CS-113 · AI Security

Shadow AI

The unsanctioned use of AI tools inside an organization, outside of governance and oversight.

Executive Summary

Shadow AI is the use of artificial intelligence tools by employees or teams without the knowledge, approval, or oversight of the organization. People adopt these tools to work faster, but doing so outside governance can leak sensitive data, create compliance gaps, and introduce ungoverned automated decisions. Managing it is less about banning tools and more about providing safe, sanctioned options and clear rules.

What It Is

Shadow AI is the AI-specific form of shadow IT, the long-standing pattern of staff using technology that has not been reviewed or approved. It includes pasting company information into a public chatbot, adopting an unvetted AI feature inside an everyday app, connecting an outside AI service to internal data, or building a small automation that quietly makes decisions. The tools are often genuinely useful, which is exactly why adoption outpaces policy. The problem is not the technology itself but the absence of oversight: no one has assessed the data being shared, the vendor holding it, the accuracy of the output, or the risk of the decisions being made.

Why It Matters

When AI use happens in the shadows, the organization loses visibility into where its data goes and how automated outputs are used. Confidential documents, customer records, source code, and strategy can end up in an outside service that may retain or reuse them. Decisions influenced by unreviewed AI output can be wrong, biased, or impossible to explain later, which is a serious problem in regulated activities. Because the organization cannot govern what it cannot

see, shadow AI undermines data protection, contractual commitments, and emerging AI regulations at once. For governance, risk, and compliance professionals, it is one of the most pressing everyday AI risks, precisely because it grows through ordinary, well-meaning behavior rather than through an attack.

How It Works

Shadow AI spreads through convenience. An employee facing a deadline reaches for whatever tool helps, often a free public one, and shares whatever information gets the job done. Others copy the habit, and AI features increasingly appear inside tools the organization already uses, so adoption happens without any deliberate decision. The exposure accumulates quietly: data leaves the organization, outputs feed into work products, and small automations start influencing real outcomes, all without a record. Governing it means discovering what is actually in use, offering approved tools that meet the same needs safely, setting clear and realistic rules about what data may be used and how, training people on why the rules exist, and monitoring for continued unsanctioned use rather than assuming a policy alone will stop it.

Architecture Diagram

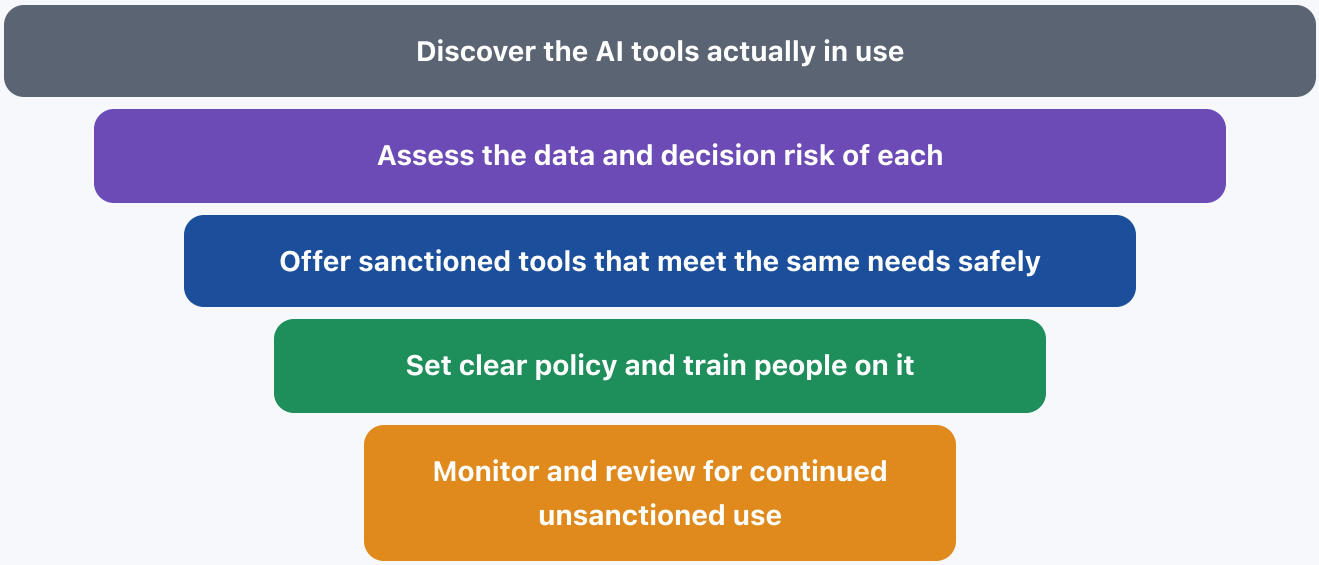

Bring shadow AI into the light: discover what is in use, offer safe sanctioned options, set clear rules, and monitor over time.

