# Shadow AI in the Enterprise

## A Comprehensive Research Analysis

Aggregated from 10+ Primary Industry Sources
Including: IBM, Netskope, Harmonic Security, UpGuard, BlackFog, LayerX, WalkMe, ISACA

Created For:
**IT Security Professionals & Tribal IT Leadership**

Created by:
Synergy Information Solutions

Created:
February 3, 2026

## Table of Contents

## Versioning

| Version | Date | Author | Change Description |
|---|---|---|---|
| **1.0** | 02.03.26 | J. Eaton | Initial Document Creation |
| **1.1** | 02.06.26 | S. Stinnett | Formatting |

## Executive Summary

This report synthesizes findings from 10+ industry research sources published between 2024-2026, analyzing the phenomenon of "Shadow AI" in enterprise environments. The data has been cross-referenced and reconciled to provide IT security professionals with accurate, actionable intelligence.

**Shadow AI Definition:**
*The use of AI tools for work purposes outside of organizational visibility, governance, or approval. This specifically refers to employees using AI to process company data, not personal use during work hours. All major research sources define Shadow AI as work-related AI usage that bypasses corporate controls.*

## Shadow AI Taxonomy

This report uses "shadow AI" to encompass four distinct but related phenomena. When reviewing statistics, note which category each finding addresses:

| Category | Description |
|---|---|
| **Unapproved Tools** | AI applications not sanctioned by IT/security (e.g., personal ChatGPT accounts used for work) |
| **Personal Accounts** | Approved tools accessed via personal (non-corporate) accounts, bypassing enterprise controls |
| **Ungoverned Usage** | Approved tools used without enterprise controls (e.g., SSO, DLP, audit logging) |
| **Embedded AI** | AI features within other SaaS applications that IT may not have visibility into |

## Key Findings at a Glance

| Metric | Finding (with denominator context) |
|---|---|
| **Shadow AI Prevalence** | 49% of all employees (BlackFog 2026); 78-81% among screened AI users (WalkMe, UpGuard 2025) |
| **Personal Account Usage** | 47% of genAI users (down from 78% YoY) (Netskope 2026) |
| **Sensitive Data in Prompts** | 8.5% of prompts (Harmonic 2024); 40% of file uploads contain PII/PCI (LayerX 2025) |
| **Free-Tier Usage** | 58% of employees using unapproved AI rely on free versions (BlackFog 2026) |
| **Breach Cost Association** | $670,000 higher average cost associated with high vs low shadow AI (IBM 2025) |
| **C-Level Risk Tolerance** | 69% of President/C-level believe speed trumps security (BlackFog 2026) |
| **AI Governance Gaps** | 63% of breached orgs lack AI governance policy (IBM 2025); only 15% have formal policies (ISACA 2024) |

# How to Interpret This Data

Before diving into the findings, it is essential to understand why different studies report different numbers and how to reconcile them for your organization.

## Why the Numbers Vary

### Sample Definition Matters

The single biggest source of variation is whether a study surveyed all employees or only those who already use AI:

- **All employees (BlackFog, Cybernews):** Report ~49-59% shadow AI usage
- **AI users only (WalkMe):** Reports ~78% usage of unapproved tools among surveyed AI users. Note: WalkMe screened respondents for AI use at work, so 100% AI adoption reflects study design, not the broader workforce.
- **Security leaders (UpGuard):** Reports 81% of employees and 90% of security leaders using unauthorized AI tools

*Both numbers are accurate for their populations. For tribal organizations assessing risk, the "all employees" figure (~50%) is more useful for estimating total exposure, while the "AI users" figure (~80%) tells you about the behavior of those already using AI.*

### Measurement Method Matters

- **Survey-based studies (UpGuard, WalkMe, BlackFog):** Ask employees what they do—subject to recall bias and social desirability
- **Telemetry-based studies (Harmonic, LayerX, Netskope):** Observe actual network traffic and prompts—more accurate but limited to enterprises with monitoring deployed

*Telemetry data is generally more reliable for behavioral statistics, while surveys are better for understanding attitudes and awareness.*

**Important Data Limitations:**
*Telemetry studies (Harmonic, LayerX, Netskope) measure environments where their security tooling is deployed, which may skew toward larger or more security-mature organizations. Vendor classifiers define "sensitive" using proprietary methods. Survey results depend heavily on question wording and respondent interpretation. Breach cost studies show correlation, not necessarily causation.*

