Yes, you have Shadow AI.

This is how you fix it.

The Insight

The catalog is the middle ground CIOs are looking for.

Departments bypass IT and Security **because the approved path is slower than the unapproved one.**

Two models. One underlying asset.

Centralized:
Safer, but too slow for AI velocity.

Decentralized:
Fast, but breeds shadow IT and audit exposure.

The catalog:
Curated tools. Tiered ownership. Both problems solved.

1/2 of your employees are using personal AI apps to perform work.*

Both centralization and decentralization depend on the same underlying asset: a curated AI service catalog with tiered ownership. It's the solution every governance approach needs.

* Netskope cloud telemetry (2026) suggests ~47% of GenAI users rely on personal AI apps, outside corporate monitoring

The Insight

Speed of approval is now the single biggest driver of shadow AI.

Every month without a catalog is a month where departments standardize on unmanaged tools, unvetted contracts, and ungoverned data flows.

82% of enterprises uncovered shadow AI agents in the last year, often in internal automation and LLM platforms.*

CIOs who treat **shadow AI** as a culture problem spend months writing policy. CIOs who treat it as an **operational opportunity** ship a catalog in weeks.

*Cloud Security Alliance, April 20, 2026.

Step 1. Stand up the catalog in 60 days.

A catalog turns an ownership argument into a procurement workflow.

Approved tooling

A curated set of LLMs, copilots, vector stores, and connectors that meet security baselines.

Enterprise contracts

DPAs, BAAs, data residency, and model training clauses negotiated once at scale.

Built-in controls

SSO, MFA, central logging, retention, and DLP wired into every tool by default.

One intake front door

A published request form with a service-level commitment for review and response.

Step 2. Tier the catalog. Match ownership to risk.

Three tiers. Three ownership models. One catalog.

TIER 1
Enterprise platforms

IT owns fully

Productivity AI used by everyone. IT procures, secures, and contracts. Mandatory for general use cases.

Examples: *M365 Copilot, enterprise ChatGPT, internal copilots.*

TIER 2
Domain and BU tools

Co-owned

IT approves from a security-reviewed list. Business unit funds, operates, and owns the use case.

Examples: *Coding copilots, marketing content tools, legal review aids.*

TIER 3
Sandbox and pilots

BU-led, IT-gated

Low-risk data only. 60 to 90 day review. Tools graduate to Tier 2, converge with Tier 1, or retire.

Examples: *New vendors, experiments, proof-of-concept use cases.*

Step 3. Install five non-negotiable guardrails.

01	**One identity plane**	SSO and MFA for every AI vendor. No local accounts, no shared logins.
02	**Central logging**	Who used which tool, against which data class, for which integration.
03	**Data classification map**	Green / yellow / red tags with concrete examples. Published and enforced.
04	**Standard contract clauses**	No training on your data. Clear retention. Export on exit. Regional residency.
05	**Fast exception path**	A documented, time-bound process for legitimate edge cases outside the catalog.

The Work

Step 4. Federate the governance.

A small steering group outperforms a large committee.

Stand up a cross-functional group: IT, security, legal, risk, business units.

Meet monthly to approve new catalog entries, review Tier 3 pilots, and retire overlapping tools.

Publish an AI playbook so managers know what is approved, how to request additions, and what data must stay out of generative tools.

Who sits at the table

- CIO or IT leadership
- CISO or security lead
- Legal and privacy
- Risk and compliance
- Business-unit leaders
- Finance as a standing observer

The Compliance Reality

Federation has hard limits for regulated data.

65% of organizations experienced an AI-agent-related incident last year, most commonly data exposure (61%) and operational disruption (43%) *

For organizations in CMMC, HIPAA, PCI, or SOC 2 scope, CUI, PHI, and PII must be restricted to Tier 1 tools only.

Department-level purchasing of generative AI for regulated data creates contract breach risk and audit failure.

Central logging is the audit artifact.

It proves which data left your environment, which tool it went to, and which user sent it. Without it, you cannot defend scope, close customer security reviews, or pass a CMMC, HIPAA, or SOC 2 audit.

*Cloud Security Alliance, April 20, 2026.

We build the catalog and the governance.

We design the operating model based on your business needs and compliance obligations.

We facilitate the steering group until it runs on its own

AI governance program design

Catalog, intake, tiering, RACI, and playbook built for your business model and speed targets.

Compliance-mapped data classification

Green / yellow / red rules aligned to CMMC (CUI), HIPAA (PHI), PCI, GLBA, and GDPR obligations.

Contract clause library

Standard AI vendor language on training, retention, export, residency, and indemnification.

Steering group facilitation

CISO-led decision rights and a monthly operating cadence that converts shadow AI into a managed portfolio.

Let's discuss your AI catalog.

ResilientTech Advisors

inquiries@resilienttechadvisors.com
(804) 621-1136
resilienttechadvisors.com

Scan to send us a message.

Cybersecurity that accelerates your business.