Shadow AI, Data, and Governance: How Rapidly Growing and Mid-Sized Companies Protect Their Business Know-How

What shadow AI puts at risk, why bans fail, and the practical governance that lets you keep the productivity and close the leak

This article brings together key questions and answers from my executive briefings, conference speaking, participants in my training programs and additional research addressing shadow AI, data security, private AI deployment, employee usage policies, and the governance of AI-driven decisions.

James Proctor, The Inteq Group | November 2026

KEY TAKEAWAYS

1. Shadow AI is the use of AI tools outside organizational approval and oversight. It can be well intentioned, widespread, and difficult for leadership to see.
2. The larger risk isn't only confidential data. It's know-how: the pricing logic, exception handling, and customer judgment that make your company competitive.
3. Restrictions alone may not prevent unauthorized AI use. Combine clear rules, approved tools, employee education, and proportionate technical controls with secure options for sensitive work.
4. Properly configured private AI can increase control over data, documents, and decision logic. Keeping information within approved boundaries requires deliberate architecture, security controls, and restrictions on external integrations.
5. AI governance for rapidly growing and mid-sized companies should be practical and proportionate: business ownership, documented rules, human escalation, auditable decisions, access controls, and regular review.

What is shadow AI?

Shadow AI manifests as employees using AI tools the company hasn't approved or set up, usually free or personal accounts of popular chat assistants. It is almost always well intended. People are trying to write proposals faster, summarize contracts, or analyze a spreadsheet.

Shadow AI is widespread. Microsoft and LinkedIn's 2024 Work Trend Index found that 75% of surveyed knowledge workers used generative AI at work, and 78% of those users brought their own tools. Gartner reported in 2025 that 69% of surveyed cybersecurity leaders had observed or suspected use of prohibited public generative AI.

More recently, PagerDuty's 2026 survey of 1,250 office professionals at large-revenue organizations found that 66% had used AI at work despite believing it was not permitted by company policy; 34% reported entering customer information into public AI tools. Protiviti's 2026 survey found that 68% of medium-sized organizations had only partial or no visibility into employees' AI tool usage. These surveys cover different populations and should not be interpreted as a trend series.

The productivity isn't the concern. The concern is that pricing logic, customer details, contracts, and internal know-how may be going to an outside service under terms nobody at the company has read.

What is at risk?

The obvious risk is confidential data leaving your control: customer records, contract terms, personnel matters. The less obvious risk is your know-how. When someone pastes in a pricing model, a proposal strategy, or the reasoning behind how you handle a difficult account, they are sharing the judgment that makes the company competitive. That information may be retained, reviewed, or used under terms you didn't negotiate, and it may breach commitments you've made to customers about how their data is handled.

IBM's 2025 Cost of a Data Breach study examined 600 organizations that had experienced breaches. Approximately one in five reported a breach involving shadow AI. Organizations with high levels of shadow AI incurred average breach costs about $670,000 greater than organizations with little or no shadow AI.

Across the breached organizations studied, 63% lacked an AI governance policy or were still developing one. Gartner separately forecast that, by 2030, more than 40% of organizations would experience security or compliance incidents related to unauthorized AI tools; that is a projection, not an observed incident rate.

None of this requires bad intent. It happens one helpful shortcut at a time, which is why the answer is a clear policy and a safe alternative, not blame.

Shadow AI leaks know-how one helpful shortcut at a time.

Do AI providers train on the data employees enter into the chat?

It depends on the provider, product tier, account settings, and applicable agreements; these terms can change. Consumer AI services may allow content to be used for model improvement subject to their settings, while business and enterprise offerings commonly provide stronger contractual protections. OpenAI, for example, states that content from its business offerings and API is not used to train models by default. Organizations should verify the current data-use terms for each service rather than assuming all accounts are treated the same.

Training is not the only question. Terms also cover how long data is retained, whether provider staff can review it, and where it is stored. And if employees use personal accounts, the company's business agreement doesn't apply to that use at all. Read the current terms for every tool your people actually use, and treat the answers as an intellectual property question that belongs in the risk register, not only in procurement.

Should you ban public AI tools?

Bans alone may not eliminate unauthorized AI use, particularly when employees find significant productivity benefits in tools they cannot access through approved channels. Samsung restricted generative AI tools in 2023 following reports that employees submitted sensitive code to an external AI service. Restrictions may be necessary for certain data or settings, but they are more effective

when combined with secure alternatives, clear communication, employee education, and technical controls.

A better approach channels use rather than depending on prohibitions alone. Provide an approved option for everyday work, define the information that cannot be entered into outside tools, and offer a properly secured option for sensitive data and core know-how. These steps can reduce risk, but monitoring and enforcement remain important. Start by understanding what tools employees actually use through appropriate surveys and technical discovery; explain that the goal is safer work, not punishment.

How do you write an AI usage policy without a compliance department?

Keep the employee-facing policy short enough that people will read it. One or two pages can cover the essentials, supported as needed by more detailed security, privacy, vendor approval, incident reporting, and records-retention procedures:

- **Approved tools:** which tools and account types employees may use.
- **Never share:** the categories of information that never go into an outside tool, such as customer data, pricing, contracts, and personnel matters.
- **Check the output:** a person stays responsible for anything AI helped produce.
- **Ask first:** where to go with a new tool or use case.
- **Private option:** how to handle sensitive work safely.

Assign one owner, review the policy every six months, and explain the reason behind each rule. People follow rules they understand.

What does private AI keep inside your walls?

