nudge

AI Adoption in Practice

What enterprise usage data reveals
about risk and governance

Introduction: AI adoption has outpaced security governance

Over the past two years, AI adoption across the enterprise has accelerated faster than any prior technology shift. What began as employee experimentation with general-purpose chat tools has evolved into widespread use of AI across every corner of the digital workplace: meetings, documents, code, customer support, and internal automation. Today, AI systems are not just responding to conversational prompts—they are embedded into workflows, integrated with core business platforms, and increasingly capable of taking autonomous action on behalf of or as delegated by individual employees.

This shift is occurring at the same time cybersecurity teams are grappling with a broader breakdown of traditional security boundaries. As modern work—and corporate data along with it—has migrated to SaaS and now AI-native app environments, perimeter-based controls struggle to maintain visibility and control. AI compounds these challenges by introducing new data flows, new integration patterns, and new behavioral risks that are often invisible to legacy security tooling.

AI governance has emerged as a top priority for security and risk leaders, but many programs remain too narrowly focused on model-level risks, establishing AI gateways, or masking sensitive data from prompts. While necessary, these controls alone are insufficient. The most consequential AI risks now stem from how employees actually use AI tools day to day—what data they share, which systems AI is connected to, and how deeply AI is embedded into other tools and operational workflows.

What makes AI governance uniquely challenging is not model behavior in isolation, but how AI tools intersect with existing SaaS ecosystems and data. AI is now a *data amplifier*: it accelerates how quickly information moves across systems, how broadly it can be shared, and how easily human judgment can be bypassed. Understanding these intersections—between people, permissions, and platforms—is the foundation of effective AI security.

This report examines real-world AI adoption and usage patterns observed across real enterprise environments to help security leaders ground AI security and governance decisions in empirical data. By understanding where AI is present, how it is used, and what data flows into these systems, organizations can move from theoretical AI risk to practical, enforceable governance.

About this report

This report is based on anonymized and aggregated telemetry collected across Nudge Security customer environments. The data reflects observed AI vendor presence, integrations, browser extensions, and user interactions with AI tools during the measurement period.

Rather than relying on surveys or self-reported usage, this analysis is grounded in direct observation of AI activity within enterprise SaaS environments. This approach provides a more accurate view of how AI is actually being adopted and used in practice, including tools that may not be centrally procured or formally approved.

To protect customer privacy and confidentiality:

- Our analysis used only anonymized and aggregated data
- No company- or user-identifying information was analyzed
- Metrics are expressed as percentages rather than absolute counts
- No specific company identities can be inferred from the results
- No individual user behavior can be inferred from the results

The findings are intended to highlight directional trends and common patterns, not to represent any single organization. Adoption rates reflect the presence of AI tools within environments as of the first week of January of 2026. Adoption rates do not reflect the extent or quality of usage, unless otherwise noted. Prompt behavior data was analyzed for the preceding year (2025).

Key findings

Usage of core LLM providers is nearly ubiquitous.

OpenAI is present in 96.0% of organizations, with Anthropic at 77.8%.

The most-used AI tools are diversifying beyond chat.

Meeting intelligence (Otter.ai at 74.2%, Read.ai at 62.5%), presentations (Gamma at 52.8%), coding (Cursor at 48.4%), and voice (ElevenLabs at 45.2%) are now widely present.

Agentic tooling is emerging.

Agent tools like Manus (22%), Lindy (11%), and Agent.ai (8%) are establishing an early footprint.

Integrations are prevalent and varied.

OpenAI and Anthropic are most commonly integrated with the organization's productivity suite, as well as knowledge management systems, code repositories, and other tools.

Usage is concentrated.

Among the most active chat tools observed, OpenAI accounts for 67% of prompt volume.

Data egress via prompts is non-trivial.

17% of prompts include copy/paste and/or file upload activity, and 72% of uploads come from local files.

Sensitive data risks skew toward secrets.

Detected sensitive-data events are led by secrets and credentials (47.9%), followed by financial information (36.3%) and health-related data (15.8%).

In summary

AI adoption is no longer experimental. The data suggests AI is now embedded in everyday workflows—browsers, meeting tools, and developer environments—and connected to sensitive systems like email, documents, code, and tickets. Effective AI governance must therefore focus on data flows, integrations, and employee behaviors, not just vendor allowlists.

