TERAMIND

TERAMIND RESEARCH

The Shadow AI Behavior Report

How Employees Move Data Into and Out of Unapproved AI Tools

Methodology Note

This report synthesizes findings from Teramind research, anonymized and/or aggregated behavioral observations, survey data, third-party research, and publicly available industry sources. Teramind does not identify any specific customer, employee, organization, or monitored individual in this report. Findings are directional and intended to illustrate observed and reported Shadow AI trends; they should not be interpreted as statistically representative of all organizations, industries, regions, or deployment environments unless expressly stated. Where Teramind-derived observations are referenced, such observations are based on aggregated, anonymized, or otherwise de-identified research outputs and are not intended to disclose or describe any specific customer environment, employee activity, customer content, or confidential customer data.

Published: 2026 | Teramind Research & Intelligence Team
Based on behavioral telemetry, survey data, and cross-industry benchmarking from 2025–2026

Executive Summary

89%
Workplace AI
Use occurs outside enterprise-governed channels

$670K
Breach cost
Average additional breach cost tied to Shadow AI activity

86%
Organizations
Lack visibility into how data flows to/from AI tools

20%
Organizations
Have already suffered a Shadow AI–related breach

32%
Employees
Only 1 in 3 employees with governance policies undergo regular AI audits

48%
Employees
Say they would keep using AI tools even if explicitly banned

Section 1 · Prevalence and Sprawl - How Widespread Is Shadow AI?

Shadow AI appears to be widespread across many enterprise environments. Multiple research efforts suggest a meaningful gap between official AI adoption and actual employee AI usage. The gap between official AI adoption and actual AI usage is significant, and organizations are largely unaware of how deeply unapproved tools have embedded themselves in daily workflows.

40%
Companies
Have an official LLM subscription

90%
Companies
Have employees regularly using AI tools

98%
Organizations
Have employees using unsanctioned apps including Shadow AI

1.2 Shadow AI Tool Proliferation

The volume and diversity of unsanctioned AI tools in use creates a significant discovery and inventory challenge. Research consistently shows that organizations are vastly underestimating both the number of tools in use and the velocity at which new tools appear and disappear across their employee base.

Teramind survey

269
Distinct tools

Shadow AI tools per 1,000 employees (small business)

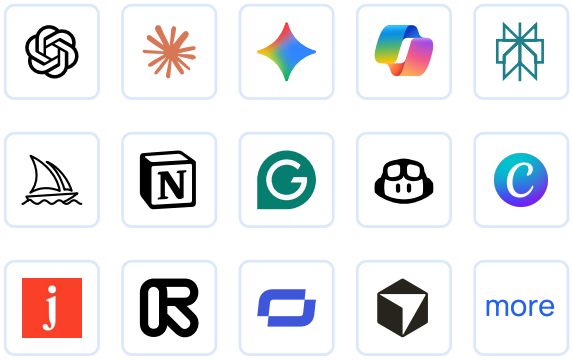

66

GenAI apps 10% high-risk

Detected per organization

Teramind survey

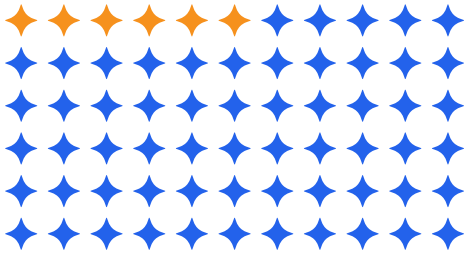

68 /100

Avg. security score

Shadow AI tools

Teramind survey

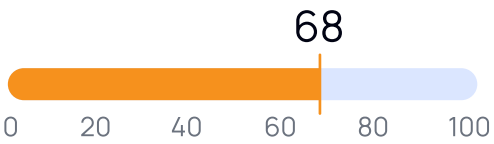

~76%

Tools

Failing SOC 2 compliance

Teramind survey

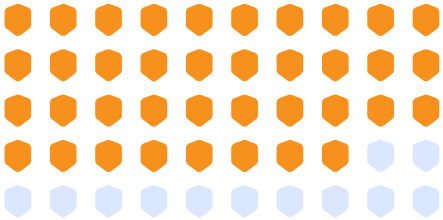

98%

Organizations with employees

Using unsanctioned apps

Teramind survey

69%

Organizations

Suspect employees use prohibited GenAI

Gartner 2025

1.3 Persistence: Shadow AI Is Not a Temporary Experiment

One of the most operationally significant findings in recent Shadow AI research is tool persistence. Unlike other unauthorized software that employees might use briefly before abandoning, Shadow AI tools become embedded in workflows. Once adopted, they are rarely removed — in many organizations, they may remain undetected for extended periods of time.

