thinkia

GOVERNANCE AND RISK · INTELLIGENCE ERA 2026

Shadow *AI*.

The invisible risk is not that your employees use AI. It is that they use it where you cannot see it. This is not a problem of prohibition: it is a problem of governance.

AUTHOR'S LETTER

AI has already come in through the back door. The question is *who governs it.*

When an executive committee asks me when their company will start using artificial intelligence, my answer usually makes the room uncomfortable: it already does, every day, and almost no one there knows exactly how. Adoption did not wait for the strategy. It arrived from the bottom up, in browser tabs and private conversations with models that live outside the corporate perimeter. We call this Shadow AI, and it is the accelerated version of an old acquaintance: shadow IT.

The difference is speed and surface area. Shadow IT required installing something; Shadow AI only requires a free account and an idea for saving half an hour. That is why it spreads faster and leaves a lighter trail. And that is why the reflex to ban it, understandable as it is, is also the most counterproductive: what is prohibited does not disappear; it becomes invisible.

This document is not meant to alarm. It is meant to reframe. Shadow AI is not a moral failure by employees, nor a gap to close with a memo. It is a signal that value arrived before control, and that the leader's job is no longer to prevent, but to govern: see what is happening, provide the right tool, and set the right boundaries. That is what follows.

David Alejano

Chief Strategy Officer · Thinkia · Madrid, June 2026

linkedin.com/in/dalejano

EXECUTIVE SUMMARY

Shadow AI is the use of artificial intelligence by employees outside IT supervision. It is not a future scenario: it is already happening in most organizations. Roughly three out of four knowledge workers use AI in their work, often with tools and data that no one has approved.

The thesis of this document is direct: the invisible risk is not that your employees use AI, but that they use it where you cannot see it, and banning it does not solve that. What is not governed is not controlled; and what is not controlled creates exposure.

These pages diagnose the phenomenon, inventory its real risks, explain why banning makes it worse, and propose a governance framework in three moves: make visible, govern, and enable. The goal is not to slow adoption, but to turn it from a silent liability into a governed capability.

75%

KNOWLEDGE WORKERS

already use AI at work, often without IT supervision.

3

GOVERNANCE MOVES

make visible, govern, and enable, in that order.

0

REQUIRED INSTALLATIONS

a free account is enough for data to leave the perimeter.

CONTENTS

In these pages.

Section 01

What Shadow AI is and why it is already *inside.*

It is not a threat on the horizon. It is a reality already operating inside your organization, outside the IT dashboard.

Shadow AI is the use of artificial intelligence tools by employees without approval, supervision, or visibility from technology and security teams. An analyst pasting a contract into a public chatbot to summarize it, a salesperson generating a proposal with customer data, a developer completing code with an unauthorized assistant: none of them is acting with bad intent, and all of them unknowingly expand the company's risk surface.

The phenomenon inherits the pattern of **shadow IT** - the applications teams adopted without going through IT - but accelerates it. Shadow IT required installation, configuration, sometimes payment. Shadow AI only requires opening a tab. That near-zero friction explains why it spreads faster than any policy can keep up, and why it leaves a much lighter trail.

> **Shadow IT required installing something. Shadow AI only requires *a tab*.**

The consequence is counterintuitive: the better the tool, the more it is used in the shadows. Immediate usefulness beats distant policy whenever an employee has a task in front of them and a deadline overhead. Denying that this happens does not stop it; it only guarantees that it happens without witnesses.

Section 02

Anatomy of the phenomenon: *why they do it.*

Understanding Shadow AI requires an uncomfortable concession: the employees who use it are often right in the moment. They solve a real task, save real time, and improve a real output. The problem is not their local judgment, but the sum of many local judgments without a common framework. The motivation is productivity, not rebellion.

Three forces feed the phenomenon. The first is the provisioning gap: the corporate tool, when it exists, arrives late, requires permissions, or performs worse than the public option, and the employee does not wait. The second is incentive asymmetry: people are measured on delivery, not on compliance with a policy they see as a brake. The third is cost invisibility: the risk of leaking data is diffuse and future, while the benefit of finishing sooner is concrete and present.

Read this way, Shadow AI stops being a discipline problem and becomes a design problem. Where the organization does not offer a good, governed, and fast path, the shadow fills the gap. Closing that gap - not punishing those who fall into it - is the only intervention that scales.

THE THREE FORCES

Why it grows

- **Provisioning gap.** The corporate option arrives late or performs worse than the public one.
- **Incentive asymmetry.** Delivery is rewarded, not policy compliance.
- **Invisible cost.** The benefit is present and concrete; the risk is future and diffuse.

None of these is fixed with a memo. All three are fixed with governance and good provisioning.