1. Vendor landscape: AI is everywhere—and increasingly specialized.

Across enterprise environments, AI usage spans general-purpose models, meeting intelligence, content generation, and developer tools. The leading vendors by adoption rate show that organizations are standardizing on a small set of core providers while simultaneously experimenting with specialized tools.

Top 10 Most-Adopted AI Tools

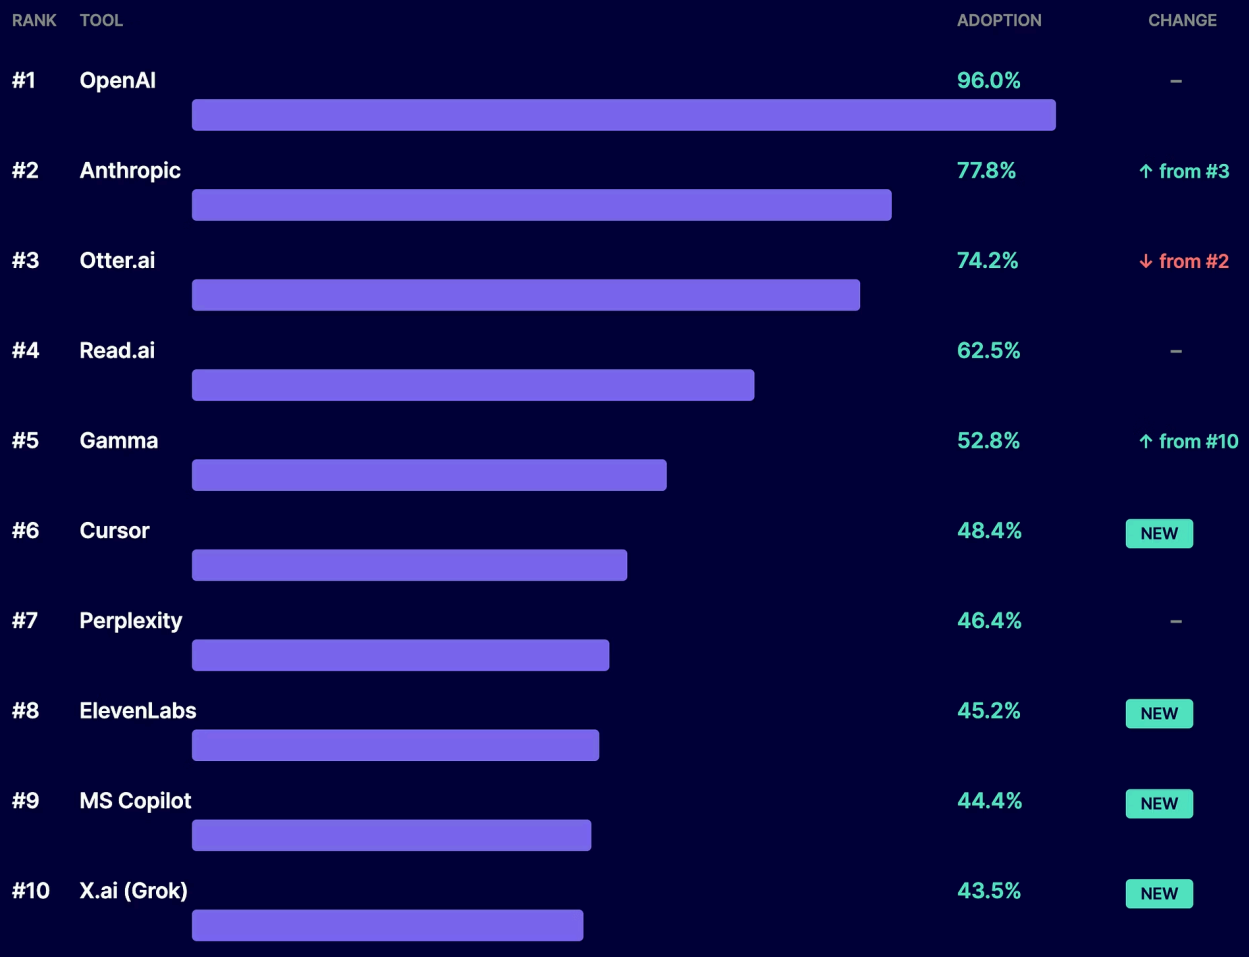

Percent of organizations that have adopted each tool as of January of 2026, and change from 1 year ago

