Shadow AI

The Invisible Threat from your Most Productive Employees

John Wunderlich, Chief Privacy Officer, JLINC

VIPSS 2026 | Session 6G

Agenda

What we will talk about

- What is Shadow AI?
- The Productivity Paradox
- Risks
- Next Steps

What we won’t talk about

These are real issues, but not for this session

- Environmental Costs of AI
- Data Sweatshops
- IP and training data

What is Shadow AI?

Understanding the scope of the problem

WHAT IS SHADOW AI?

> *The use of AI tools, models, and services by employees without explicit organizational approval, oversight, or governance — extending shadow IT into a domain where data flows are opaque and outputs are non-deterministic.*

Shadow IT

Unapproved software
and hardware.
Known risk category.
Established detection.

Shadow AI

Unapproved AI tools
with opaque data flows.
Non-deterministic outputs.
Harder to detect.

Sanctioned AI

Governed AI tools
with data controls.
Audit trails.
Aligned to policy.

Opaque Data Flows

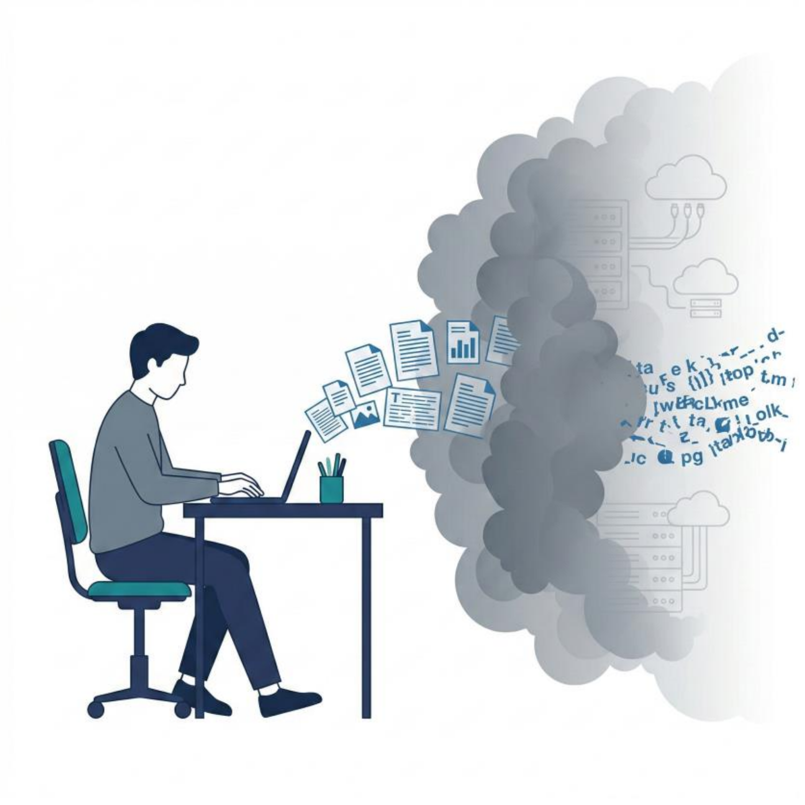

- Unknown channels through your perimeter
- No audit trails – no records of processing activities
- You remain legally responsible