400+

Days after first appearing in an organization's environment

Unsanctioned AI tools persist

- Remaining active
- Unmonitored
- Data-processing

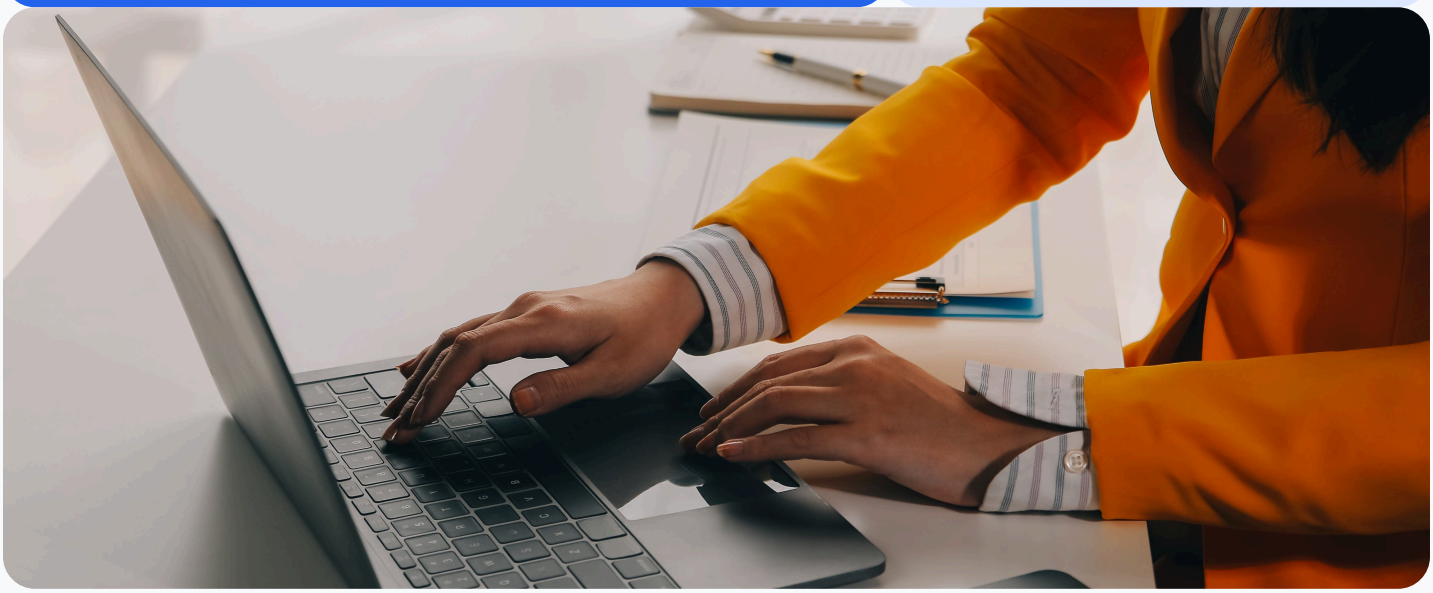

This persistence means that a single governance failure - an employee signing up for a new AI writing tool on a Monday - can translate into more than a year of ungoverned data movement before any remediation occurs, if it occurs at all.

Section 2 · Data Exposure - What Information Is Entering AI Tools?

The prevalence of Shadow AI would be a manageable governance challenge if employees were using these tools only for low-risk tasks - drafting personal emails, generating creative content, or exploring general-purpose queries. The data shows something categorically different. Employees are routinely inputting sensitive, regulated, and competitively critical organizational data into unapproved AI systems.

2.1 Volume and Growth of Sensitive Data Flowing to AI

27.4%
Corporate data
Sent to AI tools qualifies as sensitive

2.5×
Increase in GenAI DLP incidents
Year-over-year

15%
DLP incidents
Now involve generative AI

485%
Increase in corporate data sent
To AI tools from 2023 to 2024

40%
GenAI file uploads
Include PII or PCI data

223 per month
Sensitive-data-to-AI incidents
At the average organization

2.2 Categories of Sensitive Data Exposed

Categorizing the types of sensitive data employees are sending to AI tools reveals a risk profile that intersects with nearly every compliance framework - GDPR, HIPAA, PCI-DSS, SOX, and IP protection law.

