olakai

WHITE PAPER

Shadow AI: The Hidden Risk and the Opportunity

Why unauthorized AI isn't a security failure — it's a measurement failure. And why the organizations that respond to both signals win.

June 2026 · Published by Olakai Research · olakai.ai

EXECUTIVE SUMMARY

Shadow AI — the use of AI tools outside of official enterprise procurement and approval — is growing faster than any IT department can monitor with conventional controls. But the most important thing about shadow AI is not that it is a threat. It is that it is a signal. The employees using unapproved tools are telling you exactly where AI creates enough value to justify circumventing IT process. Organizations that act on both the risk and the signal outperform those that address only one.

Key findings:

78% of knowledge workers report using AI tools that were not provided or approved by their employers, according to Microsoft's 2025 Work Trend Index.[1] The compliance exposure is material: consumer-grade AI tools have no enterprise data handling commitments, and employees routinely submit PII, financial projections, legal materials, and source code into interfaces with no contractual protections. At the same time, shadow AI is a remarkably precise adoption signal — it identifies the workflows where AI assistance is creating enough real value that employees will go outside sanctioned channels to get it. This paper gives CISOs, CFOs, and boards a framework to address the risk and capture the signal simultaneously.

What Shadow AI Actually Is

Shadow AI is the enterprise version of a phenomenon IT teams have managed for decades: shadow IT. But where shadow IT typically meant running an unapproved SaaS tool for project management or file sharing, shadow AI carries a qualitatively different risk profile — because AI tools do not just process data, they ingest it, generate outputs from it, and in many cases, retain it to improve their models. The surface area for data exposure is larger, the tooling changes faster than procurement cycles can keep pace with, and the employee demand is structurally stronger than it has ever been for any prior category of shadow technology.

The tools involved range widely: ChatGPT and Google Gemini accessed through personal accounts, Claude.ai or Perplexity on individual subscriptions, purpose-built AI tools for legal, finance, or sales work deployed without procurement review, and at the high end, employees with engineering backgrounds using direct API access to build personal workflows on top of foundation models. The common thread is not the tool category. It is the gap between the AI assistance employees need to do their jobs well and the AI assistance their organization has officially made available to them. When that gap exists — and for most enterprises it does, across most roles — shadow AI fills it, regardless of IT policy.

The Risk Side: Data Exposure, Compliance, and Liability

The Data Exposure Problem

The most immediate risk from shadow AI is data exposure. When an employee pastes a contract summary into a consumer ChatGPT interface, drafts a board memo in a personal Claude.ai account, or uses an unapproved AI writing tool to process client correspondence, the data they submit enters an environment with no enterprise data handling agreement, no contractual prohibition on model training use, and no audit trail. In most cases, the employee does not know this. They see a useful tool and use it. The CISO finds out later — if at all.

The data types that flow through unauthorized AI interfaces are not low-risk: they include personally identifiable information (PII) about customers and employees, nonpublic financial projections and board materials, legal analyses and privileged communications, M&A due diligence materials, and proprietary source code. IBM's 2024 Cost of a Data Breach Report found that the average breach now costs enterprises $4.88 million — a figure that does not capture the full cost of regulatory penalties, reputational damage, or the loss of competitive advantage when proprietary materials are exposed through AI interfaces with inadequate data handling controls.[2]

The asymmetry of this risk is important for boards to understand. Employees who use shadow AI tools rarely do so with malicious intent. They are trying to be productive, often in the absence of adequate sanctioned alternatives. But intent does not matter in a regulatory audit, and it does not limit liability in a client data agreement. The organization's exposure is the same whether the data left through an accident or a deliberate act.

> **The training data risk:** Many consumer AI tools — including some that position themselves as enterprise-friendly — have default data handling settings that permit use of submitted content for model improvement. Even when opt-out mechanisms exist, they require affirmative action by the user. Employees who access AI tools through personal accounts, or who use tools without reviewing enterprise data terms, may be unknowingly contributing proprietary materials to the training pipelines of external AI systems.

