RESEARCH
Influence and insight through social media

August 2026

THE HIDDEN AI RISK

Most IT Teams Never See Coming

WHITE PAPER

Prepared by
Zeus Kerravala

ZK Research
A Division of
Kerravala Consulting

INTRODUCTION

Artificial intelligence (AI) is entering small and mid-sized organizations through everyday tools, not just through formal transformation programs or centralized platform decisions. The use of unsanctioned chatbots, mobile apps, browser extensions, AI meeting assistants, and software-as-a-service (SaaS) copilots that employees connect on their own, often well before IT has established a governance model for them, is referred to as shadow AI. As a result of this trend, AI risk is an operational reality for many organizations rather than a future issue.

That urgency is evident in the data. In the 2026 Mobile AI Security Report, organizations surveyed by ZK Research estimated that employees use an average of 27 AI-powered applications, and 14% said more than 50 tools are already in use (Exhibit 1). At the same time, 44% of employee AI interactions were estimated to occur on mobile devices, and 63% of organizations reported investigating incidents in which generative AI contributed to attempted or successful theft or leakage of sensitive data. For lean IT teams, the question is no longer whether AI is entering the environment but whether the organization can create enough visibility and control to guide it responsibly.

Exhibit 1: Widespread Adoption of AI Challenges IT and Security

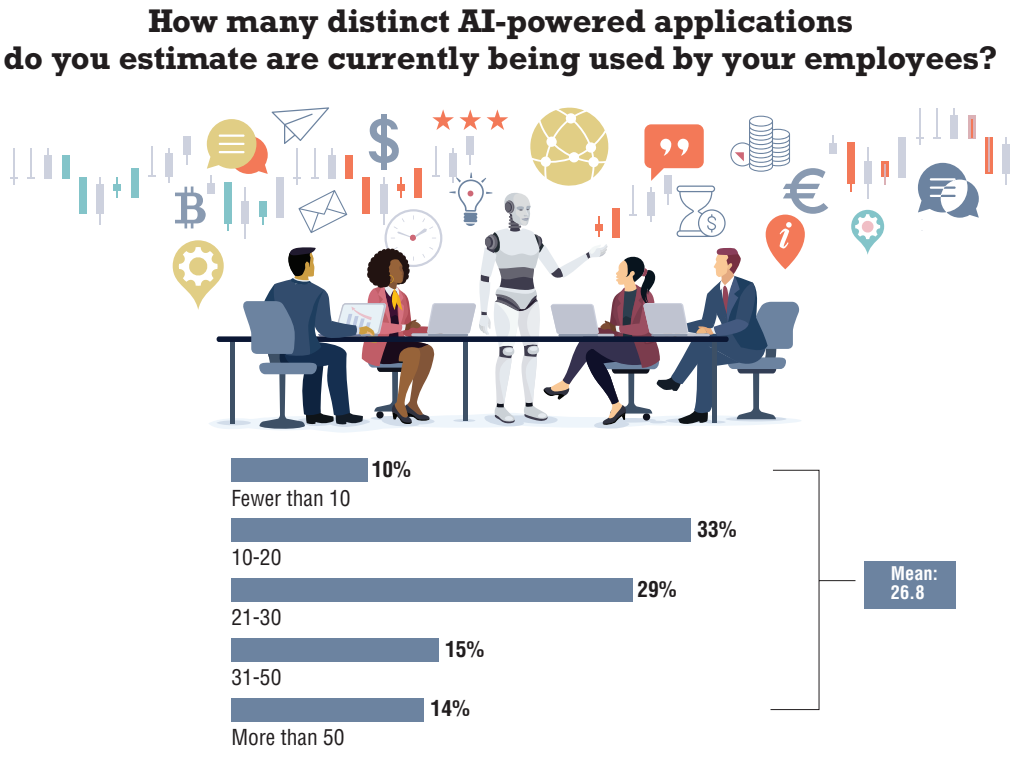

Note: The total does not equal 100% due to rounding.

ZK Research 2026 Mobile AI Security Report

ABOUT THE AUTHOR

Zeus Kerravala is the founder and principal analyst with ZK Research. Kerravala provides tactical advice and strategic guidance to help his clients in both the current business climate and the long term. He delivers research and insight to the following constituents: end-user IT and network managers; vendors of IT hardware, software and services; and members of the financial community looking to invest in the companies that he covers.

Only 36% of respondents said they have full real-time visibility into the specific data uploaded to unsanctioned AI tools.