Notable movements in the top tier

- **LLM competition is active.** Anthropic adoption is rising, demonstrating broader Claude usage across teams. Claude Code has also helped Anthropic to establish a foothold within enterprises.
- **Productivity helpers are showing strong traction.** Tools like Otter.ai, Read.ai, and Microsoft Copilot are some of the most popular tools outside of the core LLM providers.
- **AI-assisted creation is accelerating.** Presentation tooling (Gamma) is now present in more than half of organizations.
- **Developer momentum is clear.** Cursor has gained rapid adoption, breaking into the top 10 tools and showing up in nearly 50% of orgs, signaling growth in AI-native coding workflows.
- **Voice and audio AI is entering enterprise workflows.** ElevenLabs' presence suggests experimentation with narration, dubbing, and synthetic voice interfaces.
- **Niche tools show high churn.** Tools focused on narrow workflows—such as video editing, copywriting, or noise cancellation—are more likely to cycle in and out of the top tier as platforms consolidate features.

What security teams often miss: Vendor consolidation ≠ risk reduction

While AI adoption appears to be consolidating around a small number of dominant vendors, risk does not consolidate in the same way. A single LLM provider can underpin dozens of extensions, integrations, and internal workflows—each with its own permissions and data access. Focusing governance efforts solely on approved vendor lists often obscures the true sources of exposure, which live downstream in how AI tools are connected and used.

2. Emerging categories: Agents and “vibe-code” platforms

Two categories stand out as fast-moving: AI agents that can take actions across systems, and AI-native development platforms designed for rapid prototyping and application creation.

AI agents

Agent tools represent a shift from *assistive* AI (answering questions) to *agentic* AI (planning and executing tasks). This transition introduces new governance needs, including least-privilege access, approval workflows, action logging, and clear guardrails around what agents can do.

Leading AI Agents

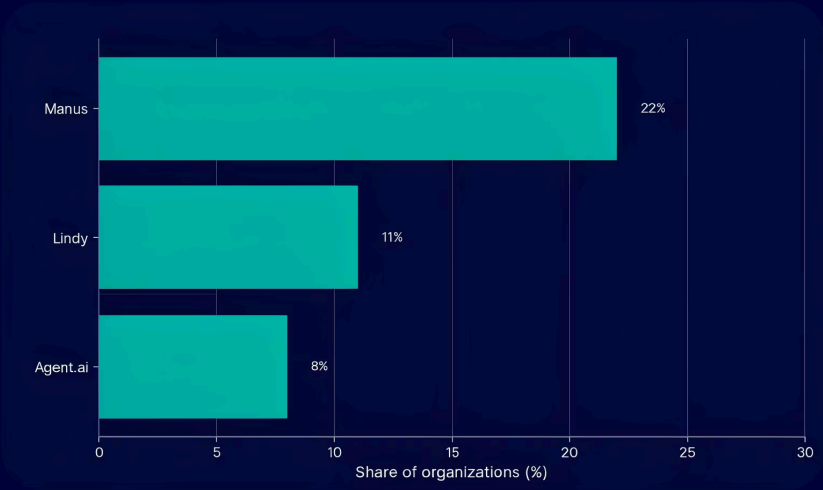

Percent of organizations using AI agents

What security teams often miss: Agents accumulate permission debt

Early AI agents are often granted broad permissions to “just make them work,” especially during experimentation. Over time, these permissions persist even as ownership, use cases, or organizational context changes. This creates permission debt: access that is no longer well understood, actively monitored, or aligned with least-privilege principles—yet capable of taking autonomous action across systems.

Agent guardrails

- Define which actions require human approval (for example, sending emails, modifying tickets, or writing to code repositories).
- Log agent actions centrally.
- Require scoped access tokens rather than shared credentials.

AI-native development and vibe-coding platforms

AI-native development environments enable non-traditional builders to create prototypes quickly and allow engineers to iterate faster. Their growing presence underscores the need to treat these tools as part of the software supply chain.

Percent of organizations using vibe code and AI-native development tools

Recommendation: Adopt a supply-chain mindset with AI-powered development tools

- Apply standard SDLC controls to AI-generated code paths.
- Encourage peer review.
- Enforce repository access boundaries.
- Scan outputs for embedded secrets before deployment.

3. AI in the browser: Where adoption meets daily work

Browser extensions are a practical signal of “AI in the flow of work” because they surface directly inside email, documents, and web applications. The leading extensions tilt toward writing assistance, screen recording, translation, and AI-powered documentation.