Visual Workflow

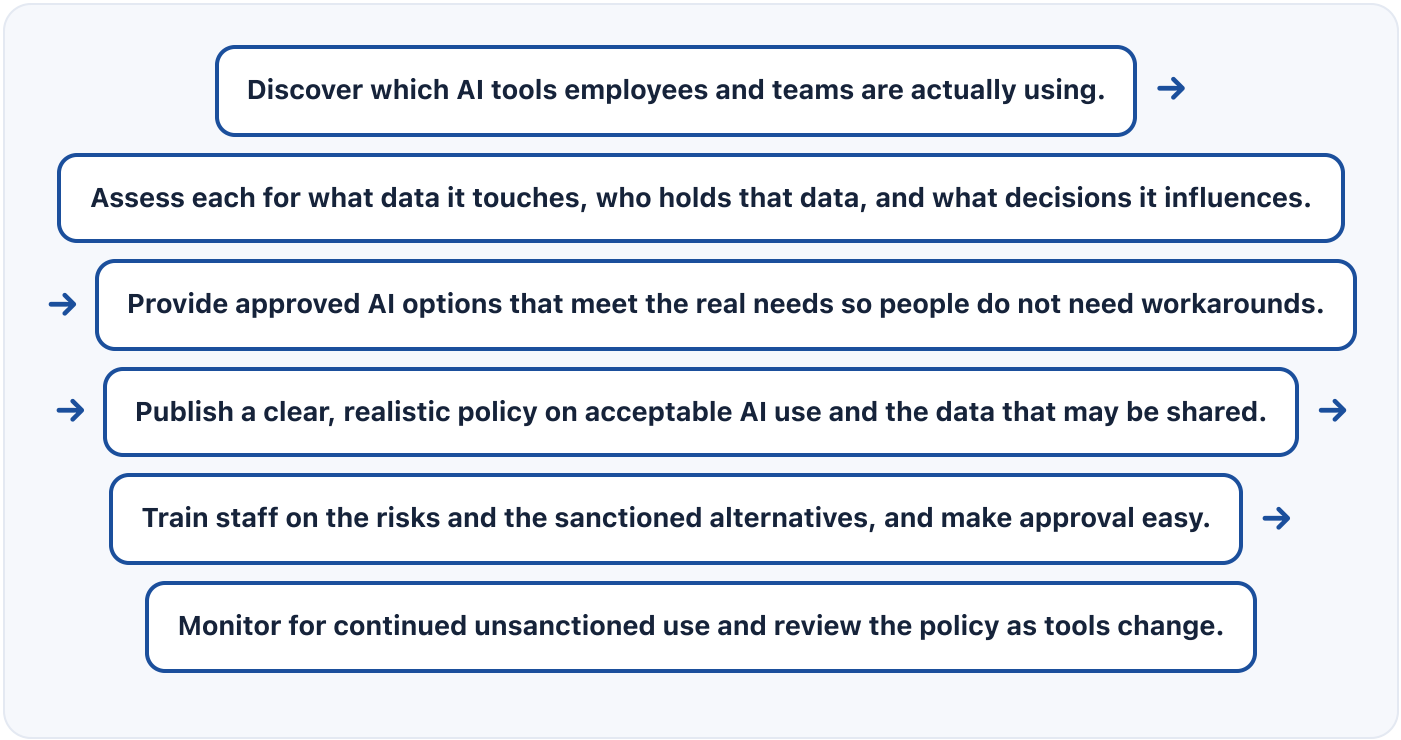

Common Attacks

- Sensitive data leaking into an outside AI service that may retain or reuse it
- Confidential output from a public tool being exposed or shared beyond intended limits
- An unvetted AI vendor becoming an uncontrolled path for company information
- Ungoverned automated decisions that are biased, wrong, or impossible to explain
- Attackers targeting the weak, unmonitored AI tools staff adopt on their own