The Compliance and Regulatory Exposure

For regulated industries, shadow AI creates a second layer of risk beyond data exposure: compliance violation. Healthcare organizations operating under HIPAA cannot have patient information passing through AI tools without a Business Associate Agreement in place — and consumer AI tools do not offer BAAs. Financial services firms operating under SEC recordkeeping requirements, MiFID II obligations in Europe, or FINRA rules face similar constraints: AI-assisted communications must be captured, retained, and available for audit in ways that consumer tools do not support.

The EU AI Act, now in enforcement as of 2025, adds a classification dimension: enterprises operating in EU markets must assess the risk tier of every AI system in use and apply appropriate governance controls. An AI tool deployed without procurement review has, by definition, not been assessed. That gap is a compliance finding waiting to happen. In the United States, the FTC's expanded enforcement posture on data mishandling and the SEC's guidance on AI use in financial services create parallel exposure for US-based organizations.

78%

of knowledge workers use AI tools not provided or approved by their employer (Microsoft Work Trend Index, 2025)[1]

$4.88M

average cost of a data breach in 2024 — before regulatory penalties and reputational impact (IBM Security)[2]

25%

of organizations have comprehensive visibility into shadow AI tool usage across their workforce[1]

The Liability Gap: What You Cannot See, You Cannot Govern

The deepest structural problem with shadow AI is not any specific tool or incident — it is the visibility gap that makes proactive governance impossible. An organization cannot apply data handling policies to AI tools it does not know employees are using, cannot assess compliance risk against a tool that was never submitted for procurement review, and cannot conduct a post-incident audit of data flows through a system with no enterprise logging.

Forrester's research on enterprise AI adoption found that organizations consistently underestimate unauthorized AI tool use by a factor of three to five when relying on self-reporting or network proxy logs alone — because employees using web-based AI interfaces on managed devices generate traffic that is often indistinguishable from ordinary browser use, and employees using personal devices generate no detectable signal at all.[3] The consequence is a liability gap: the organization is responsible for all data employees process in their work capacity, regardless of whether that processing occurred on approved systems — but cannot document its governance posture with respect to tools it cannot see. In a regulatory review or litigation discovery process, that gap is not just an operational problem; it is a legal one.

> **The audit question your legal team is preparing for:** "Did you have controls in place to prevent employees from submitting client data to unauthorized AI systems?" If your answer depends on employees following an AI policy they may not have read, you do not have controls. You have a policy.

The Opportunity Side: Shadow AI as an Adoption Signal

Here is what the risk-only framing misses: the employees who bring their own AI tools to work are not the problem. They are the leading indicator. They are showing you, at zero research cost, which workflows benefit enough from AI assistance that employees will circumvent institutional friction to get it. That is exactly the intelligence a Head of AI or CFO needs to make defensible investment decisions about where to expand sanctioned AI access.

Consider what shadow AI usage patterns tell you at the workflow level. A team of sales managers using personal Claude.ai accounts to process call transcripts and draft follow-up emails is telling you that AI assistance in the post-call workflow creates enough value to pay for it personally. A group of financial analysts using personal ChatGPT to summarize analyst reports is telling you that AI-assisted research synthesis is a high-leverage task in their daily workflow. A legal team using AI writing tools outside of procurement to accelerate contract review is telling you that AI-assisted legal work is ready for enterprise deployment — the evidence is already in the usage data.

The conventional IT response to shadow AI is to detect it and block it. The strategically superior response is to detect it, classify it by risk level, and treat the low-and-medium-risk usage as a procurement roadmap. The tools employees are adopting organically are the ones where AI-assistance is creating real, day-to-day productivity value — as opposed to the tools IT or the AI team assumed would be valuable based on vendor demos and analyst reports. Organic adoption is a more reliable signal than a proof-of-concept pilot, because it is self-selecting: employees who do not find a tool useful stop using it. The ones who go through the effort of setting up personal accounts and working around corporate controls are the ones who found real value.

THE TWO-QUESTION FRAMEWORK

For each shadow AI tool detected, ask: (1) **Does it create data exposure risk at the usage volume observed?** If yes, address the risk vector immediately — either sanction the tool with proper data controls, or block and replace with an equivalent that meets data handling requirements. (2) **Does it signal a workflow where enterprise AI could create measurable ROI?** If yes, add it to the sanctioned tool roadmap and measure the value once it is under management.