AI Browser Extensions

WRITING **Grammarly** 89.3% adoption	VIDEO **Loom** 81.3% adoption	DOCUMENTATION **Scribe** 57.3% adoption
TRANSLATION **Google Translate** 48% adoption	SEARCH **ChatGPT Search** 40% adoption	ASSISTANT **Claude** 34.7% adoption
TEXT-TO-SPEECH **Read Aloud** 30.7% adoption	VIDEO **Vidyard** 29.3% adoption	DOCUMENTATION **Tango** 26.7% adoption
SALES **Apollo.io** 25.3% adoption	EMAIL **Superhuman** 22.7% adoption	

Percent of organizations where each browser extension was detected

Governance note

Browser extensions are often installed without procurement or centralized review. Consider:

- Extension inventorying
- Policy controls for high-risk permissions (cookie access, network request content access)
- Clear guidance on which extensions are approved for sensitive workflows

What security teams often miss: Browser AI bypasses traditional controls

Browser-based AI tools frequently operate outside traditional security control points. They are installed directly by users, interact with sensitive content in real time, and often bypass network-based inspection or centralized approval workflows. As a result, some of the highest-risk AI interactions occur in the least-visible part of the environment: the employee's browser.

4. Integrations: AI is being wired into core business systems.

Integrations are a key maturity signal. They indicate that AI is not just used interactively, but embedded into workflows across calendars, documents, tickets, and code. This increases productivity—and expands the potential blast radius of misconfigurations or data leakage.

A misconfigured scope, compromised token, or overly permissive agent can expose entire document repositories, issue trackers, or codebases—often without triggering traditional alerts. In this model, governance failures scale at the speed of automation.

When organizations integrate AI tools, the most common endpoints are collaboration suites and work hubs, particularly Google Workspace and Microsoft 365. Developer and knowledge tools (GitHub, Notion, Atlassian, Figma) also appear frequently, especially for LLM providers.