WHY SHADOW AI IS A DAILY IT ISSUE

For small and mid-sized organizations, shadow AI does not typically start with a large internal AI initiative. It starts when employees use public large language models (LLMs) to draft documents, connect AI agents to SaaS applications, rely on meeting recorders to summarize calls, or access AI tools on personal or unmanaged mobile devices. This pattern makes shadow AI harder to govern than earlier forms of shadow IT because it combines rapid experimentation with direct access to sensitive business data.

The AI visibility challenge is especially important because many organizations appear more confident than their control posture warrants. While 95% of our respondents reported having full or partial visibility into the organization's AI usage, only 36% said they have full real-time visibility into the specific data uploaded to unsanctioned AI tools. That gap matters because effective AI governance depends less on knowing that a tool exists than on understanding what data flows into it, who authorized access to it, and what permissions an AI agent may have inherited along the way.

UNDERSTANDING THE EXPOSURE

Lean IT teams should begin by assuming shadow AI exposure already exists across devices, applications, and workflows. The most practical first step is to build an inventory using existing tools, such as firewall logs, web gateways, cloud access security broker (CASB) platforms, mobile threat defense (MTD) or mobile device management (MDM) consoles, and SaaS administration dashboards. Even a basic first-pass inventory can help an organization group AI usage into manageable categories, such as public chatbots, browser extensions, mobile apps, meeting assistants, and SaaS-connected AI agents.

From there, the next step is to identify the highest-risk data flows. That means focusing on where employees may copy, upload, or expose regulated data, financial data, intellectual property, customer records, or source code to external AI tools. This is also where shadow permissions become dangerous: AI agents connected to systems such as Microsoft OneDrive, Jira, Salesforce, Slack, or Zoom may inherit broad user access originally granted for convenience rather than reviewed through a security lens.

Several findings from our survey indicate where lean teams should look first. As noted earlier, roughly 44% of AI interactions were said to occur on mobile devices. Organizations also estimated that more than a third of their core SaaS tools have been integrated with third-party AI agents without formal IT review (Exhibit 2), and 89% of leaders reported concern about shadow permissions. In other words, the problem is not limited to employees visiting a public chatbot site from a desktop browser; it now spans mobile operating systems, embedded AI features, SaaS connectors, and agentic workflows that often fall outside traditional security review cycles.

Exhibit 2: The Use of Shadow AI Is Rampant in Organizations

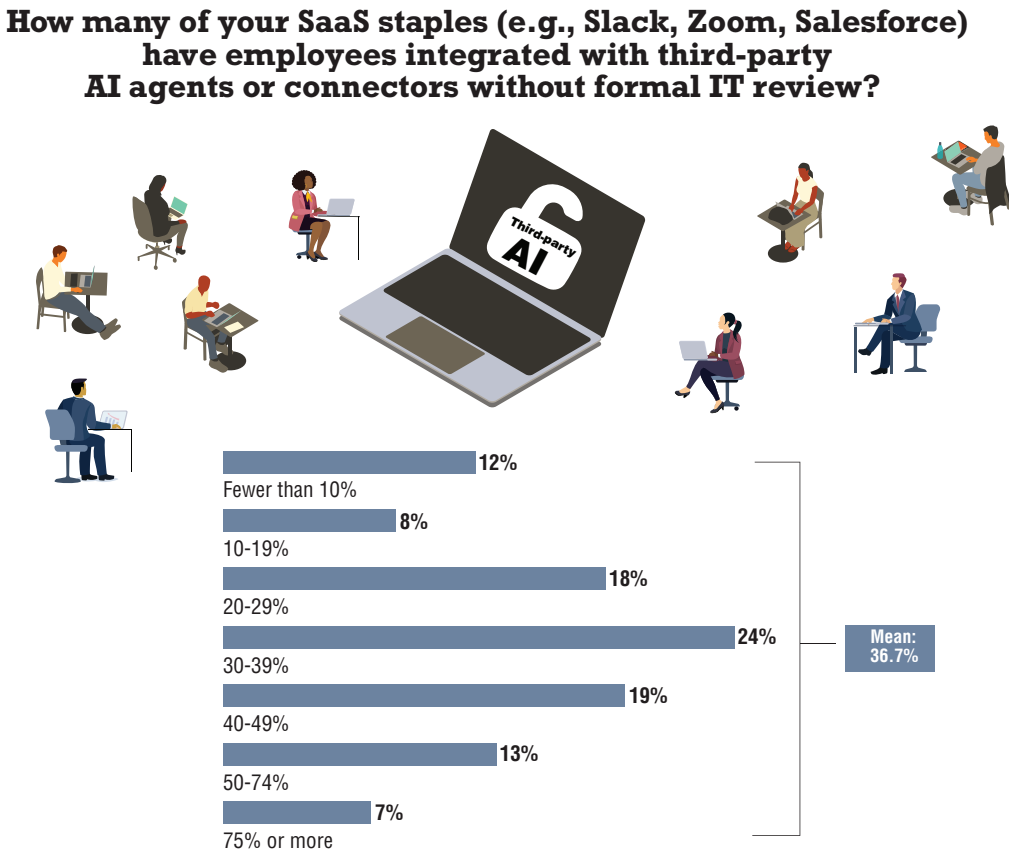

Note: The total does not equal 100% due to rounding.