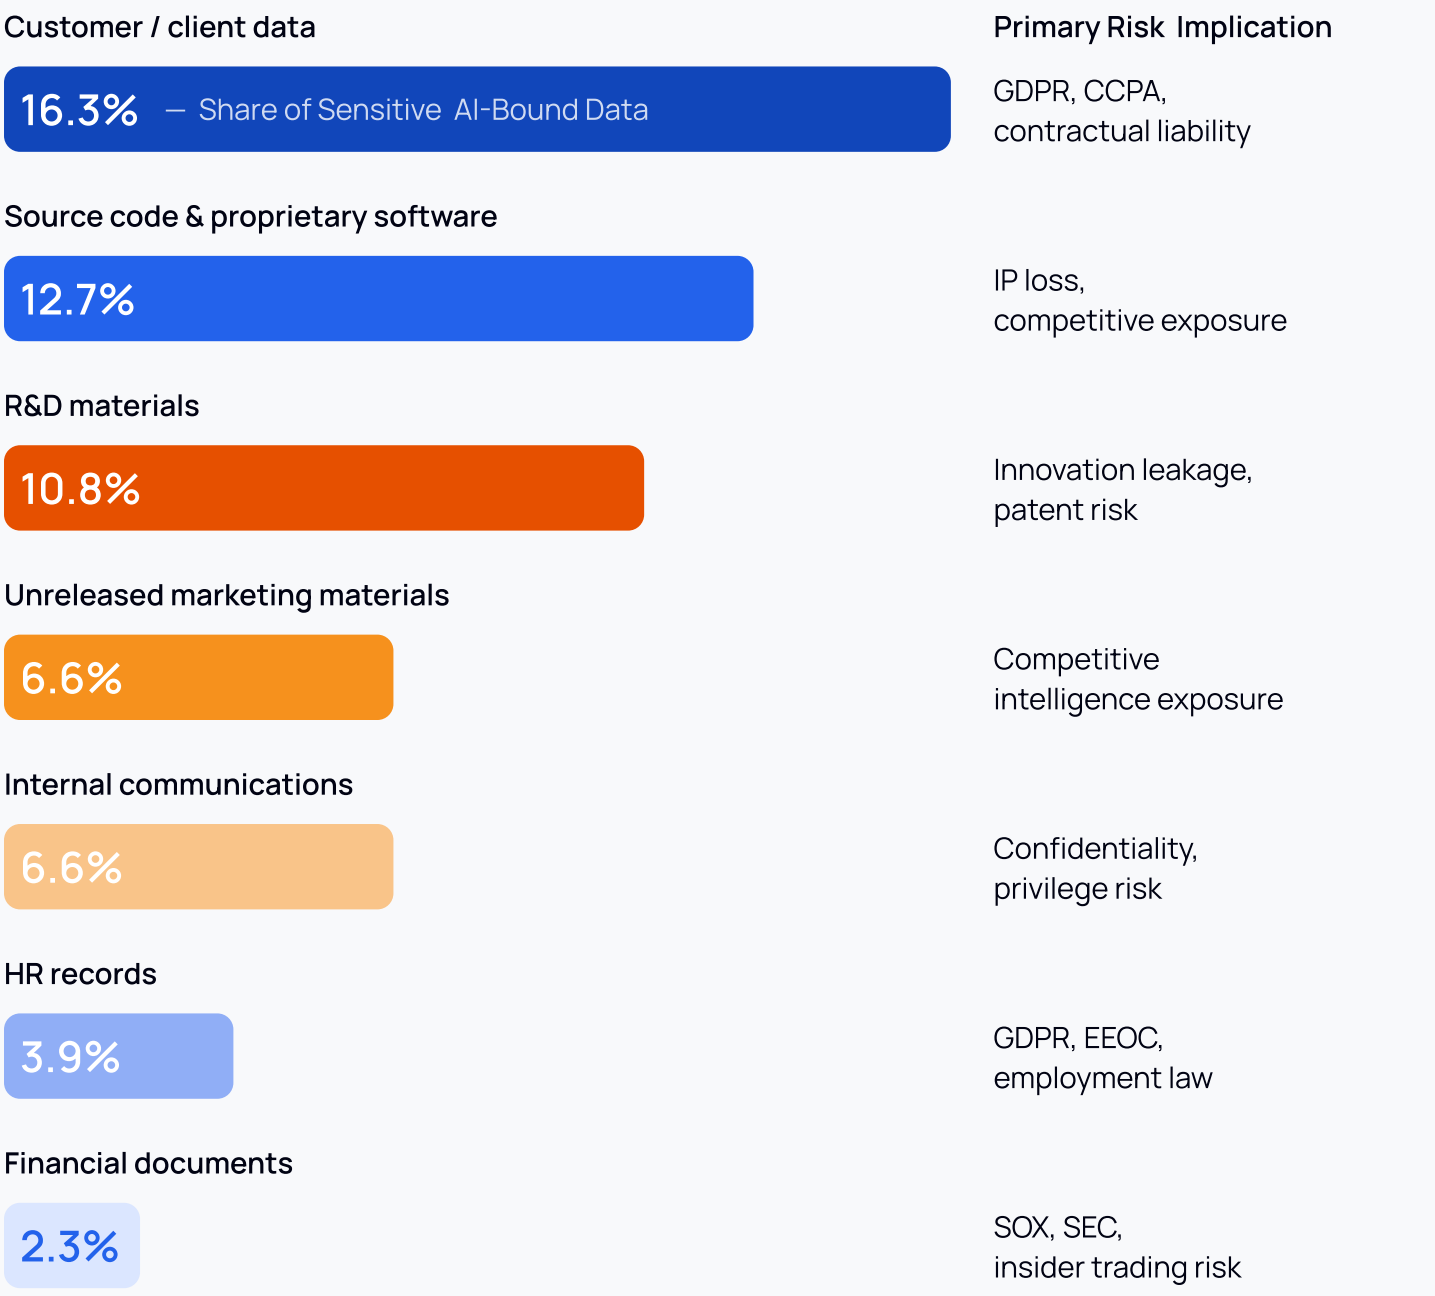

*Risk implications listed in this report are illustrative and not intended as legal conclusions. Applicability depends on the facts, jurisdiction, data type, industry, contractual obligations, and how the relevant technology is configured and used

Among employees using unsanctioned AI tools, 33% have shared research or datasets, 27% have shared employee data (including salary and performance information), and 23% have shared company financial information.

2.3 The Copy/Paste Problem

Data does not always travel to AI tools through file uploads or API calls. The most common pathway is also the hardest to detect through traditional DLP controls: manual copy-and-paste. This behavioral pattern represents a structural gap in most existing security architectures.

Teramind research

40%

IT leaders

Use occurs outside enterprise-governed channels

Teramind research

14 pastes per day

Employees make

Using non-corporate accounts

Teramind research

3 / 14

Daily pastes

Contain sensitive data

Teramind research

32%

GenAI Accounts for Data Exfiltration

Of all corporate-to-personal data exfiltration

Section 3 · Governance Bypass - How Shadow AI Evades Enterprise Controls

The governance challenge posed by Shadow AI is not simply that employees are using unapproved tools. It is that the mechanisms those employees use to access AI — personal accounts, unmanaged browsers, non-SSO authentication — systematically bypass the identity, access, and data controls that enterprise security architectures depend on.

The Personal Account Problem

The single most consequential governance failure in the Shadow AI landscape is the use of personal accounts to access AI tools on corporate devices. Personal accounts are invisible to enterprise DLP, unmanaged by IT, excluded from SSO and MFA enforcement, and not subject to corporate data retention or deletion policies.

Employees using GenAI on corporate devices, 72% authenticate with personal email accounts, 17% use corporate emails without SSO, and only 11% use properly governed corporate channels.

Share of GenAI Users (Corporate Devices)