Common Integration Endpoints

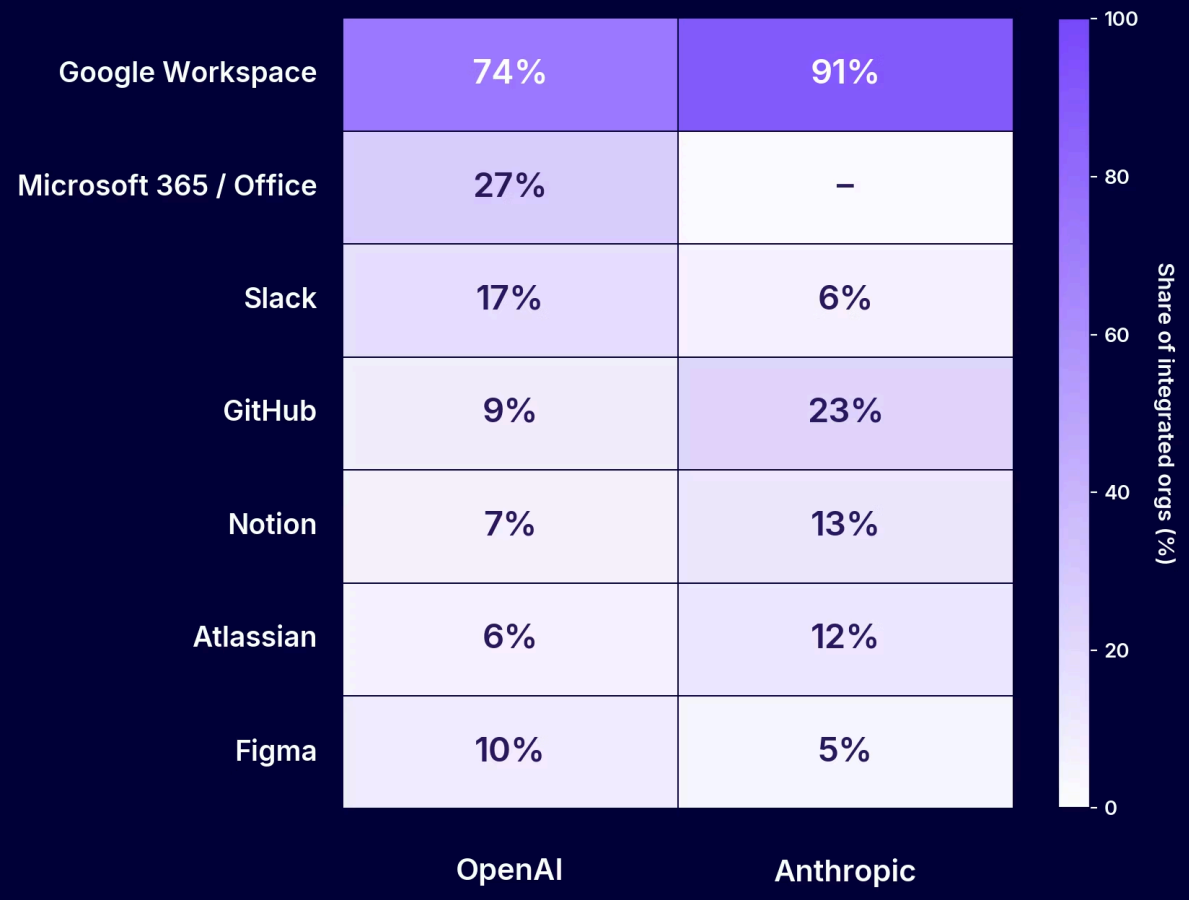

Prevalence of AI integrations with other business systems

Key observations

- **Google Workspace is the dominant endpoint.** It appears in a large majority of integrations for each leading AI tool.
- **Developer-adjacent endpoints are common for LLM providers.** GitHub, Notion, Atlassian, and Figma appear more frequently alongside LLMs than meeting tools.
- **Slack is a common integration point.** Integrating Slack with ChatGPT or Claude can transform Slack from a chat tool into a thinking + execution layer for teams.

What security teams often miss: Integrations define AI blast radius

An AI tool's real risk profile is defined less by its model and more by what it is connected to. Integrations create trusted pathways for data to travel, meaning a single misconfigured scope or compromised token can expose entire document repositories, ticketing systems, or codebases. As AI becomes more deeply integrated, blast radius—not prompt quality—becomes the dominant security concern.

Best practice

Treat AI integrations like OAuth apps:

- Require review before connecting AI tools to email, documents, ticketing, or code.
- Enforce least-privilege scopes.
- Monitor access to sensitive repositories or shared drives.
- Regularly rotate credentials used by automations.

5. Usage and data flows: Concentration, sharing, and spillover risk

Adoption answers *what is present*; usage answers *what is happening*. Two patterns stand out:

- Prompt activity is concentrated, with OpenAI accounting for 67% of prompt volume.
- A meaningful share of prompts include direct data sharing via copy/paste or file uploads.