# Shadow AI Prevalence & Trends

## Current State (2025-2026)

Cross-referencing multiple sources, we can establish the following baseline:

- Approximately 49-50% of all employees use AI tools without corporate approval for work tasks (BlackFog 2026, Cybernews surveys)
- Among employees who use any AI tools, approximately 78-81% use at least some unapproved tools (WalkMe 2025, UpGuard 2025)
- 47% of genAI users use personal AI apps (down from 78% year-over-year) (Netskope 2026)
- Shadow AI is widespread; multiple studies suggest it is rarely absent in enterprise environments

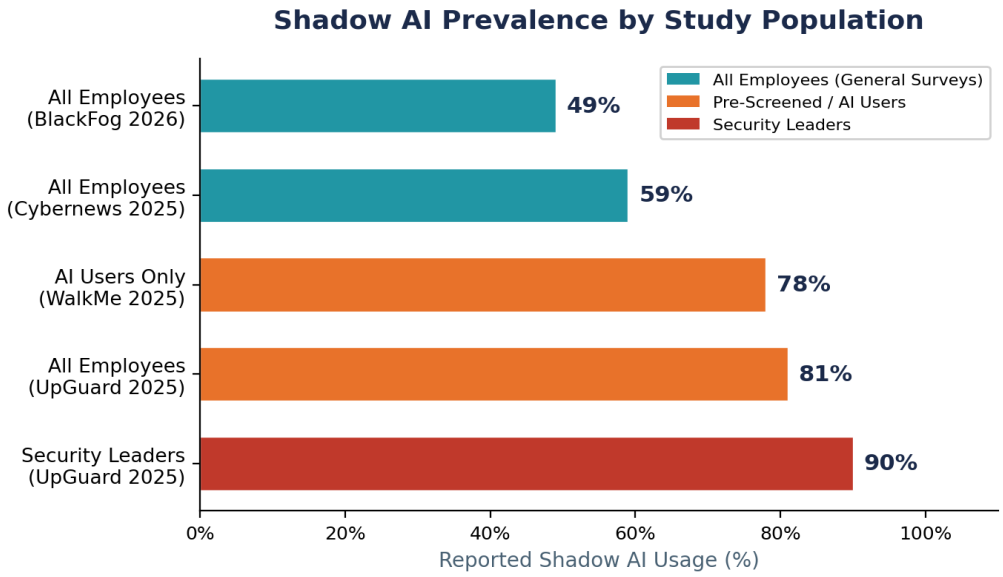


## Year-Over-Year Progress

The good news: organizations are making progress. Netskope's longitudinal data shows significant improvement in governance adoption. Personal account usage dropped 31 percentage points year-over-year as organizations adopted managed AI solutions.

**Interpretation for Tribal IT:** This trend suggests that organizations providing sanctioned AI alternatives successfully shift usage away from shadow channels. However, the 47% still using personal accounts represents substantial ongoing risk—nearly half of AI activity remains outside organizational visibility.

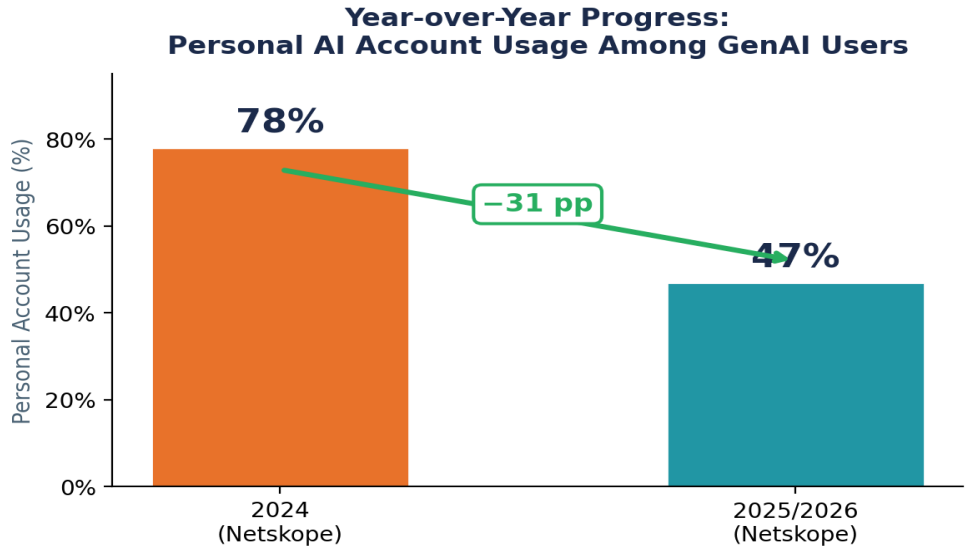

## The Leadership Paradox

One of the most significant findings across multiple studies is that senior leadership—the people responsible for setting security policies—are among the most likely to circumvent them.