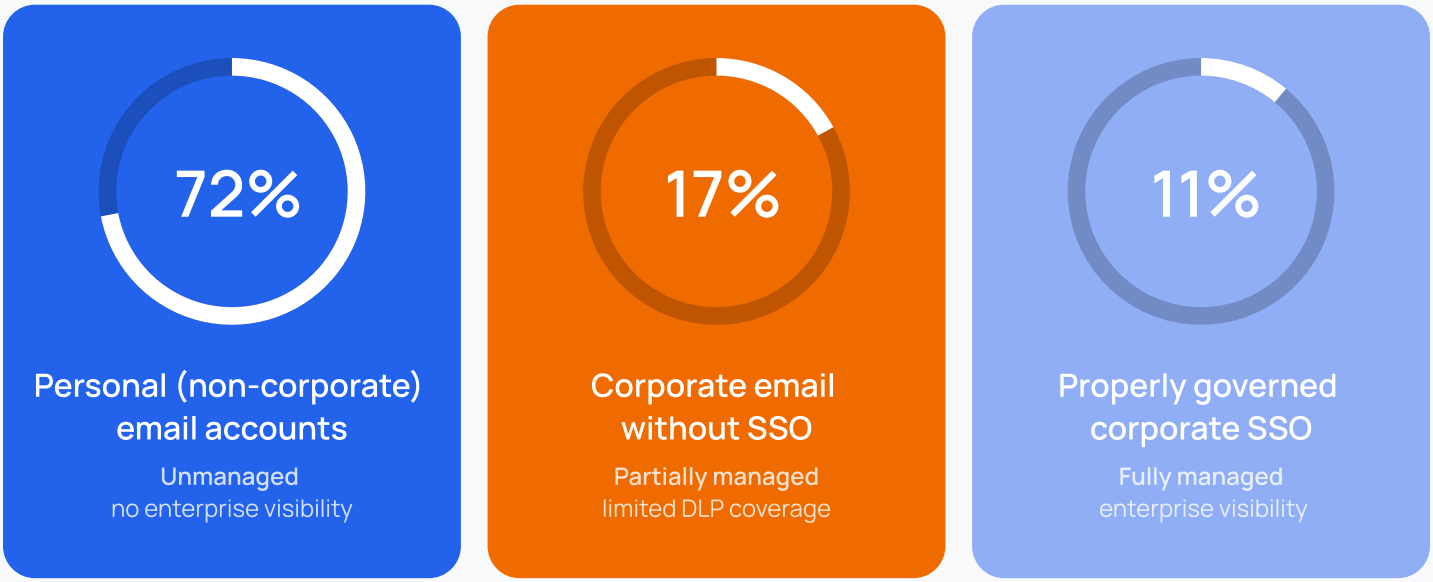

67% of all AI usage in enterprise environments occurs through unmanaged personal accounts. The ChatGPT-specific figure from (73.8% non-corporate) and the Gemini figure (92.4% non-corporate) suggest that even the most widely adopted consumer AI platforms remain largely outside enterprise governance.

3.2 Organizational Visibility Gap

Even when organizations suspect Shadow AI activity is occurring, most lack the instrumentation to confirm, quantify, or respond to it.

86%
Organizations
Lack visibility into how data flows to and from AI tools

60%
IT teams
Cannot see specific prompts employees make using GenAI tools

60%
IT teams
Lack confidence in their ability to identify unapproved AI use

3.3 The Evasion Problem: Blocking Redirects, It Does Not Stop

Perhaps the most operationally significant finding for security teams is what happens when organizations attempt to block access to specific AI tools. Rather than stopping Shadow AI behavior, blocking tends to redirect it — pushing employees toward alternative tools, workarounds, or personal devices that are even harder to monitor.

Organizations that respond to Shadow AI solely through prohibition are not solving the problem - they are displacing it into less visible channels. Effective governance requires visibility, behavioral controls, and approved alternatives working together.

Section 4 · Workflow Dependence - AI Output Is Re-Entering the Business

The Shadow AI risk conversation has historically focused on what employees put into AI tools. An equally significant - and less frequently measured - risk is what comes out. AI-generated content is flowing back into business systems, external communications, regulated workflows, and customer-facing outputs, often without disclosure, verification, or attribution.

4.1 AI Output Insertion Map

Research from Cyberhaven provides a cross-industry view of where AI-generated content is appearing in business operations. These insertion rates represent known floor estimates — actual rates are likely higher given the governance visibility gap documented in Section 3.

Business System / Workflow	AI-Originated Content Share	Risk Implication
Professional network posts (e.g., LinkedIn)	5.1%	Brand, disclosure, authenticity
R&D materials and documentation	3.4%	IP integrity, patent attribution
Source code insertions	3.2%	Security vulnerabilities, licensing
Graphics and design content	3.0%	Copyright, IP, disclosure
Customer communications	Estimated material	Accuracy, trust, compliance
Legal documents and contracts	Estimated material	Privilege, enforceability

4.2 The Integrity and Compliance Dimension

Once AI-generated content enters regulated workflows - legal documents, financial filings, clinical notes, audit reports - the organization assumes legal and compliance liability for that content regardless of its origin. In sectors governed by strict accuracy, disclosure, or provenance requirements, the unchecked insertion of AI-generated content creates audit risk, regulatory exposure, and potential liability.