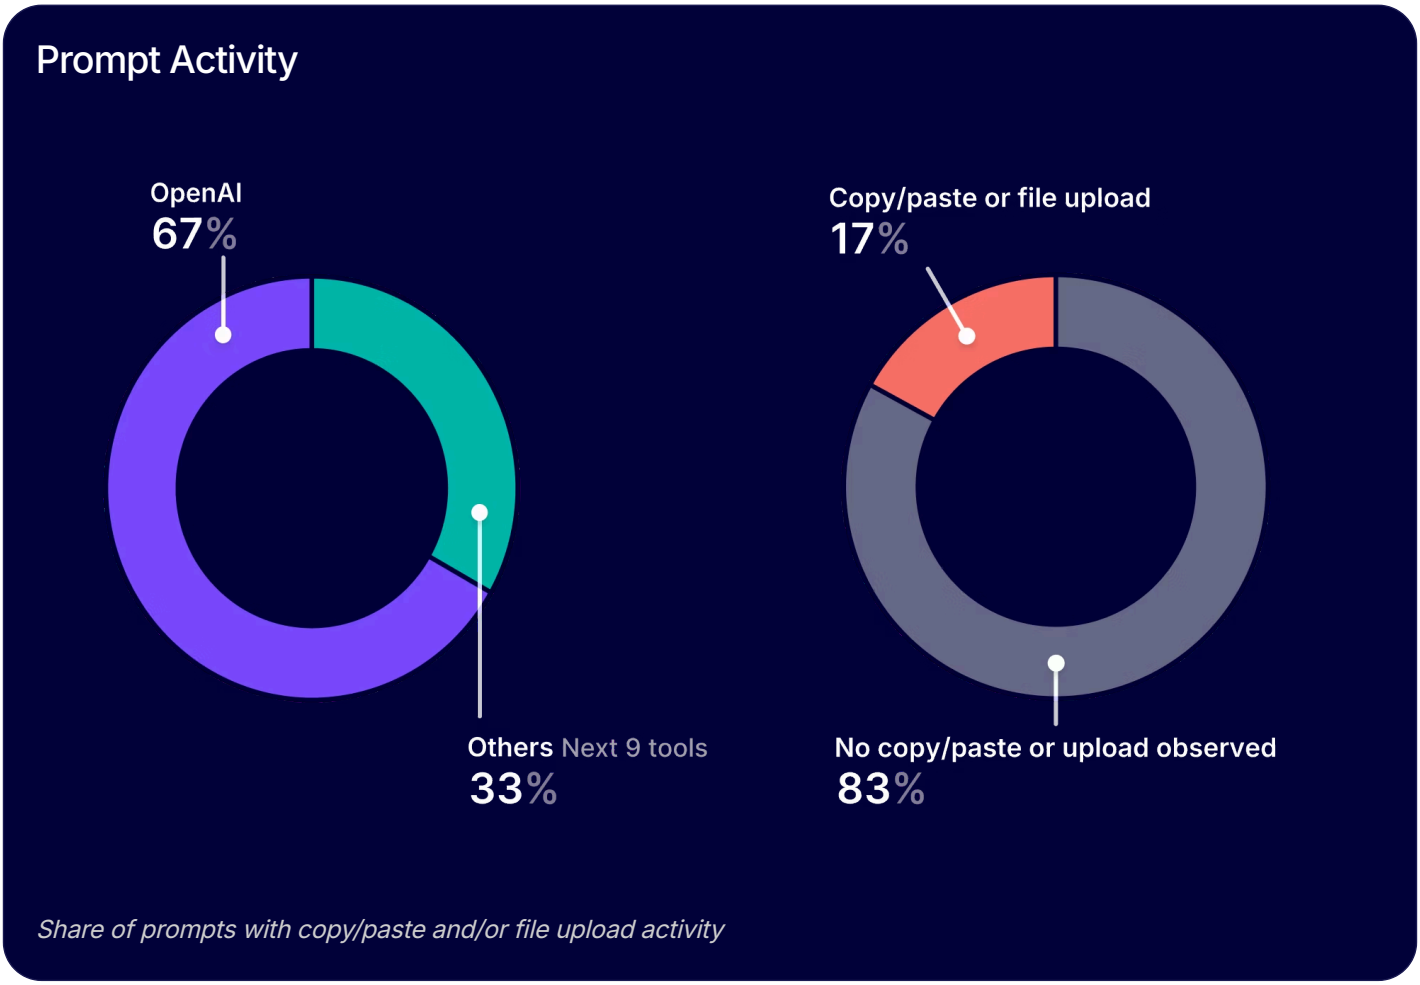

Interpretation

Concentrated usage can simplify governance by reducing the number of platforms to secure—but it also increases dependency risk. At the same time, copy/paste and file uploads are the primary pathways for sensitive data to enter AI systems, making behavior-focused controls essential.

Where data is coming from

Copy/paste behavior is a leading indicator of data egress into AI tools. File uploads are often higher risk than plain text prompts due to richer embedded data, including metadata, tables, customer exports, and attachments.