ZK Research 2026 Mobile AI Security Report

THE FIVE BIGGEST RISKS FOR SMALL AND MID-SIZED ORGANIZATIONS

1. **Fragmented visibility:** Shadow AI often spans managed laptops, unmanaged personal phones, sanctioned SaaS platforms, and third-party extensions, making it difficult for smaller teams to maintain a unified operational picture. When usage is distributed so broadly, incident response slows because IT may know AI is being used but not which users, data, or workflow introduced the exposure.

2. **Data leakage and IP loss:** Employees often adopt AI to improve productivity, but that convenience can lead them to paste sensitive data into public tools or to authorize AI agents to access systems in a broader manner than intended. That concern is reflected in our survey findings, which show that AI-generated malware and intellectual property theft are the top shadow AI concerns for security leaders.

AI governance must be treated as an operating discipline, not a one-time project.

3. **Mobile and SaaS sprawl:** AI usage is no longer confined to the corporate desktop, and having system-level AI built into mobile operating systems introduces a new class of monitoring and policy challenges. Respondents also indicated they feel least prepared to defend against system-level AI that processes corporate data locally on devices, suggesting that many teams still have a blind spot regarding mobile AI.

4. **Governance friction:** Nearly all surveyed organizations have a formal AI acceptable-use policy, but only 47% primarily enforce it with AI-aware security tools, while many still rely on legacy filtering or employee honor systems. For lean IT teams, that gap between policy and enforcement creates a dangerous middle ground where leadership assumes AI is governed, but technical controls have not kept pace with employee behavior.

5. **Capacity and budget:** Our survey respondents said their organizations are allocating an average of 19% of their 2026 security budget to AI governance, visibility, and compliance solutions, and 42% reported allocating more than 20%. For smaller organizations, those numbers reinforce a practical truth: AI governance must be treated as an operating discipline, not a one-time project. However, many lean teams do not have the headcount to sustain that effort alone.

FIVE STEPS IT LEADERS SHOULD TAKE NOW

1. **Establish a lightweight AI governance process.** Small and mid-sized organizations do not need a large formal committee on day one, but they do need clear ownership, a basic acceptable-use policy, and clear rules regarding which data can and cannot be used with public or third-party AI tools. The most effective starting point is often a small, cross-functional working group that includes IT, security, and a limited set of business stakeholders who can make decisions quickly.

2. **Create a living inventory of AI tools, integrations, and data flows.** This need not begin as a major platform deployment; a disciplined spreadsheet or lightweight system of record can suffice if it tracks which AI tools are in use, who owns them, what data they touch, and the SaaS permissions they require. The key shift is to treat AI usage as an asset inventory and risk register rather than as isolated exceptions discovered only after a problem occurs.

3. **Prioritize visibility into mobile and bring your own device (BYOD).** Because nearly half of AI interactions were estimated by our respondents to occur on mobile devices, organizations should tune their MTD or MDM controls to identify AI-related traffic, risky app behavior, operating system–level AI use, and data movement between enterprise apps and external assistants. This area deserves immediate attention because only 20% of our

