Shadow AI in the Enterprise: Governing the AI Tools Your Organisation Didn’t Approve

Cetus AI — White Paper
June 2026
hello@cetusai.com.au | cetusai.com.au

Executive Summary

Most Australian organisations are operating under a significant and growing governance blind spot: the AI tools their employees are using every day that were never reviewed, approved, or even known to IT and security teams. This phenomenon — commonly called Shadow AI — represents one of the fastest-growing compliance risks in enterprise technology today.

Unlike Shadow IT of previous decades, Shadow AI carries a qualitatively different risk profile. When an employee uses an unapproved SaaS application, the primary concern is data leakage. When an employee uses an unapproved AI tool, the risks extend to privacy law exposure under the Australian Privacy Principles, potential breach of the APS AI Policy for government agencies, regulatory liability under the Corporations Act for directors, and the reputational consequences of AI-generated outputs that cannot be attributed, audited, or recalled.

This white paper examines the Shadow AI risk landscape for Australian enterprise and government, the regulatory frameworks that create direct exposure, and how a structured discovery-declare-govern workflow — as implemented in the Songlines Control Shadow AI module — closes the gap between what organisations know about their AI usage and what is actually happening.

1. The Scale of the Problem

The proliferation of consumer and prosumer AI tools has outpaced the ability of most organisations to govern them. Tools such as ChatGPT, GitHub Copilot, Grammarly AI, Midjourney, Notion AI, and dozens of others are available to any employee with a credit card or a free account. Many are actively used for tasks that involve sensitive organisational data — drafting contracts, summarising board papers, generating code from internal specifications, or analysing customer data.

Research from multiple enterprise security vendors in 2025 found that the average large organisation has between 40 and 80 AI tools in active use that are not sanctioned by IT or

security. In financial services and government, where data classification requirements are most stringent, the gap between sanctioned and actual AI tool usage is often widest — precisely because the formal procurement process is slowest.

The challenge is not that employees are acting maliciously. In most cases, they are using the most effective tools available to them to do their jobs. The challenge is that organisations have not built the governance infrastructure to keep pace with the rate at which AI capabilities are becoming embedded in everyday work.

---

2. Why Shadow AI Is Different from Shadow IT

Shadow IT governance frameworks developed over the past decade focused primarily on data residency, access controls, and contractual protections. The question was: does this application store our data somewhere we haven't approved, and can we get it back if we need to?

Shadow AI introduces three additional dimensions that existing Shadow IT frameworks do not address.

Inference as data processing. When an employee submits a prompt to an AI model, they are not simply storing data — they are processing it. The model may use that data to generate outputs that are shared externally, stored in training pipelines, or made available to other users. The Australian Privacy Principles treat this as a disclosure of personal information if the data includes information about identifiable individuals, regardless of whether the employee intended to share it.

Output attribution and liability. AI-generated outputs — code, contracts, analysis, communications — may be used in consequential decisions without any record of how they were produced. When something goes wrong, the organisation cannot reconstruct the AI interaction, cannot demonstrate that appropriate oversight was applied, and cannot establish whether the output was reviewed before use. This creates direct exposure under the Corporations Act for directors who have approved AI use policies without the governance infrastructure to enforce them.

Regulatory inventory requirements. The APS AI Policy 2024 requires Australian Government agencies to maintain an inventory of AI tools in use and to assess each tool against a risk framework before deployment. An agency that cannot account for the AI tools its employees are using is in direct breach of this requirement, regardless of whether any specific harm has occurred.

---

3. The Regulatory Exposure

3.1 Privacy Act 1988 — Australian Privacy Principles

The Australian Privacy Principles (APPs) apply to any organisation with an annual turnover above $3 million, all Australian Government agencies, and certain other entities regardless of size. APP 6 governs the use and disclosure of personal information: an organisation must not use or disclose personal information for a purpose other than the primary purpose for which it was collected, unless an exception applies.

When an employee submits personal information — customer records, employee data, patient information — to an unapproved AI tool, the organisation has disclosed that information to a third party without the individual's consent and without a contractual framework governing how the third party will handle it. This is a direct APP 6 breach. The Office of the Australian Information Commissioner (OAIC) has signalled that AI-related privacy breaches will be a priority enforcement area from 2025 onwards.

3.2 APS AI Policy 2024

The Australian Public Service AI Policy, updated in 2024, requires agencies to apply a risk-based approach to AI adoption. Key requirements include maintaining a register of AI tools in use, conducting risk assessments before deploying AI in significant decisions, and ensuring that AI use is consistent with the APS Values and Code of Conduct. An agency that has employees using AI tools that are not on the register — and that has not conducted risk assessments for those tools — is non-compliant with the policy, irrespective of whether any specific harm has occurred.

3.3 ISO/IEC 42001 — AI Management System Standard

ISO/IEC 42001, published in 2023, is the international standard for AI management systems. It requires organisations to establish, implement, maintain, and continually improve an AI management system that includes policies for AI use, risk assessments for AI applications, and controls to manage AI-related risks. Shadow AI — by definition — falls outside the scope of any AI management system that does not have a mechanism for discovering and governing unapproved tools.

3.4 Corporations Act 2001 — Director Liability

Directors of Australian companies have a duty of care and diligence under section 180 of the Corporations Act. Where an organisation has adopted AI use policies but lacks the governance infrastructure to enforce them, directors may face personal liability if AI-related harm occurs that could have been prevented by adequate oversight. The Australian Securities and Investments

Commission (ASIC) has identified AI governance as an emerging area of director liability in its REP 798 report on technology governance in financial services.