When an open-weight model is deployed within infrastructure the organization controls, the company can exercise greater control over where data is processed, stored, and accessed. This can include business data, documents, decision rules, and records of agent activity. A properly configured private open-weight deployment can keep sensitive information within approved security boundaries without sending prompts to external model providers, but this is not automatic. Network controls, access management, integration restrictions, logging, and scrutiny of external services are necessary.

Private open-weight deployment can also help an organization preserve knowledge developed as agents are designed and refined. Documented rules, validated examples, corrections, evaluation results, and process improvements become reusable organizational assets when systematically captured and maintained; they do not automatically retrain the underlying model. Organizations should also review external providers' contractual terms, retention practices, training policies, and intellectual-property provisions. Properly configured enterprise services may provide meaningful protections, while private deployment can offer greater direct control over sensitive workflows.

What does AI governance look like for rapidly growing and mid-sized companies?

Practical and proportionate. It doesn't need a governance office or an enterprise framework. It needs a few things done consistently:

- **An owner for each agentic process,** who answers for its decisions.

- **Written decision rules and limits,** stating what each agent may and may not do.
- **Escalation to people,** with clear thresholds and named roles.
- **A decision log** that ties every action to its basis.
- **Regular review,** sampling decisions and adjusting the rules.
- **A usage policy** for everyday AI tools.

In IBM's 2025 breach study, 97% of organizations experiencing AI-related security incidents reported lacking appropriate AI access controls. Access governance therefore matters alongside process governance. The free NIST AI Risk Management Framework provides a scalable reference organized around Govern, Map, Measure, and Manage. Effective controls are easier to design into the process from the start than to add after deployment.

Is it safe to let AI agents make decisions?

AI agents can make mistakes, as people do. Safe decision-making depends on clearly defined authority limits, validated decision criteria, required supporting evidence, and escalation procedures. Cases involving missing information, conflicting evidence, unusual conditions, or decisions outside approved boundaries should go to a person. Confidence scores can provide an additional signal when properly evaluated and calibrated, but should not be the sole basis for autonomy.

Initial deployments should begin with limited decision authority and human review where additional judgment is needed. The scope of autonomous decisions should expand only after testing demonstrates reliable performance on representative transactions, including exceptions and higher-risk cases. Ongoing monitoring should confirm that the agent continues to operate within its approved boundaries.

How do you audit the decisions an AI agent makes?

Design an audit trail into the agentic process from the beginning. Record the decision made, the relevant evidence and business rules, actions taken, and any escalation to a person, including the outcome. Where available and validated, confidence indicators may also be recorded. Protect confidential information in audit records through access controls, appropriate retention, and data minimization. Regular review of decisions, especially exceptions and higher-risk cases, helps identify performance issues and improve decision rules.

Governance without prohibition

Rapidly growing and mid-sized companies need not choose between unrestricted AI use and blanket prohibition. Start by understanding what is in use, publish a concise policy, provide approved tools and secure options for sensitive work, and incorporate proportionate governance into agentic processes from the outset. Consistent safeguards can help protect know-how, improve reliability, and demonstrate accountability to customers, lenders, and insurers.

Questions answered in this paper

Each question below has a short, direct answer in Shadow AI, Data Security, and AI Governance: 10 Key Questions Answered for Rapidly Growing and Mid-Sized Businesses.

- What is shadow AI, and is it happening in our company?

- What's actually at risk when employees paste company information into ChatGPT?
- Do AI providers use the data our employees put in to train their models?
- What about our data? Where does it go, and is it secure?
- What does private or on-premises AI actually keep inside our walls?
- Should we just ban public AI tools?
- How do we write an AI usage policy without a compliance department?
- What does AI governance look like for a mid-sized company?
- Is it safe to let AI make decisions? What happens when it gets something wrong?
- How do we audit the decisions an AI agent makes?

Sources

1. Microsoft and LinkedIn, 2024 Work Trend Index (May 8, 2024). https://blogs.microsoft.com/blog/2024/05/08/microsoft-and-linkedin-release-the-2024-work-trend-index-on-the-state-of-ai-at-work/
2. Gartner, "Gartner Identifies Critical GenAI Blind Spots" (November 19, 2025). https://www.gartner.com/en/newsroom/press-releases/2025-11-19-gartner-identifies-critical-genai-blind-spots-that-cios-must-urgently-address0
3. IBM, "Cost of a Data Breach Report 2025" (July 30, 2025). https://newsroom.ibm.com/2025-07-30-ibm-report-13-of-organizations-reported-breaches-of-ai-models-or-applications%2C-97-of-which-reported-lacking-proper-ai-access-controls
4. OpenAI, "How your data is used to improve model performance". https://help.openai.com/en/articles/5722486-how-your-data-is-used-to-improve-model-performance
5. The Register, "Samsung bans generative AI after data leak" (May 2, 2023). https://www.theregister.com/2023/05/02/samsung_generative_ai_ban/
6. NIST, AI Risk Management Framework (AI RMF 1.0). https://www.nist.gov/itl/ai-risk-management-framework
7. Protiviti, "AI Pulse Survey, Vol. 4: No Visibility, No Confidence" (May 2026). https://www.protiviti.com/sites/default/files/2026-05/aipulse26-vol4-survey-booklet-0426-iz-en-protiviti.pdf
8. PagerDuty, "Shadow AI Survey" (June 11, 2026). https://investor.pagerduty.com/news/news-details/2026/PagerDuty-Report-Finds-Two-Thirds-66-of-Office-Professionals-Have-Used-Unauthorized-AI-Tools-at-Work/default.aspx

About The Inteq Group

The Inteq Group helps organizations strengthen business performance through business process reengineering, business systems analysis, and agentic AI solution analysis.

inteqgroup.com | 800-719-4627 | info@inteqgroup.com