### Supporting Data

- 69% of President/C-level executives believe speed trumps security (BlackFog 2026)
- 90% of security leaders report using unapproved AI tools at work (UpGuard 2025)
- 69% of CISOs incorporate unauthorized AI into daily workflows (UpGuard 2025)
- Senior leadership is 50% more likely to use shadow AI than junior staff (UpGuard 2025)
- In contrast, only 37% in administrative roles believe speed trumps security (BlackFog 2026)

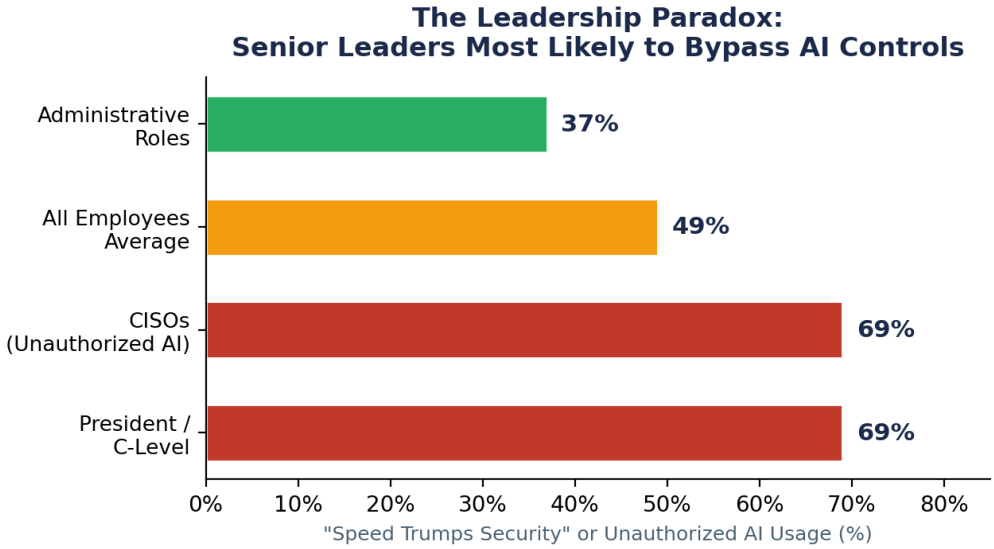


## Implications for Governance

This pattern indicates that shadow AI is fundamentally a cultural problem, not merely a training or awareness issue. Policies that target only frontline employees while exempting leadership (formally or informally) will fail. Effective governance requires visible executive commitment and accountability.

**Tribal Governance Consideration:**
*In tribal organizations where leadership may include elected officials or commissioners with varying technical backgrounds, explicit AI governance training for leadership—including tribal council members and commission chairs—is essential. The "tone at the top" directly influences organizational compliance culture.*

# Sensitive Data Exposure

## How Much Data is at Risk?

Different studies measure sensitive data exposure in different ways, leading to apparently conflicting statistics:

| Measurement | Rate | Source & Context |
|---|---|---|
| **Prompts containing sensitive data** | 8.5% | Harmonic analysis of 2024 prompt data (published Apr 16, 2025)—includes potentially concerning content |
| **File uploads containing PII/PCI** | 40% | LayerX 2025 browser telemetry—enterprise customers |
| **Uploaded files with sensitive content** | 26.4% | Harmonic Q3 2025, US/UK enterprise sample—up from 22% in Q2 |

*The variance reflects different definitions of "sensitive." File uploads (documents, spreadsheets) are more likely to contain sensitive data than text prompts.*

## What Types of Data Are Exposed?

Based on Harmonic's analysis of 22.4 million enterprise prompts (January-December 2025):

- **Code: 26.5%** — Significant for organizations with proprietary systems
- **Legal discourse: 22.3%** — Client contracts, case materials
- **M&A data: 12.6%** — Strategic business intelligence
- **PII: 15%** — HR records, employee data