The workflow dependence problem also has a talent and institutional knowledge dimension. When employees outsource judgment, analysis, and drafting to AI tools without organizational awareness, organizations lose visibility into which outputs are human-validated and which are not - degrading the reliability of internal knowledge assets over time.

Section 5 · Human Factors - Why Employees Use Shadow AI

Understanding why employees adopt Shadow AI tools is essential to designing governance strategies that work. Framing Shadow AI solely as a compliance problem — or as employee misconduct — misses the central dynamic: employees are rational actors responding to productivity pressure, policy ambiguity, and the absence of adequate approved alternatives.

5.1 Productivity Is the Primary Driver

80%
AI-using employees

Say AI improves their productivity

60%
Employees

Agree that using unsanctioned AI is worth security risks if it helps meet deadlines

63%
Employees

Believe it is acceptable to use AI tools without IT oversight if no approved option exists

5.2 Leadership Sets the Tone - and It Is Not Governance-First

One of the most important and counterintuitive findings in recent research is that Shadow AI risk tolerance is highest at the top of organizational hierarchies, not the bottom. The common assumption — that Shadow AI is primarily a frontline worker problem driven by technical curiosity or carelessness — is not supported by the data.

Seniority Level	"Speed over security" on AI tools	Shadow AI Implication
C-Suite / President	69%	Executive behavior normalizes ungoverned AI
Directors / Senior VPs	66%	Middle management reinforces high-risk norms
Managers	~50% (est.)	Supervision gap: managers may enable or model Shadow AI
Administrative staff	37%	Frontline workers are actually more risk-conscious

5.3 The Policy Awareness Gap

Employees frequently do not know what is or is not permitted. This is not primarily an enforcement problem — it is a communication and training failure that precedes enforcement.

Policy/Training Finding	Stat
Employees who received conflicting guidance on AI use	51%
Employees who lack clear guidance on AI usage policies	56%
Employees who received extensive AI training	7.5%
Employees who received no AI training at all	23%
Orgs that experienced a breach and had no AI governance policy	37%
Orgs that experienced a breach and had no AI governance policy	63%

Source: Teramind research

5.4 Concealment: Employees Know the Risk, and Hide It

A significant share of employees are not oblivious to the governance implications of Shadow AI — they are actively aware that their behavior may be problematic and are concealing it.

49%

of employees claim not to use AI to avoid judgment from colleagues

45%

pretend to know how to use AI while actually relying on it secretly

62%

of Gen Z employees actively hide their AI use at work

This concealment dynamic has a compounding effect on governance. Organizations that rely on self-reporting or manager observation to identify Shadow AI activity will systematically undercount the problem - particularly among younger, more AI-fluent employees who are also the highest-velocity AI adopters.

Additionally, 21% of employees believe their employer will simply look the other way on Shadow AI use as long as work is completed on time - a perception that reflects the risk-tolerance signals being sent from leadership described in Section 5.2.

Section 6 · The Financial Cost of Shadow AI

Shadow AI has moved from a theoretical risk to a quantifiable financial exposure. The IBM Cost of a Data Breach Report 2025 - conducted across 600 organizations by the Ponemon Institute - provides the most rigorous financial analysis to date of Shadow AI's impact on breach costs.

20%

of studied organizations experienced a breach caused directly by Shadow AI

IBM 2025

$670K

average additional breach cost at high Shadow AI organizations vs. low

IBM 2025

$4.2M

average total cost of a Shadow AI–related data breach

Organizations that experienced Shadow AI–related breaches were significantly more likely to expose customer PII (65% of cases) and intellectual property (40% of cases) compared to standard breach patterns - a distribution that intersects with the highest-consequence regulatory and legal liability frameworks.

Section 7 · The Governance Maturity Gap

The research consistently reveals that organizational AI governance is lagging far behind the pace of employee AI adoption. This is not merely an awareness problem — it is a systemic failure of policy, technology, and organizational accountability structures.