Non-deterministic outputs

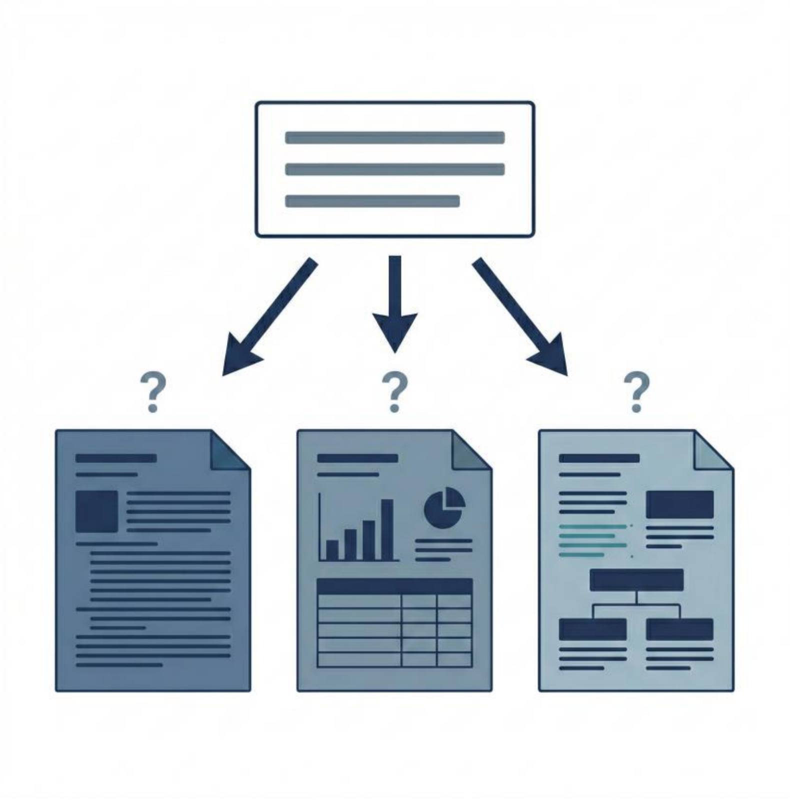

- Payroll calculations are deterministic
- Credit scores are based on probabilistic models
- AI queries to summarise or analyze contracts can produces different outputs

Harder to Detect

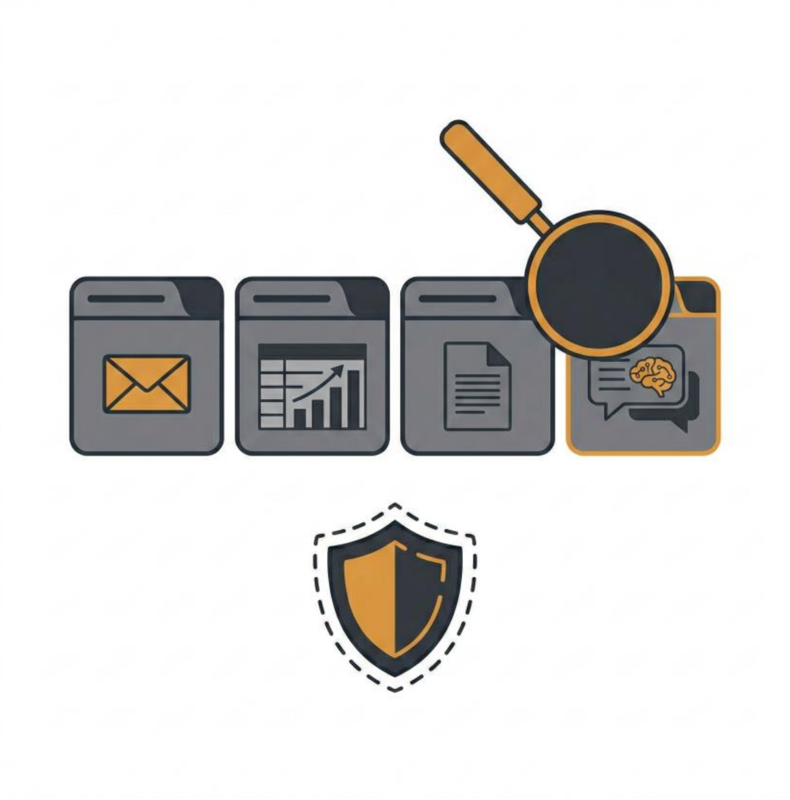

- No inventory or software management
- Missing security tooling
- Detection only possible after an incident.

The Productivity Paradox

Why do employees choose to use AI.