*Note: Harmonic uses both top-level categories and subcategories. The figures above are subcategory-level; some sources report higher-level rollups (e.g., "Legal Content: 35%" which includes multiple subcategories).*

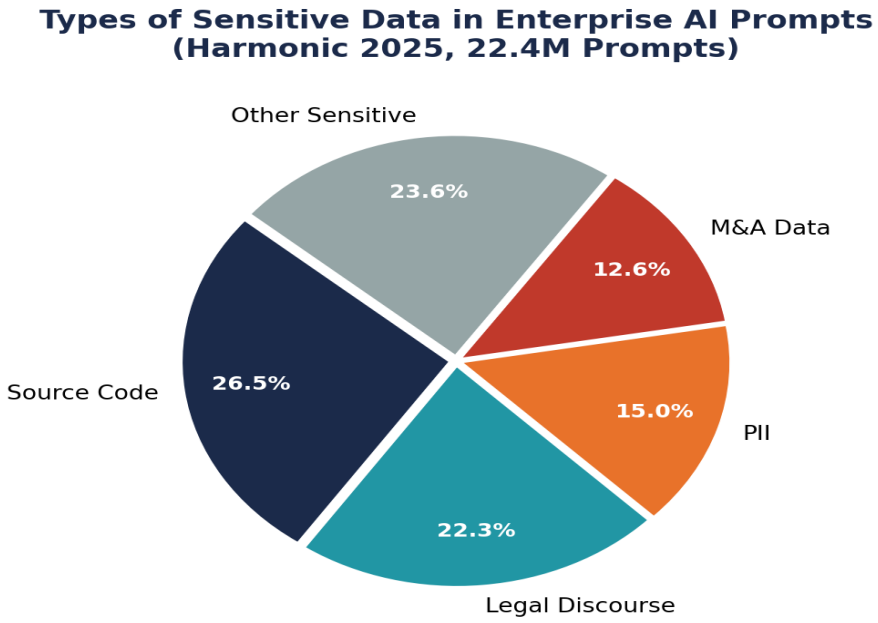

## The Free-Tier Risk Factor

A critical but often overlooked factor is the distinction between free and enterprise AI tiers. Free-tier AI services present elevated risks:

- Training on inputs: Many free tiers allow use of inputs for service improvement/training; terms vary by provider and account type
- **No data retention controls:** Organizations cannot enforce deletion policies
- **No audit logs:** Complete invisibility into what data was shared
- **No enterprise security features:** Missing SSO, DLP, access controls

Enterprise tiers can provide contractual and administrative controls (e.g., no-training commitments where available, retention/deletion controls, audit logs, SSO) that materially reduce risk compared with free/personal accounts.

### Where Free-Tier Risk Concentrates

Personal/free-tier accounts often bypass enterprise audit, DLP, and retention controls, creating a visibility gap.

- 16.9% of detected sensitive exposures flowed through personal/free-tier accounts (Harmonic 2025)
- Among employees using unsanctioned AI tools, 58% rely on free AI tools (BlackFog 2026)
- 34% of all employees admit to using free versions of company-approved AI tools (BlackFog 2026)

The implication is clear: while blocking alone can push usage to personal devices, prioritizing visibility into free/personal usage and migrating work activity to managed enterprise accounts (SSO, logging, retention controls) materially reduces data exposure risk and improves auditability.

## Financial Impact

### Breach Cost Association

**Important Methodological Note:**
*The IBM findings show associations observed in their breach dataset, not necessarily causation. Organizations with high shadow AI may share other characteristics that contribute to higher breach costs. The data demonstrates correlation, not proof that shadow AI directly causes the cost differential.*

### IBM 2025 Key Findings

- 20% of breaches (1 in 5) involved shadow AI as a contributing factor
- Organizations reporting high levels of shadow AI had $670,000 higher average breach costs than those with low/no shadow AI ($4.74M vs $4.07M).
  Breaches that included unsanctioned/shadow AI averaged $4.63M.
- 65% of shadow AI-related breaches exposed customer PII (vs. 53% for standard breaches)
- IP was compromised in 40% of shadow-AI-related incidents (vs 33% overall)
- 97% of organizations reporting AI-related breaches lacked proper AI access controls
- Global breach lifecycle averaged 241 days in 2025; IBM attributes faster containment in part to AI-powered defenses.