Common Mistakes

- Banning AI outright, which pushes usage further underground
- Assuming a written policy alone will change everyday behavior
- Never discovering what tools are actually in use across the organization
- Offering no sanctioned alternative, so staff keep reaching for unapproved tools
- Ignoring AI features quietly added inside apps the organization already uses

Best Practices

- Discover and inventory the AI tools already in use before setting rules
- Provide safe, approved options that genuinely meet employee needs

- Write clear, realistic policy about acceptable use and the data that may be shared
- Classify data so people know what must never go into an outside tool
- Train continuously and make requesting a new tool simple and fast
- Review vendors, monitor usage, and update the policy as the tool landscape shifts

Quick Checklist

- [x] An inventory of AI tools in use exists and is kept current
- [x] Sanctioned AI options are available and meet common needs
- [x] A clear acceptable-use policy defines what data may and may not be shared
- [x] Data classification tells people what is off-limits for outside tools
- [x] Staff are trained and the approval path is easy to use
- [x] Usage is monitored and vendors are reviewed on a schedule

Recommended Tools

AI and SaaS discovery tooling
Reveals which AI services and features are in use across the organization

Data loss prevention (DLP)
Detects and limits sensitive data flowing to outside services

Cloud access security broker (CASB)
Applies policy to cloud and AI application usage

Sanctioned enterprise AI platform
Gives staff an approved tool with controls in place

Industry Standards

NIST AI Risk Management Framework (AI RMF, AI 100-1)
Its Govern function supports policy, oversight, and accountability for AI use

ISO/IEC 42001
Standard for an AI management system that formalizes governance of AI use

ISO/IEC 27001
Anchors the information security controls that shadow AI can bypass

Career Relevance

Shadow AI sits squarely in the work of AI governance analysts, who build the policies and inventories that bring it under control, and AI risk managers, who weigh its exposure across the business. GRC analysts assess it during data protection, vendor, and compliance reviews, and AI security engineers help discover and constrain the tools involved. Because it is a governance-first risk driven by human behavior, it is a natural fit for the AI-Governance-Jobs.com audience and a common topic in interviews for these roles.

Interview Questions

- What is shadow AI, and how is it related to shadow IT?
- Why can banning AI tools make the shadow AI problem worse?
- What are the main data protection risks when staff use unsanctioned AI?
- How would you build an inventory of AI tools already in use?
- How do sanctioned tools and clear policy work together to reduce shadow AI?

Related Certifications

ISACA AI governance offerings (as available)
ISO/IEC 42001 lead implementer training (as available)
IAPP privacy certifications (for the data angle)

Further Reading

- NIST AI Risk Management Framework
- OWASP Top 10 for LLM Applications
- CISA: Cybersecurity Best Practices

Key Takeaways

- Shadow AI is unsanctioned AI use that grows through convenience, not attack.
- Its biggest risks are data leakage, vendor exposure, and ungoverned automated decisions.
- You cannot govern what you cannot see, so discovery comes first.
- Safe sanctioned tools plus clear, realistic policy beat outright bans.
- It is a governance-first risk that AI governance, risk, and GRC roles are built to manage.

FAQ

▶ **Is shadow AI always malicious?**

▶ **Should we just block all outside AI tools?**

▶ **How do we even find shadow AI?**

Related Careers

RELATED ROLES

Ai Governance Analyst

Ai Risk Manager

Grc Analyst

Ai Security Engineer

RELATED CERTIFICATIONS

ISACA AI governance offerings (as available)

ISO/IEC 42001 lead implementer training (as available)

IAPP privacy certifications (for the data angle)

CURRENT OPENINGS

Live roles are on the job board.

Browse all jobs

GUIDES & SALARY

Career guides

Role roadmaps

Salary data

Career tools

SUGGESTED LEARNING PATH

1. Ground the basics with CS-001 Cybersecurity
2. Study this sheet: **Shadow AI**
3. Go deeper: Secure AI Adoption
4. Go deeper: Prompt Injection
5. Validate it: work toward **ISACA AI governance offerings (as available)**
6. Find the role: browse current openings

RELATED SHEETS

CS-115 Secure AI Adoption
CS-111 Prompt Injection
CS-116 OWASP Top 10 for LLM Applications
CS-001 Cybersecurity
CS-104 NIST Cybersecurity Framework (CSF)