Well Intentioned Dedicated Employees

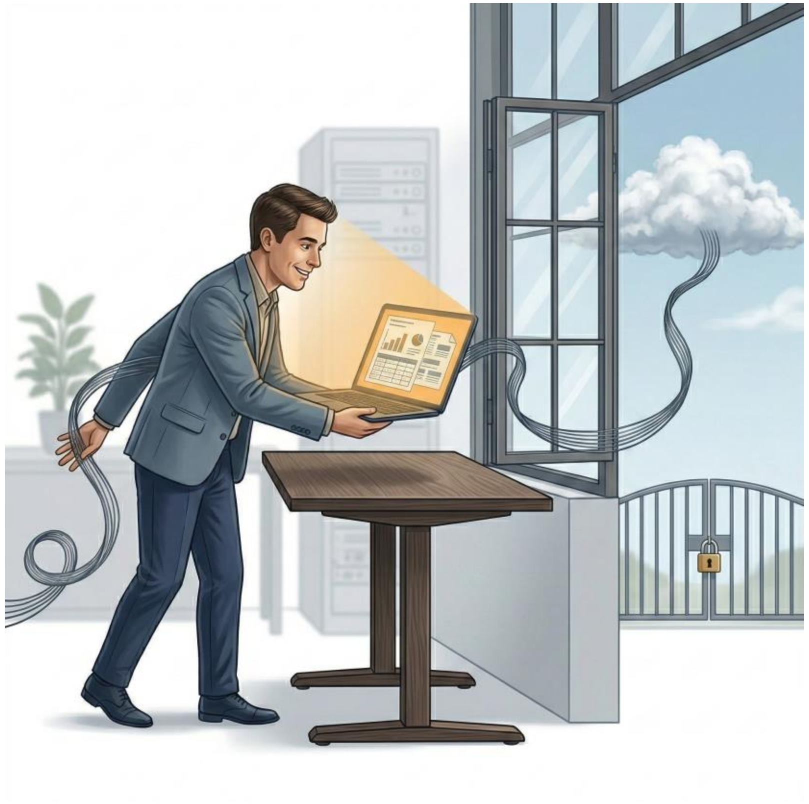

The WIDE error:

The use of unauthorised AI tools by employees to do their job better, faster, or more thoroughly — and the organization has not provided a governed alternative that meets their needs.

Why employees use Shadow AI

Shadow AI is a demand signal. It tells you where your governance has created a vacuum that employees are filling themselves.

The question is not how to stop them. The question is how to meet the need they're expressing.

Tool Friction

- Approved tools are slow, limited, hard to access, or nonexistent

Productivity Pressure

- Deadlines outpace governance approval timelines

AI Capability gap

- IT hasn't provided AI tools that match employee needs

Blanket bans backfire

- Prohibition drives usage underground, not away

Well Intentioned? Dedicated Soldier

General Jack Ripper

- Acting in the national interest
- Has the tools to bypass command
- No personal gain

Result: Nuclear Confrontation

Risks

What’s at stake for your organization?

The Risk Landscape

Data Leakage

46% of organizations report internal data leaks through GenAI

IP Exposure

Proprietary code, strategy docs, and trade secrets pasted into public models

Compliance Gaps

Ungoverned use violates data protection, sectoral, and contractual obligations

Model Reliance

Hallucinated outputs entering business decisions without validation

Audit Blindness

No logs, no lineage, no way to respond to regulator inquiries

Third-Party Risk

Free-tier AI tools with permissive data retention and training policies

Agentic AI amplifies every risk

Agentic AI systems don't just respond to prompts — they take autonomous action within systems. They can browse, query databases, send emails, and execute code.

When these agents operate outside governance, the blast radius of a single incident expands from data exposure to active system manipulation.

Prompt-based AI

Data in → Data out

Agentic AI

Data in → Actions taken

Shadow Agentic AI

Unknown data → Unknown actions

Identity governance for AI entities: only 48% of organizations have controls

Navigating a Regulatory Patchwork

DEAD	**AIDA (Bill C-27)**	Died with prorogation Jan 2025. Minister signalled "light, tight, right" replacement — AI regulation will be decoupled from privacy reform.
IN FORCE	**PIPEDA**	Still the federal private-sector privacy law since 2000. No AI-specific provisions, but consent and accountability principles apply to AI data processing.
IN FORCE	**Quebec Law 25**	Automated decision-making transparency requirements. Right to explanation. Fines up to $25M or 4% global turnover.
UPDATED	**TB Directive on ADM**	Federal public-sector directive. Updated 2025 with stronger accountability, mandatory Algorithmic Impact Assessment. Compliance deadline June 2026.
PUBLISHED	**CAN/DGSI 101:2025**	National standard for ethical AI design and use. Risk management, ethics by design, continuous monitoring. Tailored for organizations <500 employees.