Governance Indicator	Current State
Organizations with an AI governance policy	34%
Orgs without a policy and not developing one	61%
Governed orgs that conduct regular AI audits	31%
Orgs that experienced AI security incident lacking access controls	97%
Orgs where GenAI access is completely unrestricted	22%

Source: Teramind research

7.1 The Training-Behavior Gap

Training programs exist, but they are not closing the behavior gap. Recall of training does not translate reliably into lower-risk behavior, particularly when the structural conditions that drive Shadow AI - productivity pressure, unclear approved alternatives, low perceived risk of detection - remain in place.

40% of employees recalled receiving AI training.
Yet 40% of those same employees reported daily use of unapproved AI tools. Training without behavioral controls and approved alternatives is insufficient as a standalone governance strategy.

However, training at scale does show measurable impact: organizations with formal AI training programs see 40% fewer security incidents than those without. The training gap is therefore a meaningful risk lever but only when accompanied by clear policy, approved tools, and behavioral monitoring.

7.2 A Framework for Effective Shadow AI Governance

Based on the research synthesized in this report, organizations may consider the following components as part of a Shadow AI governance program:

01 **Behavioral visibility**
Real-time monitoring of data movement to and from AI tools — copy/paste, file upload, and AI output insertion.

02 **Account governance**
Detection and policy enforcement on personal-account AI usage across corporate devices and networks

03 **Tool inventory and risk scoring**
Continuous discovery and classification of AI applications in the environment, with risk scoring by data-retention policy, SOC 2 compliance, and account type

04 **Approved alternatives**
Provision of enterprise-grade AI tools for common use cases — without approved alternatives, prohibition will fail

05 **Clear, specific policy**
Policy documents that specify which tools are approved, which data types are prohibited, and what employees should do when their needed use case is not covered

06 **Training with behavioral reinforcement**
Training that is followed by observable behavioral standards, not just completion tracking

07 **Leadership accountability**
Governance programs that apply equally to executive users - and that treat leadership-level Shadow AI with the same scrutiny as frontline behavior

Compliance Note

Monitoring employee activity, AI usage, copy/paste behavior, file uploads, application usage, prompts, outputs, or data movement may be subject to jurisdiction-specific employment, privacy, labor, works council, collective bargaining, notice, consent, proportionality, and data protection requirements. Organizations should consult counsel and implement appropriate policies, notices, access controls, retention limits, and governance safeguards before deploying monitoring or enforcement controls.

Disclaimer

This report is provided for general informational and educational purposes only. The information, analysis, observations, trends, benchmarks, and recommendations contained in this report are not intended to constitute, and should not be relied upon as, legal, regulatory, compliance, security, employment, privacy, financial, or professional advice. Teramind makes no representation or warranty, express or implied, as to the accuracy, completeness, timeliness, suitability, or reliability of the information contained in this report. Any references to industry practices, risk trends, security concerns, workforce monitoring practices, insider threat indicators, data loss prevention considerations, or compliance-related topics are provided for general discussion only and may not apply to every organization, jurisdiction, industry, workforce, or use case.

Organizations are solely responsible for evaluating their own legal, regulatory, employment, privacy, security, and operational obligations, including any requirements relating to employee monitoring, workplace surveillance, data protection, consent, notice, collective bargaining, works councils, cross-border data transfers, retention, access controls, and use of monitoring or security technologies. Readers should consult with their own legal, privacy, security, HR, and compliance professionals before implementing or relying on any practice, control, technology, policy, recommendation, or conclusion discussed in this report.

Any examples, scenarios, statistics, or descriptions included in this report are illustrative only and should not be interpreted as guarantees of results, predictions of future events, legal conclusions, or representations regarding any specific customer, employee, organization, threat actor, incident, or deployment environment. To the extent this report references third-party sources, tools, standards, frameworks, or publications, such references are provided for informational purposes only and do not constitute endorsement, sponsorship, or approval by Teramind or any third party.

Use of Teramind products and services remains subject to the applicable agreement between Teramind and the customer, including Teramind's Terms of Service, Data Processing Addendum, order form, and any other written agreement executed by the parties. Nothing in this report modifies, supplements, expands, or supersedes any contractual rights, obligations, warranties, disclaimers, limitations of liability, data processing terms, or security commitments applicable to Teramind's products or services.

Teramind does not warrant that the use of any product, feature, configuration, workflow, recommendation, or practice described in this report will ensure legal compliance, prevent all security incidents, eliminate insider threats, stop data loss, satisfy regulatory requirements, or produce any particular business, security, compliance, or operational outcome. Customers and other readers remain solely responsible for configuring and using any technology in a lawful, appropriate, and context-specific manner.