Organizations that treat shadow AI as purely a threat to eliminate will spend compliance resources blocking tools and writing policies that employees route around. Organizations that treat shadow AI as both a threat to manage and a signal to act on will use the same detection infrastructure to build a more productive, more compliant AI stack — because they are making procurement decisions on the basis of real usage data, not assumptions.

5–10×

typical ratio of actual shadow AI tools in use vs. tools on approved list (Forrester research)[3]

76%

of enterprise organizations report shadow AI usage increased in 2025 — up from 52% in 2024[4]

3.6×

total shareholder return advantage for enterprises that act on AI measurement data vs. those that don't (BCG)[5]

The Detection and Response Framework

Addressing shadow AI effectively requires a three-phase approach: discover what is in use, assess the risk profile of each tool and usage pattern, and respond — either by sanctioning the tool under proper governance or by replacing it with a compliant alternative that meets the same workflow need. The goal is not zero shadow AI. That is an unachievable and counterproductive target. The goal is complete visibility, accurate risk assessment, and a governance posture that can be demonstrated to regulators, auditors, and board members.

PHASE 1

Discover: Build a Complete AI Tool Inventory

Effective shadow AI discovery requires visibility at the usage level, not the network level. Network proxy logs catch web-based tool usage on managed devices; they miss personal device access, native desktop AI applications, and API-level usage by developers. A comprehensive shadow AI picture requires combining network-layer detection with endpoint telemetry and identity-level usage analytics. The goal is a complete AI tool inventory: every tool in active use, by which teams, at what volume, and with what data. Most organizations discover their shadow AI footprint is three to five times larger than assumed — and that the largest concentrations of usage are in the highest-risk data categories: customer data, financial data, and legal materials. Visibility alone is a governance win: you cannot address what you cannot see, and most organizations have been ungovernable with respect to AI tools not because their policies are wrong, but because their detection is absent.

PHASE 2

Assess: Classify Every Tool by Risk and Signal Value

Not all shadow AI tools carry the same risk. A team using an approved AI writing assistant through a personal account instead of the corporate license is a license management issue, not a data security incident. A team submitting customer PII to a consumer AI interface with no enterprise data agreement is a compliance exposure. The risk classification needs to happen at the usage-pattern level, not just the tool level: the same tool can present different risk profiles depending on what data types flow through it, how frequently, and in what volume. Alongside risk classification, assess the opportunity signal: which tool categories appear across multiple teams or functions? Which workflows show high adoption of a specific tool type? Those are the inputs to your next-generation AI procurement roadmap — built on actual demand rather than vendor analysis. The assessment phase typically produces three outputs: a prioritized risk remediation list, a list of tools ready for rapid sanctioning with proper data controls, and a procurement roadmap for AI tool investment based on observed demand.

PHASE 3

Respond: Govern, Sanction, or Replace

Response follows one of three paths. For tools with material data handling risk and no compliant equivalent, replace: procure an enterprise-grade tool that meets the same workflow need under proper data governance. For tools with manageable risk that employees are accessing through personal accounts, sanction: add proper licensing and data controls, often at lower cost than the informal subscription spend already occurring. For genuinely low-risk tools with limited data exposure, acknowledge: formally note the tool is in use and meets baseline requirements. This third category is smaller than most CISOs expect, but acknowledging it builds credibility with the business — the governance program is managing risk proportionally, not blocking AI use categorically.

Role Takeaways

Shadow AI creates different problems and opportunities depending on your seat at the table. The following takeaways are designed to be actionable within 30 days for each role.