Also: Voluntary Code of Conduct on Generative AI | Canadian AI Safety Institute (CAISI) | Alberta PIPA, BC PIPA

Addressing a Governance Gap

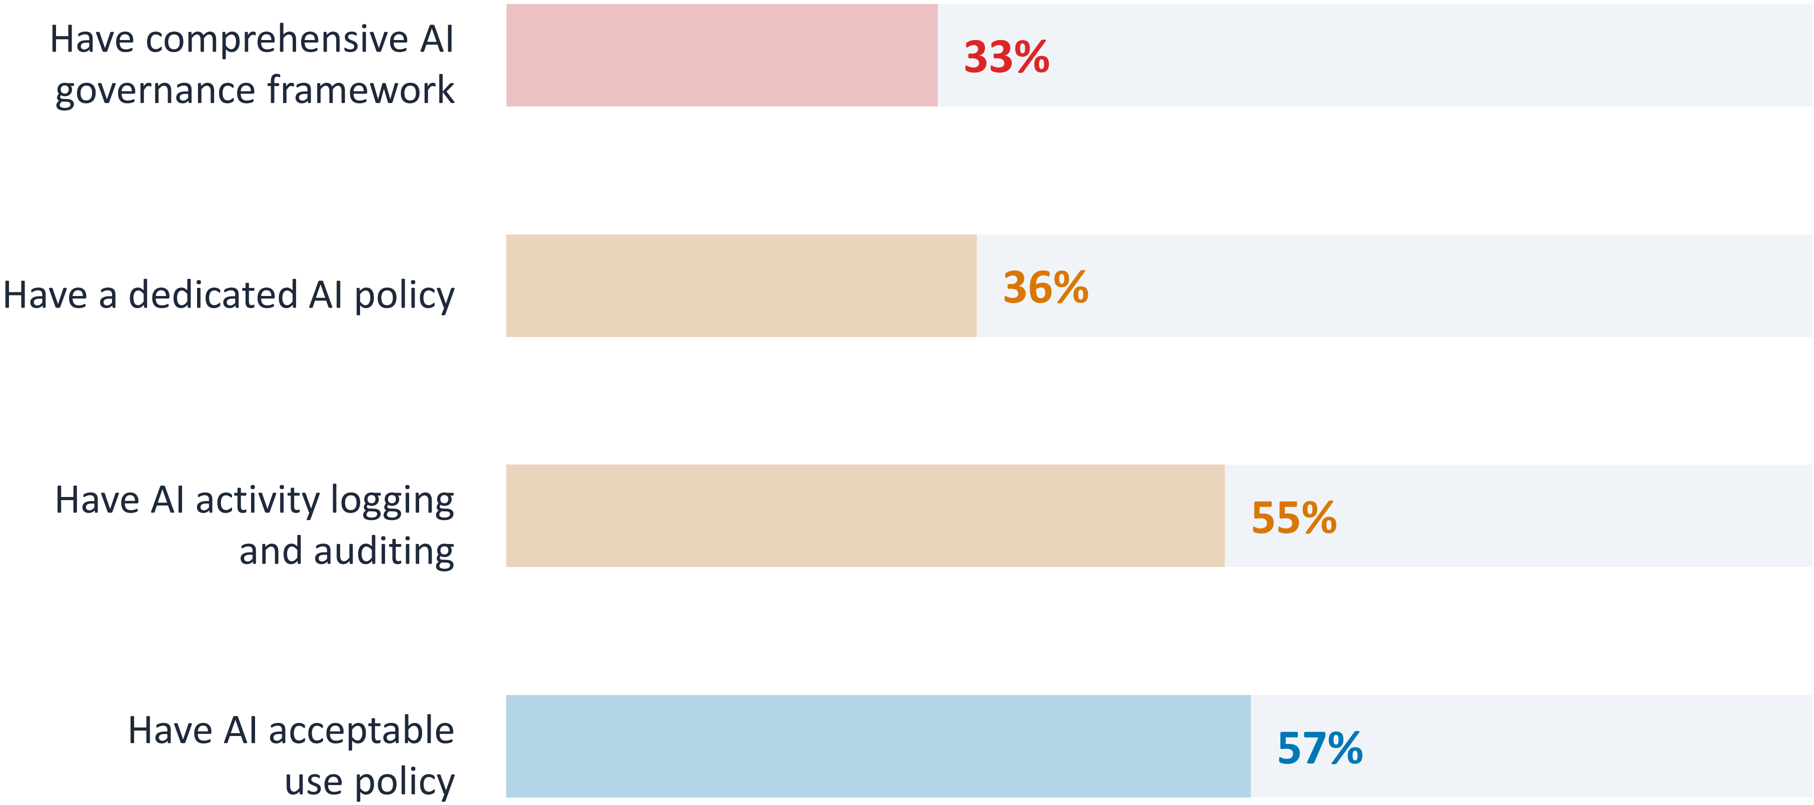

Sources: [6] ISACA 2025, [7] S&P Global Research 2025

Next Steps

Mitigating risk to effectively use AI

1. Technical Detection

AI Discovery & Inventory

Automated scanning for AI tool usage across cloud services, endpoints, and browser extensions. Build a living inventory.

Network & DLP Monitoring

Monitor traffic to known AI endpoints. Deploy DLP rules to detect and block sensitive data in prompts and uploads.

AI Gateway / Control Plane

Route all AI traffic through a central gateway. Apply policy, log interactions, and enforce data classification rules.

Endpoint & Browser Controls

Browser extension policies, managed device restrictions, and application allow-listing to control access at the edge.

2. Organisational Governance

AI Governance Committee

Cross-functional body: privacy, security, legal, business, HR. Centralized policy, federated execution. Meets monthly minimum.

Acceptable Use Policy

Approved tool catalogue, data sharing rules, prohibited use cases. Aligned with CAN/DGSI 101:2025 and PIPEDA principles.

Risk-Based Classification

Tier AI use by data sensitivity and impact. Map to privacy impact assessments required under Quebec Law 25.

Training & Awareness

Mandatory AI literacy program. Not just rules — teach employees why governance matters and how to use sandboxes.

3. Board-Level Accountability

Shadow AI is not an IT problem. It is a governance risk that belongs at the board table.