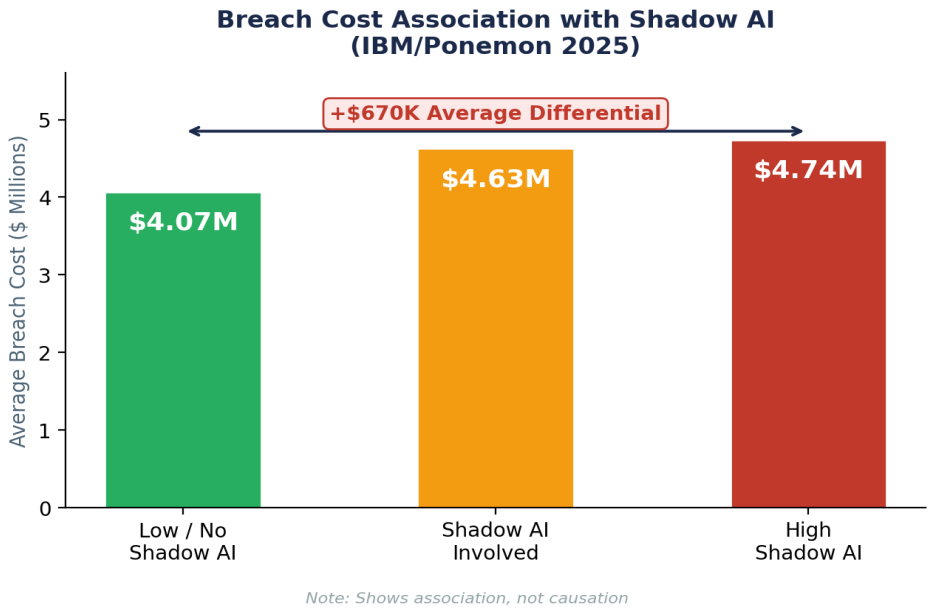


## Beyond Breach Costs

Shadow AI also contributes to operational inefficiency:

- $104 million average enterprise loss from underutilized technology investments (WalkMe 2025)—note: this refers to digital adoption generally, not shadow AI specifically
- Nearly 60% of surveyed employees said it often takes longer to figure out how to use an AI tool than to complete the task without it (WalkMe 2025)

# The Governance Gap

## Policy Gaps

- 63% of breached organizations lack AI governance policy or are still developing one (IBM 2025)
- Only 15% of organizations have formal AI policies; 40% offer no AI training (ISACA 2024, n=3,270 global IT professionals, March 2024)
- Only 34% perform regular audits for unsanctioned AI usage—among those with AI governance (IBM 2025)
- Only 37% have policies to manage AI or detect shadow AI (IBM 2025)

## Training Gaps

- Only 7.5% of employees received extensive AI training (WalkMe 2025)
- 23% received no AI training at all (WalkMe 2025)
- 51% receive conflicting guidance about AI usage from their employer (WalkMe 2025)

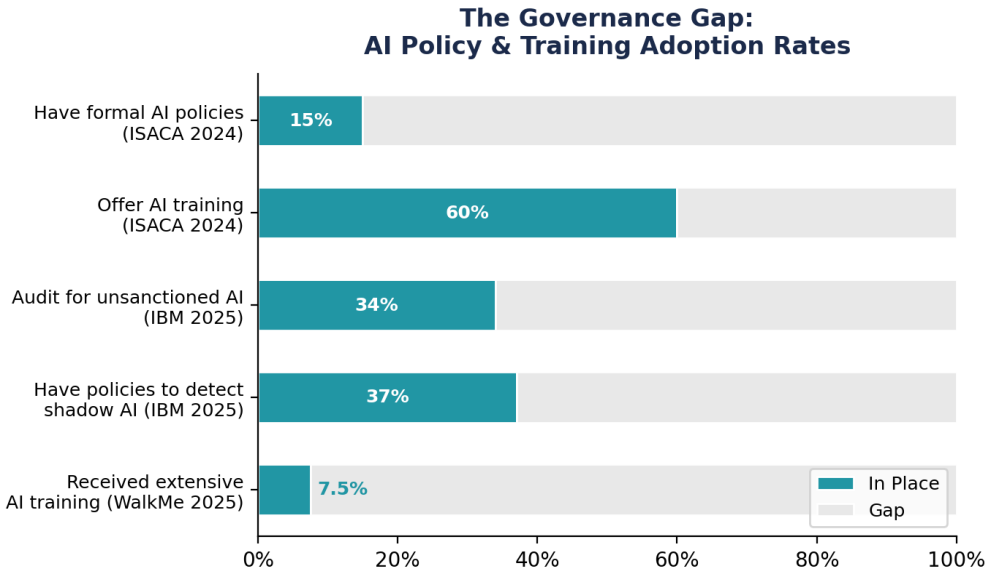


## The Training Paradox

Critically, UpGuard found that AI training does not appear to reduce shadow AI usage—employees who report understanding AI security requirements are more likely to use unapproved tools, apparently because knowledge increases confidence in managing risk independently.