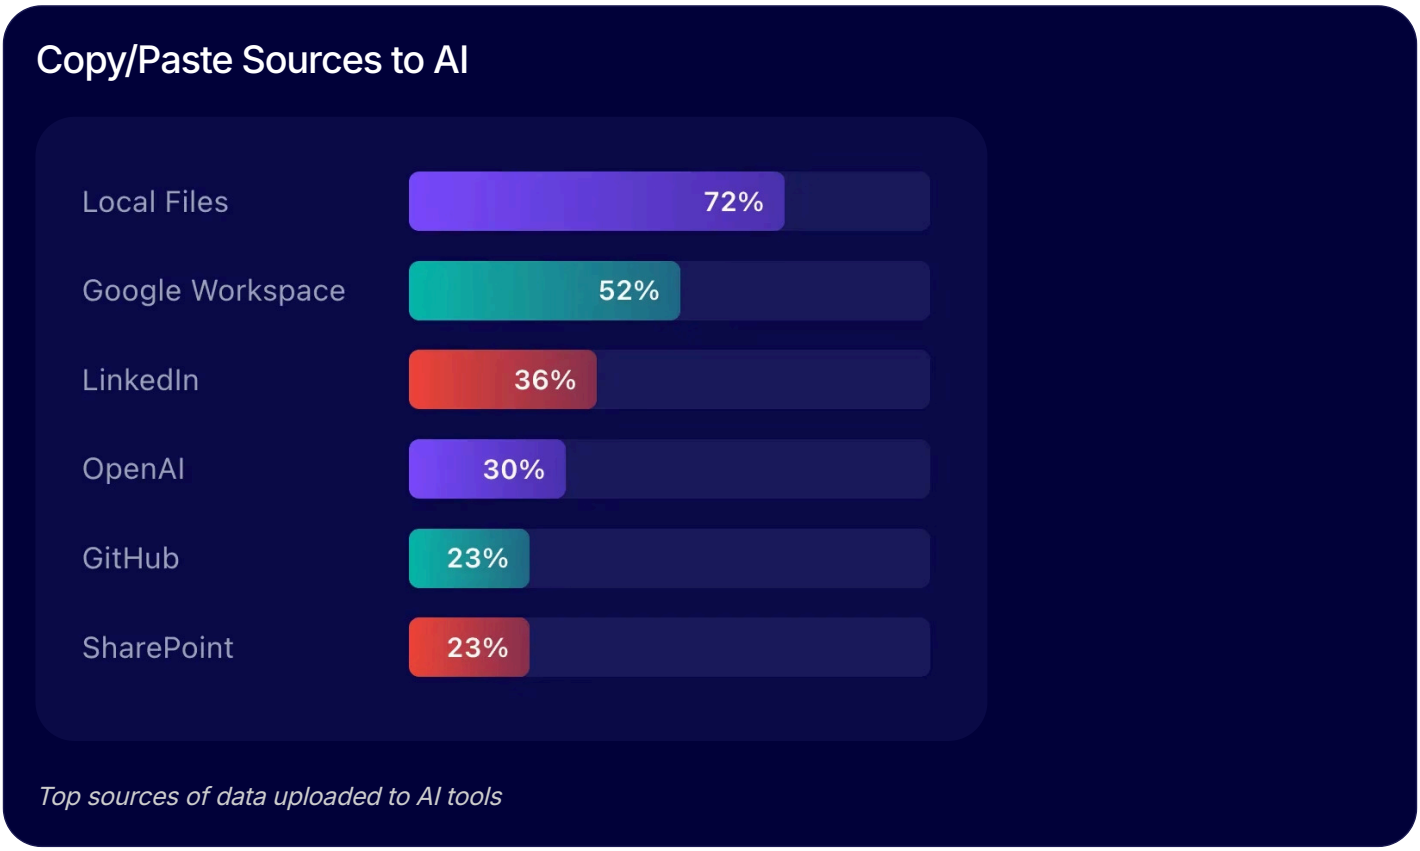

Top sources of data uploaded to AI tools

Practical controls

- Provide approved workflows for sharing documents with AI (for example, redacted templates, internal tools, or approved connectors).
- Implement just-in-time warnings when users paste from high-risk sources or attempt to upload files likely to contain sensitive content.

6. Sensitive data exposure: What shows up in AI prompts

Sensitive-data detections provide a practical lens into AI risk. In this dataset, the dominant category relates to secrets and credentials, suggesting that employees sometimes paste or upload tokens, keys, or authentication artifacts into AI tools.

Sensitive Data Detection

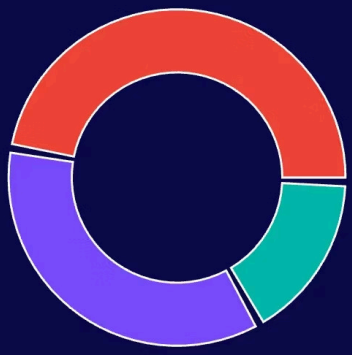

Secrets & Credentials	48%
Financial Information	36%
PHI (Health Info)	16%

Top sensitive data types: JWT tokens, webhooks, IBAN numbers, access keys, API keys

Top categories of sensitive data uploaded to AI tools

Common sensitive elements observed

- **API keys and access keys:** Often during debugging, scripting, or integration setup
- **JWT tokens and webhooks:** Frequently copied from developer tools, logs, or automation platforms
- **IBAN and other financial identifiers:** Commonly surfaced in customer support or billing workflows

These patterns reinforce a key governance lesson: AI risk cannot be managed solely at the perimeter or vendor level. Effective controls must live where decisions are made—inside browsers, chat interfaces, and developer workflows—so that users receive guidance at the moment data is about to be shared.