- CPO/CISO reporting directly on AI risk metrics
- AI risk integrated into enterprise risk register
- Quarterly shadow AI exposure reports to audit committee
- Fiduciary duty framing: directors' duty of care extends to AI governance

4. Zero-Trust Architecture for AI

Never Trust, Always Verify

Every AI interaction authenticated. No implicit trust based on network location or prior access.

Least-Privilege Access

AI tools get only the data access required for each specific task. No blanket permissions.

Continuous Monitoring

Real-time logging of all AI prompts, responses, and data flows. Anomaly detection on usage patterns.

Granular Conditional Access

Restrict by user role, device posture, data classification, and feature (e.g., block file uploads to AI tools).

5 steps to a governed AI capability

01

Discover

Audit your AI landscape. You can't govern what you can't see.

02

Assess

Risk-tier by data sensitivity and impact. Map to PIAs.

03

Govern

Policy, committee, acceptable use. Align to CAN/DGSI 101:2025.

04

Enable

Deploy sandboxes. Enterprise AI with DLP, logging, and guardrails.

05

Monitor

Continuous detection. Zero-trust. Quarterly board reporting.

Key Takeaways

- Shadow AI is a demand signal, not just a threat. Meet the need.
- Start with detection. You can't govern what you can't see.
- Move from sanctions to providing governed AI alternatives that work – starting with sandboxes you control.
- Board accountability is non-negotiable. This is a fiduciary issue.
- Canada's regulatory patchwork means you must govern proactively — don't wait for AIDA's replacement.

Thank you!

Q & A

john@jlinc.com