Exhibit 3: Five Steps IT Leaders Should Take Now

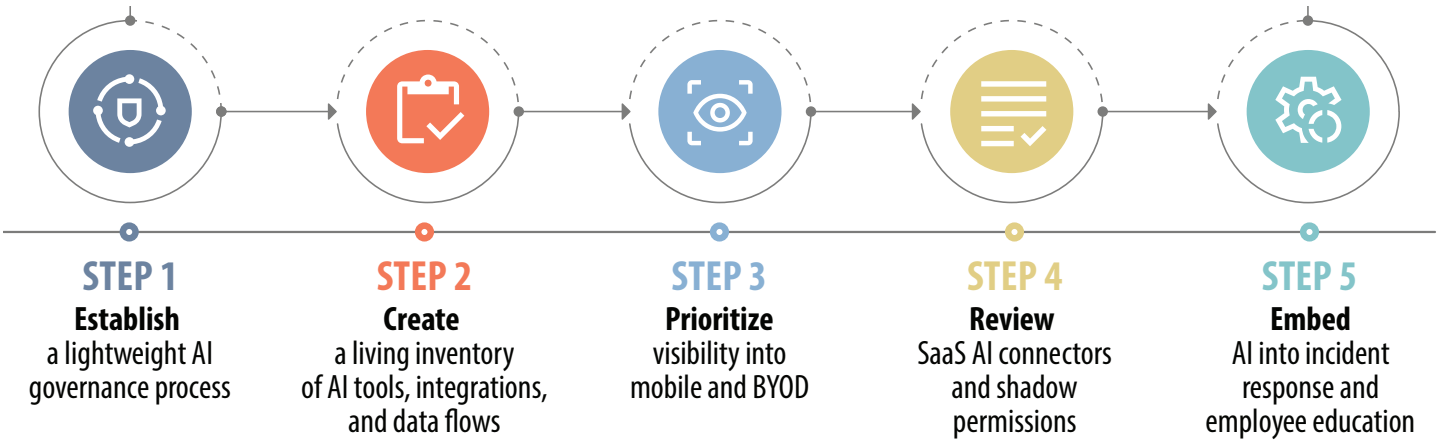

ZK Research, 2026

respondents reported having granular visibility into data flows between enterprise apps and third-party AI assistants through their existing mobile controls.

4. **Review SaaS AI connectors and shadow permissions.** Every organization should maintain a lightweight approval process for connecting AI agents or copilots to business applications, with particular attention to high-value systems such as productivity suites, customer relationship management (CRM) platforms, ticketing systems, collaboration tools, and document repositories. Access rights should be reviewed periodically to ensure that inherited access does not quietly expand beyond business needs.

5. **Embed AI into incident response and employee education.** Although 88% of organizations in our survey said AI data leaks are addressed within incident response, only 54% said they have a dedicated AI data leak playbook. Lean teams should develop at least one short runbook for scenarios such as sensitive data pasted into a public LLM, an unsanctioned AI meeting assistant capturing confidential discussions, or a SaaS AI agent accessing data beyond its intended scope.

WHY ONGOING MANAGEMENT MATTERS

Shadow AI is not a one-and-done remediation exercise. New AI tools appear constantly, employees change workflows quickly, and SaaS vendors continue to embed AI into products that may already be widely deployed across the business. As a result, any governance model that relies on a single policy rollout or a one-time cleanup will fall behind almost immediately.

This is why ongoing monitoring and periodic governance reviews are so important for lean IT teams. Organizations need a repeatable cadence to review new AI use cases, reassess data exposure, check shadow permissions, update policies, and revisit which controls are actually working.

The goal is not to stop AI adoption but to establish enough operational discipline so the business can benefit from AI without continuously introducing unmanaged risk.

HOW DATAPRISE SUPPORTS LEAN IT TEAMS

For many small and mid-sized organizations, the challenge is not recognizing the need for AI governance but finding the capacity to build it while also running the day-to-day IT environment. That is where Dataprise can fit in as a managed services and advisory partner focused on helping organizations create structure for AI adoption.

Dataprise's AI Readiness Assessment helps leadership teams evaluate their current posture across strategy and governance, data enablement, technology and tools, operating model, and workforce capability. The engagement is intended to produce a clear maturity baseline, identify gaps and risks, prioritize use cases, and build a practical 12- to 24-month roadmap that leadership can use to sequence actions.

Dataprise's AI Strategy & Governance Advisory offering extends that work by providing executive-level guidance without requiring a full-time internal AI leader. The service is built around defining an AI vision, establishing governance, creating a roadmap, and supporting executive alignment through recurring advisory sessions, steering committee participation, and board-ready briefings. For lean IT teams, this model can offer a pragmatic way to move from ad hoc AI reactions to a more disciplined operating framework.

CONCLUSION AND RECOMMENDATIONS

Shadow AI has already become part of the operating environment for small and mid-sized organizations, especially through mobile devices, SaaS platforms, and employee-led experimentation. The key lesson for IT leaders is that the real risk is not merely the presence of AI but the combination of weak visibility, unmanaged data flows, and inherited permissions that can expose sensitive business information.

The right response is practical rather than theoretical. Lean IT teams should formalize a basic AI use policy, maintain an up-to-date inventory of tools and connectors, strengthen visibility into mobile and SaaS, review permissions more rigorously, and add AI-specific scenarios to incident response and user education. These steps will not eliminate all shadow AI risks, but they will create the operational foundation needed to reduce exposure while supporting responsible adoption.

Organizations that lack the internal bandwidth to sustain this work should consider outside help that combines managed services with AI readiness and governance advisory services. A partner such as Dataprise can help smaller IT teams establish a maturity baseline, define governance, create a roadmap, and maintain the ongoing review cycle that shadow AI now requires.

CONTACT

zeus@zkresearch.com

Cell: 301-775-7447

Office: 978-252-5314