*This suggests that training alone is insufficient. Technical controls (monitoring, blocking, managed alternatives) must accompany policy and education efforts.*

# Risk Concentration Analysis

Understanding where risk concentrates enables targeted governance. Harmonic's analysis of 22.4 million enterprise prompts (2025) reveals stark concentration:

## Key Concentration Points

- 6 applications account for 92.6% of all sensitive data exposure
- ChatGPT alone accounts for 71.2% of risk despite 43.9% of usage
- ChatGPT, Gemini, Claude, Microsoft Copilot, and Grok represent 85.5% of total sensitive data exposure
- 665 different AI tools detected in enterprise environments—blocking is not scalable
- Harmonic estimates blocking general chat tools would eliminate ~71% of enterprise AI value

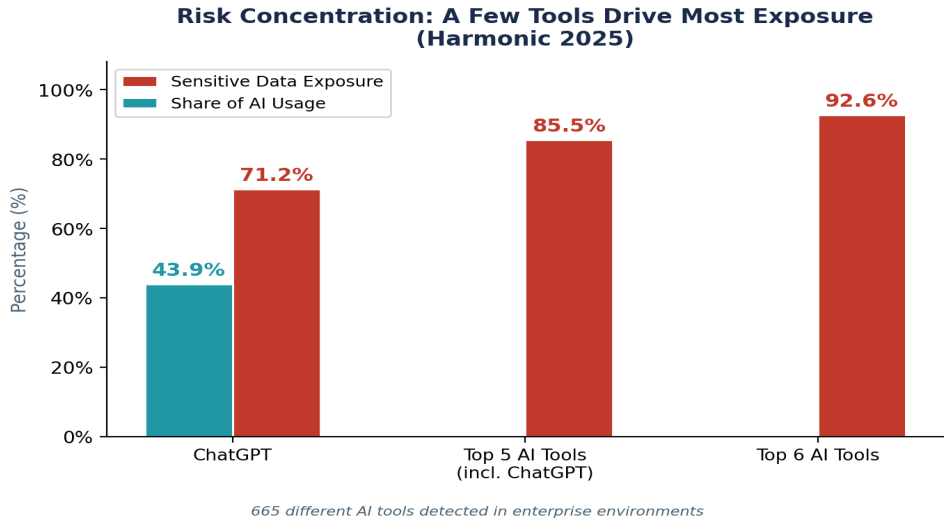


## Strategic Implication

Rather than broad blocking, focus governance on the highest-risk tools (e.g., ChatGPT, Gemini, Copilot, Claude, Perplexity) while implementing data-loss prevention at the prompt/upload level. This preserves productivity while addressing the concentration of risk. For regulated gaming systems and other high-sensitivity environments, a default-deny posture for external GenAI may be appropriate.

# Considerations for Tribal Organizations

Tribal gaming, healthcare, and government operations face unique considerations beyond general enterprise concerns.

## Data Sovereignty & CARE Principals

The CARE Principles for Indigenous Data Governance—Collective benefit, Authority to control, Responsibility, Ethics (Global Indigenous Data Alliance)—establish that tribal nations have inherent authority over their data. Shadow AI fundamentally undermines this sovereignty:

- Data uploaded to external AI services leaves tribal jurisdiction
- Training on tribal data could expose cultural, business, or governance information
- No ability to enforce deletion or audit what was shared
- Potential violation of data sharing agreements with vendors and partners

## Regulatory Considerations

### Gaming Commission Oversight

- NIGC has begun publishing AI-related resources (including an "NIGC and Artificial Intelligence" landing page and a September 2025 AI compliance plan). Note: the published AI compliance plan addresses NIGC internal federal governance requirements, not direct requirements for tribal gaming operations.
- AI-specific requirements for tribal gaming operations remain limited; however, gaming internal control standards and compact/internal control frameworks generally require system integrity, logical access

controls, and auditable processes. Unmanaged external AI use can introduce uncontrolled data flows and audit gaps.
- Gaming system data (patron information, transaction records, revenue data, and security configurations) should be treated as highly sensitive and protected under existing internal control and information-security frameworks.
- Recommendation: define a regulated system boundary for gaming systems and default-deny access to external GenAI services in those segments; adopt governed enablement (SSO, logging, DLP) for non-regulated corporate functions.

