizations turn shadow AI from a risk into a strategic asset through governance, technology, and workforce transformation using the following tools and solutions:

Trusted AI governance

We help design centralized AI governance programs that align with enterprise goals. This includes:

- Establishing AI policies and review boards
- Ensuring compliance with ethical and regulatory standards
- Developing agentic AI guidelines and policies

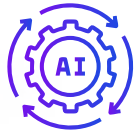

Scalable technology and inventory solutions

We support the discovery, oversight, and secure deployment of AI tools by:

- Implementing AI inventory and governance systems
- Automating AI testing and compliance reporting
- Creating secure environments for AI experimentation and scaling

AI intake and monitoring

To manage AI effectively, we:

- Build streamlined intake processes for new AI tools
- Implement automated discovery and monitoring systems
- Offer curated, role-specific AI tools within governance boundaries

AI-enabled workforce

We help organizations prepare their people for AI by:

- Integrating AI agents into workflows
- Redefining roles and governance to support AI collaboration
- Delivering training to foster a culture of responsible AI use

Authors

Bryan McGowan

Global and US Trusted AI Leader

E: bmcgowan@kpmg.com

Bryan spearheads the KPMG Trusted AI initiatives within Advisory, working across diverse sectors to drive AI innovation and implementation. He plays a key role in advancing the firm's strategic AI initiatives by integrating AI systems into business processes, and enhancing value delivery for clients worldwide. Bryan oversees a range of projects including the design of AI inventory and monitoring systems, the implementation of data-driven strategies, and the deployment of AI governance programs. As an advocate for responsible AI, he leads efforts to ensure ethical standards and scalable practices are embedded as the foundation in AI programs.

Swami Chandrasekaran

Global Head of AI & Data Labs

E: swamchan@kpmg.com

Swami leads and executes the firm's AI strategy across Tax, Audit, Advisory, and other functions, serving 200,000 knowledge professionals worldwide. He directs and oversees various R&D efforts and initiatives covering AI architecture, advanced knowledge assistants, AI agents, domain-tuned small language models (SLMs), synthetic data, enterprise discovery and search, and hardware-optimized solutions. He also chairs the KPMG AI Technology Review Board to help ensure trusted and scalable AI adoption.

We would like to thank our contributors:

Prasad Jayaraman and Aisha Tahirkheli

Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.

Please visit us: **kpmg.com** **Subscribe**