---

4. A Structured Governance Framework

Closing the Shadow AI gap requires more than a policy. Policies without enforcement mechanisms are ineffective against the rate at which new AI tools become available and the ease with which employees can access them. An effective Shadow AI governance framework has four components: discovery, declaration, assessment, and enforcement.

4.1 Discovery

Discovery is the process of identifying AI tools that are in use across the organisation, whether or not they have been approved. This includes network-level scanning for known AI endpoints, analysis of browser extension and application inventories, and review of expense claims and procurement records for AI-related subscriptions.

The Songlines Control Discovery Scanner maintains a database of 30+ known AI vendor endpoints and can simulate a network scan to identify which tools are likely in use. Discovery results feed directly into the inventory for review and risk assessment.

4.2 Declaration

Declaration is the process by which employees self-report the AI tools they are using. A self-service declaration workflow reduces the burden on IT and security teams, creates a culture of transparency rather than prohibition, and generates a more complete inventory than discovery alone — because employees often use tools that do not generate detectable network traffic.

The Songlines Control employee self-service portal allows any employee to look up the governance status of an AI tool and submit a declaration without logging in. The portal is designed to be shared organisation-wide as a link, reducing friction to the point where declaration becomes the path of least resistance.

4.3 Risk Assessment

Every declared or discovered tool should be assessed against a consistent risk framework before a governance decision is made. The Songlines Control risk scoring engine evaluates each tool across four dimensions:

Dimension	Description
Data sensitivity	Whether the tool handles PII, financial, health, or confidential data
External data sharing	Whether data submitted to the tool is processed or stored outside the organisation
Certifications	SOC 2, ISO 27001, CSA STAR — presence reduces the risk score
Vendor trust	Known enterprise vendors with established security programs receive a lower baseline score

Scores range from 0 to 100, producing four risk levels: Low (0–29), Medium (30–59), High (60–79), and Critical (80–100). The risk score is calculated automatically at the point of declaration and updated when tool details change.

4.4 Enforcement

Governance decisions — approve, conditionally approve, reject, or require further review — should be enforced consistently across the organisation. Enforcement has two dimensions: policy rules that determine what action to take based on risk level or tool category, and integration with the AI gateway to enforce those rules at the point of use.

The Songlines Control AI Risk Policy Builder allows organisations to create rules that automatically block, warn, require approval, or notify administrators when employees attempt to use tools that fall into specific risk categories. Rules are evaluated in priority order and can be scoped to specific data types, tool categories, or risk levels.

---

5. The Songlines Control Shadow AI Module

The Songlines Control Shadow AI module implements the four-component governance framework described above as an integrated, enterprise-grade capability within the Songlines Control platform.

Inventory dashboard. A searchable, filterable inventory of all declared and discovered AI tools, with risk scores, status badges, and sortable columns. Administrators can review, approve, reject, or conditionally approve tools from a single interface.

Declare a Tool wizard. A four-step declaration workflow — Tool Details, Data and Privacy, Compliance, and Review — that guides employees through the information required for a complete risk assessment. A live risk preview updates as the employee completes each step.

Discovery Scanner. A simulation of network-level scanning against the 30+ vendor endpoint database, with results fed directly into the inventory for review.

Sanctioned AI Catalogue. A browsable registry of approved tools that employees can reference before adopting a new AI capability. Employees can request access to catalogue tools directly from the portal.

AI Risk Policy Builder. An org-level rule engine with five enforcement actions (Block, Warn, Require Approval, Notify Admin, Allow) configurable by risk level, tool category, or data type.

Employee self-service portal. A public-facing page at `/shadow-ai/portal` where employees can check tool status and submit declarations without logging in to the platform.

Bulk CSV import. Pre-populate the inventory from an existing tool register or spreadsheet using a structured CSV template.

Weekly digest email. Org admins and owners receive a Monday summary of pending reviews, high-risk tools, and recent governance decisions.

PDF report export. A full Shadow AI governance report including the complete inventory, risk score breakdown, policy rules in force, and a compliance summary for Privacy Act 1988 and APS AI Policy 2024.

Guardrail Engine integration. Shadow AI policy rules can be linked to the Songlines Control Guardrail Engine, enforcing governance decisions at the API gateway level for tools that route through the platform.

---

6. Getting Started

Organisations typically begin their Shadow AI governance journey with three steps: deploying the Discovery Scanner to establish a baseline inventory, sharing the employee self-service portal link organisation-wide to seed declarations, and configuring a small number of high-priority policy rules — typically a block on Critical-risk tools and a require-approval rule for High-risk tools handling PII.

From this baseline, the governance programme can be expanded incrementally: adding tools to the Sanctioned Catalogue, refining policy rules as the inventory matures, and integrating with the Guardrail Engine as AI workloads are instrumented through the Songlines Control platform.

Conclusion

Shadow AI is not a problem that can be solved by policy alone. The rate at which AI tools are becoming available, and the ease with which employees can access them, means that governance must be built into the workflow — not bolted on after the fact. The Songlines Control

Shadow AI module provides the discovery, declaration, assessment, and enforcement infrastructure that Australian enterprise and government organisations need to close the gap between their AI use policies and the reality of how AI is being used across their organisations.

About Cetus AI

Cetus AI is an Australian AI governance company. Songlines Control® is the enterprise AI observability and governance platform built for Australian enterprise and government — delivering real-time cost control, sovereign compliance, and complete auditability across every AI interaction.

Contact: hello@cetusai.com.au
Website: cetusai.com.au
Platform: www.songlinesai.com