

SHADOW AI & BYO-AI RISK.



The same AI vendor's free and enterprise tiers can carry entirely different data training and retention terms. **The risk isn't whether an AI tool is reputable — it's which agreement is actually in place when an employee pastes sensitive data into it.**

- SHADOW AI
- CONSUMER VS. ENTERPRISE TERMS
- PROMPT DATA EXPOSURE
- AI ACCEPTABLE USE POLICY

01 THE MECHANISM



The Same Vendor, Two Risk Profiles

A consumer-tier AI account and a proper enterprise agreement from the same vendor often have **entirely different data handling terms**, so the question is which agreement is actually in place.



A Prompt Isn't Inspected Like a File

Traditional DLP and CASB tooling was built around file uploads and known API traffic; sensitive data typed directly into a chat box **doesn't always trigger the same inspection.**



AI Features Proliferate Inside Approved Tools

A vendor can quietly add an AI assistant to an already-sanctioned product, **potentially under different data handling terms** than what was originally reviewed.

02 THE FLOW

The Governed AI Adoption Path CONTAINED

1

Employees are proactively provisioned a sanctioned, enterprise-licensed AI tool with a **contractual no-training guarantee**, removing the incentive to reach for a personal account.

2

An acceptable use policy specifies which data classifications may enter which tier of tool, backed by **AI-aware DLP that inspects prompt content**, not just file uploads.

3

Newly discovered AI features inside existing sanctioned tools are **reviewed against the same data standard** before being enabled for the organization.

The Ungoverned Shadow AI Path EXPOSED

1

An employee pastes a customer contract into a personal, free-tier AI account for help summarizing it, unaware the **terms allow that data to be retained and used for training.**

2

A developer pastes proprietary source code, including a hardcoded API key, into an AI coding assistant for debugging help, **exposing both the code and the credential.**

3

Security has no visibility into either event, since **neither a file upload nor a recognized SaaS login occurred** — just text typed into a browser tab.

03 CORE PILLARS & COVERAGE

PILLAR	WHAT IT COVERS
AI Acceptable Use Policy	A written policy specifying which AI tools are approved and what data classifications may be entered into which tier of tool.
Enterprise AI Agreements	Contractual terms with AI vendors guaranteeing no training on submitted data, appropriate data residency, and audit rights, replacing ad hoc consumer accounts.
AI-Aware DLP & Discovery	Tooling specifically built to recognize AI chat interfaces, browser extensions, and prompt content, not just file uploads and known SaaS traffic.
Sanctioned Tool Provisioning	Proactively giving employees an approved AI tool so there's a legitimate alternative to a personal account.

04 DETECTION APPROACHES COMPARED

APPROACH	SUFFICIENT ALONE?	WHAT IT MISSES
Banning AI Tools Outright With No Sanctioned Alternative	No	Prohibition without a legitimate alternative typically pushes usage further underground rather than actually stopping it.
Traditional DLP With No AI-Specific Prompt Inspection	No	DLP built around file uploads and known API traffic doesn't necessarily inspect sensitive data typed directly into an AI chat interface.
Approving a Vendor Without Distinguishing Consumer vs. Enterprise Terms	No	The same AI vendor's consumer and enterprise tiers can have entirely different data training and retention terms.
One-Time Tool Approval, No Monitoring for New AI Features	No	AI capabilities can be added to an already-approved tool after the fact, under different data handling terms.
AI Acceptable Use Policy + Enterprise Agreements + AI-Aware DLP + Sanctioned Provisioning	Yes	Makes the sanctioned path the easy path while actually gaining visibility into what's being used.

SHADOW AI & BYO-AI RISK.

GOVERN · INSPECT · SANCTION

05 PROGRAM OPERATIONS

The risk isn't the AI vendor in the abstract, it's which specific tier of agreement is actually in place, since consumer and enterprise terms from the same company can be completely different.

Banning AI tools without a sanctioned alternative is a well-documented way to lose visibility entirely, not to actually stop usage.

A prompt typed into a chat box is functionally an unmonitored data-exfiltration channel unless DLP tooling is specifically built to inspect that kind of input.

Source code and credentials pasted into an AI coding assistant create the same exposure as any other unauthorized code disclosure, even though it doesn't look like a traditional breach.

06 CONTROL MAPPING

AI ACCEPTABLE USE POLICY

An organization's own written policy defining approved AI tools and what data classifications may be entered into each tier of tool.

ENTERPRISE AI VENDOR AGREEMENTS

Contractual terms distinguishing sanctioned, governed AI use from ad hoc consumer accounts, typically including no-training guarantees.

CASB / AI-AWARE DLP TOOLING

The technical control layer for discovering AI tool usage and inspecting prompt content, extending traditional CASB and DLP capability.

NIST AI RISK MANAGEMENT FRAMEWORK

General AI governance guidance relevant to how organizations should assess and manage risk from AI tool adoption broadly.

07 WHAT THIS DOES NOT SOLVE

Enforcement Without Technical Controls

A written policy alone **doesn't stop an employee from using a personal AI account** on a personal device outside any monitored channel.

AI Output Accuracy or Bias

Governing what data goes into an AI tool says **nothing about whether the tool's output is actually correct** or appropriate to rely on.

Vendor Practices Changing Over Time

An enterprise agreement's terms are only as good as the **vendor's actual practices**, which can shift as products and policies evolve.

Approved Tools Used Inappropriately

A sanctioned AI tool with proper enterprise terms can still be misused if an employee **enters data it was never intended to handle**.

08 MYTHS VS. REALITY

MYTH	REALITY
"Banning AI tools solves the shadow AI problem."	Prohibition without a sanctioned alternative typically pushes usage further underground , losing visibility rather than gaining it.
MYTH	REALITY
"If the AI vendor is reputable, the data is safe."	The same vendor's consumer and enterprise tiers can have completely different data training and retention terms .
MYTH	REALITY
"Our DLP already catches sensitive data leaving the organization."	DLP built around file uploads and known API traffic doesn't necessarily inspect text typed directly into an AI chat interface.
MYTH	REALITY
"AI risk is only about the tools we've explicitly approved."	AI features can be silently added to already-approved SaaS products , under data handling terms the organization never reviewed.
MYTH	REALITY
"Pasting code into an AI assistant for debugging help is harmless."	Source code and embedded credentials pasted into a third-party AI tool create real exposure outside the organization's control .

09 SPEAK THE PROGRAM — KEY TERMS

SHADOW AI

Unsanctioned use of AI tools by employees without security or IT oversight, the **AI-specific instance of the broader shadow IT problem**.

CONSUMER VS. ENTERPRISE AI TIER

The distinction between a free or individual AI account and a contracted enterprise agreement, which can carry **entirely different data terms**.

AI ACCEPTABLE USE POLICY

A written policy defining which AI tools are approved and **what data may be entered** into each tier of tool.

PROMPT INSPECTION

AI-aware DLP capability that **inspects text typed into an AI chat interface**, not just file uploads or known API traffic.

NO-TRAINING GUARANTEE

A contractual commitment from an AI vendor that submitted data **will not be used to train the underlying model**, typical of enterprise agreements.

AI FEATURE CREEP

The addition of AI capabilities to an already-approved SaaS product after its original review, **potentially under different data handling terms**.