CISO

- Commission a shadow AI discovery audit — network + endpoint + identity layer — and expect your footprint to be 5-10x larger than your approved list
- Classify each shadow tool by data type exposure and regulatory obligation — PII, financial, legal, and code require different treatment
- Replace the blanket-block reflex with a risk-tiered response: sanction manageable tools, replace high-risk ones with compliant equivalents
- Build a DLP policy specific to AI interfaces — prompt-level filtering is different from traditional file-transfer DLP
- Establish an audit trail for all AI tool usage that touches regulated data — your next compliance review will ask for it

CFO

- Shadow AI spend is already in your budget — hidden in expense reports, personal subscriptions, and API billing that doesn't surface until month-end
- Use shadow AI usage data as a capital allocation signal: which tools have high organic adoption in which teams? Those are your highest-confidence AI investments
- Quantify the license waste in your approved AI stack — seats unused while employees pay personally for workarounds represent a consolidation opportunity
- Model the cost of a single data breach against the cost of proper shadow AI visibility infrastructure — the insurance math typically resolves quickly
- Ask for a shadow AI report in your next quarterly AI review — if the Head of AI cannot produce one, you have a measurement gap to close

BOARD

- Ask management whether the organization has complete visibility into AI tool usage across the workforce — not just the approved stack
- Ensure that "AI governance" includes shadow AI detection, not just policy for approved tools — policy without detection is not governance
- Request a data breach risk assessment that specifically addresses AI interface exposure — this is a distinct risk category from traditional data handling controls
- Treat shadow AI measurement capability as a board-level governance question, on par with cybersecurity posture reporting
- Recognize that the organizations outperforming peers on AI value creation are those using detection data to make better procurement decisions — this is a strategic capability, not just a compliance cost

Conclusion: The Measurement Response

The enterprises that get shadow AI right are not the ones with the most restrictive AI policies. They are the ones with the most complete measurement infrastructure. They know what their employees are using because they have built the detection capability to see it — not to punish employees for exercising judgment about productivity tools, but to respond to what the data is telling them: where AI helps, where it creates risk, and where the official AI stack is not meeting real workforce demand.

Shadow AI is, at its core, a measurement failure before it is a security failure. Employees turn to unauthorized tools when the legitimate AI procurement process is too slow, too restrictive, or too disconnected from what they actually need to do their jobs. The measurement response to that problem is not faster blocks or more aggressive IT enforcement. It is better visibility: a complete picture of AI tool usage across the organization

that makes it possible to act quickly — to sanction what is safe, replace what is risky, and use the adoption signal to build an AI stack that employees actually want to use through official channels.

The data is already there. The employees who are using shadow AI are generating it every day. The organizations that can read that signal — and act on both its risk dimension and its opportunity dimension — will build better AI governance programs and better AI investment portfolios than those who cannot. Measurement is the mechanism. It is also the competitive advantage.

> **The one question to answer before your next board meeting:** Do you have a complete, real-time picture of every AI tool your employees are using — approved and unapproved — and can you demonstrate that you are acting on what that picture tells you about risk and about value creation opportunity? If the answer is no, Olakai's Assistive IQ is the fastest path to yes. Visit **olakai.ai/assistive-iq** or request a private demo link at **olakai.ai**.

About Olakai

Olakai is the enterprise AI analytics and governance platform. It gives organizations complete visibility into AI tool usage across coding tools, assistive AI, and autonomous agents — enabling measurement of ROI, detection of shadow AI and data exposure risk, and control of costs across every tool, every team, and every model. Olakai's Assistive IQ pillar is purpose-built for shadow AI discovery and governance. Learn more at olakai.ai.

References

1. Microsoft Corporation. *Work Trend Index 2025: AI at Work Is Here. Now Comes the Hard Part.* Microsoft WorkLab, 2025. Available at microsoft.com/worklab.
2. IBM Security. *Cost of a Data Breach Report 2024.* IBM Corporation, 2024. Average global cost of a data breach: $4.88 million. Available at ibm.com/security/data-breach.
3. Forrester Research. *Enterprise AI Governance and Shadow AI: 2025 Benchmark Survey.* Forrester, 2025. Data reflects enterprise survey of 800+ IT and security decision-makers across North America and Europe.
4. Olakai Research. *Shadow AI Statistics 2026: The Governance Crisis Is Already Here.* Based on aggregated platform telemetry and primary research, June 2026. olakai.ai/blog/shadow-ai-statistics-2026.
5. Boston Consulting Group. *The Widening AI Value Gap.* BCG, September 2025. Survey of 1,250 global enterprises across 12 industries. Available at bcg.com.
6. European Commission. *Regulation (EU) 2024/1689 of the European Parliament and of the Council (EU AI Act).* Enforcement phased in from August 2024 through 2026.