Control guidance

- Implement guardrails where users work, including secrets scanning, redaction prompts, and policy warnings inside chat interfaces.
- For high-risk teams, consider safer alternatives such as private LLM deployments, approved connectors with auditing, and automatic blocking of detected credentials.

What security teams often miss: Most AI risk is unintentional

The majority of sensitive data exposure in AI prompts stems from routine, well-intentioned work—not malicious behavior. Debugging code, troubleshooting integrations, summarizing documents, and resolving customer issues all create moments where secrets or sensitive data are easily shared. Effective governance must account for these everyday workflows, rather than assuming policy violations are deliberate.

7. Recommendations: A practical AI governance checklist

The following recommendations reflect a progression from visibility to control to scalable governance. Organizations do not need to implement all controls at once, but should prioritize based on data sensitivity, integration depth, and business criticality.

- **Inventory continuously.** Track AI tools, browser extensions, agents, and integrations as living assets—not one-time audits.
- **Prioritize by data access.** Focus reviews on tools connected to email, documents, ticketing, and source code, and those with broad access to production environments and critical data.
- **Consolidate where possible.** Reduce long-tail sprawl by standardizing on a small number of approved platforms for common use cases.
- **Set clear rules for agents.** Require human approval for external actions and log agent activity centrally.
- **Reduce secrets exposure.** Provide internal guidance for developers and support teams, and deploy tooling that detects and blocks credentials and other sensitive data in prompts and other AI interactions.
- **Establish an AI acceptable use policy.** Share the policy with all employees who are using AI and collect acknowledgements. Review and update your policy regularly and renew policy acknowledgements.
- **Asses configuration and detected drift.** Integrations and other configuration settings should be revisited regularly to ensure they are still aligned with changes in usage and the organization's policies.
- **Understand AI data training policies.** Ensure that reviews of contracts and MSAs include careful review of AI data training policies, not just for AI tools but for any SaaS apps with AI-enabled features.
- **Educate with examples.** Training is most effective when grounded in real workflows, such as pasting from documents, uploading spreadsheets, or sharing screenshots.
- **Measure outcomes.** Monitor adoption, integration growth, and data-sharing trends to validate whether governance controls are working.

About Nudge Security

Nudge Security provides the leading SaaS and AI security governance platform purpose-built for the modern enterprise. As software adoption has shifted from centralized IT procurement to employee-driven usage and OAuth-based integrations, traditional security tools have struggled to provide visibility and control. Nudge was designed specifically to address this new reality.

At its core, Nudge Security helps organizations discover, secure, and govern SaaS and AI tools by focusing on how software is actually used—not just how it is officially approved. The platform provides continuous visibility into SaaS applications, browser extensions, integrations, and AI tools, along with the data access and behaviors that introduce real risk.

Nudge's approach to AI security and governance emphasizes:

- **Behavioral risk over theoretical risk.** Understanding how employees share data with AI tools, rather than relying solely on policy.
- **Integration-aware security.** Treating AI tools and agents as connected systems with real permissions, scopes, and blast radius.
- **Practical guardrails.** Enabling security teams to guide safer usage through approvals, least-privilege access, and just-in-time interventions.
- **Scalable governance.** Helping organizations manage AI adoption at the speed it is occurring, without blocking innovation.

By combining deep SaaS and AI visibility with behavioral insights, Nudge Security enables security leaders to move beyond reactive controls and build AI governance programs that are grounded in reality, enforceable in practice, and aligned with how work actually gets done.

As AI continues to evolve from tool to teammate, security teams need governance models that evolve just as quickly. Nudge Security exists to make that shift possible—by turning real-world usage data into actionable guardrails, not after-the-fact alerts.

Appendix: Definitions

- **Adoption rate:** Percentage of organizations with at least one observed instance of a vendor or tool during the measurement period.
- **Integration detected:** Evidence of an AI tool connected to another system (for example, via OAuth, API keys, plugins, or workflow automation).
- **Prompt activity:** Observed chat prompts submitted by employees to AI chat tools.
- **Copy/paste and file upload activity:** Observed user actions indicating that text or files were inserted into a prompt.
- **Sensitive-data detection:** Pattern-based detection of common secret formats, financial identifiers, and health-related information. Percentages reflect category share, not severity.