### Healthcare (HIPAA)

- Using AI tools that create, receive, or transmit ePHI without a Business Associate Agreement (BAA) creates significant HIPAA compliance risk
- Tribal health programs often serve small populations where re-identification risk is elevated

### Government Operations

- Enrollment data, land records, and governance documents are sovereignty assets
- Financial data for federal grant reporting has compliance requirements

## Small Organization Dynamics

Smaller organizations face disproportionate shadow AI challenges: limited IT resources for monitoring, fewer staff may mean higher per-employee AI adoption rates, and less formal governance structures can make policy enforcement difficult. Many tribal enterprises and government operations fall into this small-to-medium category.

# Recommendations

## Strategic Recommendations for Tribal Environments

1) Regulated System Boundary: Default-deny external GenAI for gaming systems. Segment regulated gaming systems and network zones and enforce a default-deny posture for external GenAI services. For non-regulated corporate functions, use governed enablement (SSO, logging, DLP) rather than blanket blocking.
2) Leadership-First Enablement: Resolve the leadership paradox with safe, sanctioned access. Roll out approved enterprise AI to Council/Executive leadership first with required training and visible compliance, positioned as protecting deliberations and sovereignty-sensitive workflows. Apply the same rules to leadership and staff.
3) Audit and migrate free/personal accounts: Buy enforceable controls, not just software. Use measured free-tier usage to justify enterprise procurement focused on contractual and administrative controls (no-training commitments where available, retention/deletion controls, audit logs, SSO/SCIM, DLP compatibility). Frame the benefit as maintaining Tribal control and auditability of data handling.
Minimum procurement requirements: no-training commitments where available, retention/deletion controls, audit logs, SSO/SCIM, incident notification terms, and clear data processing terms.

## Immediate Actions (0-30 Days)

1. **Establish baseline visibility**

- Deploy network monitoring to identify AI tool traffic
- Survey employees (anonymously) about current AI usage patterns
- Audit browser extensions and installed applications

- Identify and document regulated/high-sensitivity system boundaries (e.g., gaming systems, ePHI workflows, enrollment/governance repositories) to guide AI control zoning.

2. **Issue interim guidance**

- Communicate that AI usage is being assessed (not necessarily prohibited)
- Identify categories of data that must never be entered into any AI tool

## Short-Term Actions (30-90 Days)

3. **Develop formal AI governance policy**

- Define approved tools and use cases by data sensitivity tier
- Define a regulated system boundary (gaming systems and other high-sensitivity environments) and implement a default-deny posture for external GenAI services in those segments.
- Address CARE principles and tribal data sovereignty explicitly

4. **Implement technical controls**

- Prioritize eliminating free/personal accounts for work: restrict free-tier access where feasible and provide managed enterprise alternatives.
- Deploy DLP rules targeting AI tool domains
- Require enterprise SSO for any approved AI tools

5. **Leadership-first enablement**

- Provide Council/commission/executive leadership with approved enterprise AI access (SSO, logging) and require short, role-specific training.
- Apply the same rules to leadership and staff; establish a clear escalation path for repeated high-risk misuse and reinforce 'tone at the top' compliance.

## Long-Term Actions (90+ Days)

6. **Deploy enterprise AI solutions**

- Evaluate enterprise AI platforms with data residency controls
- Consider on-premises or tribal-cloud-hosted AI for sensitive workloads
- Negotiate BAAs and data processing agreements (including retention/deletion, audit logs, incident notification terms, and no-training commitments where available).

7. **Establish ongoing governance**

- Implement quarterly AI usage audits
- Create feedback mechanism for tool requests
- Include AI governance in annual security assessments

# Appendix A: Source Methodology

Understanding the methodology behind each source helps assess data reliability and applicability to your organization.

| Source | Method | Sample | Key Limitations |
|---|---|---|---|
| **IBM 2025** | Breach analysis (Ponemon) | 600 breached orgs globally | Only orgs that experienced breaches; findings are associations not causation |
| **Netskope 2026** | Network telemetry | Enterprise customers | Enterprise-focused; may not reflect SMB environments |
| **Harmonic 2025** | Prompt/upload analysis | 22.4M prompts and file uploads (Jan–Dec 2025), US/UK | Security-conscious enterprises with monitoring deployed |
| **UpGuard 2025** | Survey (Dynata/Prolific) | 1,562 respondents | Self-reported; subject to social desirability bias |
| **BlackFog 2026** | Survey (Sapio) | 2,000 US/UK employees | 500+ employee companies only |
| **Cybernews 2025** | Survey | 1,000+ US employees | Media-run survey; limited methodological transparency; self-reported behavior |
| **WalkMe 2025** | Survey (Propeller) | 1,000 US AI users | Pre-screened for AI use; 100% adoption reflects study design |
| **LayerX 2025** | Browser telemetry | Enterprise customers | Limited to LayerX customer base |
| **ISACA 2024** | Survey (online) | 3,270 IT professionals | Association membership may skew results toward security-aware orgs |

# Appendix B: Bibliography

Full citations for all sources referenced in this report:

IBM (2025). Cost of a Data Breach Report 2025: The AI Oversight Gap. Ponemon Institute, sponsored by IBM. Published July 2025. https://www.ibm.com/reports/data-breach (Press release: https://newsroom.ibm.com/2025-07-30-ibm-report-13-of-organizations-reported-breaches-of-ai-models-or-applications)

Harmonic Security (2026). What 22 Million Enterprise AI Prompts Reveal About Shadow AI in 2025. Published January 15, 2026. https://www.harmonic.security/resources/what-22-million-enterprise-ai-prompts-reveal-about-shadow-ai-in-2025

**Harmonic Security (2025).** GenAI in the Enterprise: It's Getting Personal. Q3 2025 Analysis. Published November 13, 2025. https://www.harmonic.security/blog-posts/genai-in-the-enterprise-its-getting-personal

Harmonic Security (2025). New Research: The Data Leaking into GenAI Tools (2024 data). Published April 16, 2025. https://www.harmonic.security/resources/new-research-the-data-leaking-into-genai-tools

**UpGuard (2025).** State of Shadow AI Report. Published November 10, 2025. https://www.upguard.com/press/new-research-from-upguard-reveals-68-of-security-leaders-admit-to-unauthorized-ai-usage

**BlackFog (2026).** Shadow AI Threat Grows Inside Enterprises. Published January 27, 2026. https://www.blackfog.com/blackfog-research-shadow-ai-threat-grows/

Netskope Threat Labs (2026). Cloud and Threat Report: 2026. https://www.netskope.com/resources/cloud-and-threat-reports/cloud-and-threat-report-2026

Cybernews (2025). Shadow AI soaring: 59% of employees hide AI use from their bosses. Published September 30, 2025. https://cybernews.com/ai-news/ai-shadow-use-workplace-survey/

**WalkMe (2025).** AI in the Workplace Survey 2025. Published August 27, 2025. https://www.walkme.com/news-releases/employees-left-behind-in-workplace-ai-boom-new-walkme-survey-finds/

WalkMe (2025). The State of Digital Adoption 2025: Special AI Edition. Published February 25, 2025. https://www.walkme.com/the-state-of-digital-adoption-2025/ (PDF: https://www.walkme.com/wp-content/uploads/2025/03/State-of-digital-adoption-2025-AI-edition_1767c57d.pdf)

**LayerX (2025).** AI is Now the #1 Data Exfiltration Vector in the Enterprise. https://layerxsecurity.com/blog/ai-is-now-the-1-data-exfiltration-vector-in-the-enterprise-and-nobodys-watching/

**ISACA (2024).** The AI Reality: New Research from ISACA Identifies Gaps in AI Knowledge, Training and Policies. AI Pulse Poll, n=3,270, March 2024. Published May 7, 2024. https://www.isaca.org/about-us/newsroom/press-releases/2024/the-ai-reality-new-research-from-isaca-identifies-gaps-in-ai-knowledge-training-and-policies

**NIGC (2025).** NIGC AI Compliance Plan for OMB Memorandum M-25-21. September 2025. https://www.nigc.gov/nigc-and-artificial-intelligence/

**Global Indigenous Data Alliance.** CARE Principles for Indigenous Data Governance. https://www.gida-global.org/care

*This report was prepared by Synergy Information Solutions to support tribal organizations in understanding and addressing shadow AI risks.*

*For assessment services or questions about this research, contact us at your convenience.*