Adoption has already happened. It is not a decision pending at committee level, but a reality distributed across the workforce, almost always outside the technology and security dashboard.

75%

of knowledge workers already use AI at work, often without IT supervision.

Source pending independent verification.

SECTION 03 · THE RISK INVENTORY

What is at stake when AI operates outside the perimeter.

RISK	HOW IT APPEARS	IMPLICATION
Data leakage	*Sensitive information pasted into public tools that retain it or use it for training.*	Corporate data leaves the perimeter with no record and no path to recovery.
IP exposure	*Code, contracts, or strategy processed in third-party models.*	Loss of control over intellectual property and possible contractual breaches.
Non-compliance	*Processing of personal data without a legal basis or impact assessment.*	Regulatory risk under GDPR and the EU AI Act; penalties and reputational damage.
Security	*Scattered API keys, prompt injection, unsanitized outputs entering internal systems.*	A new attack surface that is hard to monitor and audit.
Hidden cost	*Individual subscriptions and token consumption without financial visibility.*	Fragmented spend, with no economies of scale or FinOps control.
Unaudited decision	*AI outputs influencing decisions with no traceability of origin.*	Inability to explain or defend a decision to an auditor or a customer.

Section 04

Why banning *makes* the problem worse.

A security leader's instinct when facing a new risk is legitimate: close the door. But Shadow AI does not have one door. It has a thousand, all the size of a browser tab, and none of them requires permission to open. A blanket ban does not eliminate usage; it pushes it into even less visible channels: personal devices, private accounts, conversations that no longer leave any trace on the corporate network.

The result is the worst possible combination: the same risk as before, now without observability. Prohibition turns a governable problem into a clandestine one. And it also imposes a silent cost: the organization gives up a productivity gain its competitors are not leaving on the table.

> **Banning does not eliminate use.**
> **It only eliminates the *witnesses*.**

This does not mean anything goes. Some uses must be forbidden and some data must never leave under any circumstances. The difference between a useful policy and a counterproductive one lies in its shape: one bans specific categories with an alternative close at hand; the other bans the entire tool and leaves the employee with no good way out. The first is followed. The second is evaded.

Section 05

The three-move governance *framework.*

Governing Shadow AI is not a choice between allowing and banning. It is a three-step sequence, and the order matters.

See

Make visible: see first, decide later

You cannot govern what you cannot see. The first move is an honest inventory of real usage: which tools, for which tasks, with which data. Without amnesty there is no reliable inventory; the goal is to understand, not to punish.

Gov

Govern: one control point

Control is built on visibility: a corporate gateway that applies authentication, data policies, quotas, and traceability before the request reaches a model. Control lives in the architecture, not in the user's goodwill.

Ena

Enable: provide the right tool

Governance holds when the official path is the easiest one. Offering a fast, capable, and secure corporate tool removes the reason to use the shadow path. Governed adoption is not imposed: it is made irresistible.

Section 06

The *CISO paradox.*

The security leader lives with a tension that seems to have no way out. If they slow things down, they kill the productivity leadership asks them to preserve and push usage underground. If they allow it without control, they accept a risk their function exists to contain. Either extreme leaves them exposed: that is the paradox.

The way out is not to choose an extreme, but to change the axis of the decision. The question is not "do I allow or ban it?" but "how do I make usage happen within a framework I can see and control?" Framed that way, security stops being a wall and becomes a channel: it does not stop the water, it directs it. The CISO who governs instead of blocking turns a risk liability into a trust asset, and stops being perceived as the department of "no."

That change of role has a political dividend as well as a technical one. A security team that enables with judgment gains the authority lost by one that only prohibits. And over time, that authority is what makes it possible to enforce the few bans that truly matter.

AXIS SHIFT

From wall to channel

The question stops being binary. Not "allow or ban," but how to make usage happen within a governed framework.

- The wall stops the water and makes it overflow elsewhere.
- The channel does not stop it: it directs it and measures it.
- Those who enable with judgment gain authority for the bans that matter.

What you do not govern, *you do not control*.

Section 07

How success is *measured.*

Governing without measuring means going blind again. Four indicators show whether the framework works.

01 **Share of governed usage**
What fraction of real AI usage goes through the corporate path versus the shadow path. This is the parent metric: if it rises, the rest improves with it.

02 **Incidents prevented and detected**
Blocked data-leak attempts, intercepted risky uses, rotated credentials. It measures the observability that prohibition destroyed.

03 **AI cost made visible**
Spend moving from scattered, invisible subscriptions to measured and budgeted consumption. What can be seen can be optimized.

04 **Corporate-tool adoption**
Real use of the official path. If the right tool is used, the shadow path empties on its own: proof that enablement worked.

TO REMEMBER

Four ideas worth taking away.

01

Shadow AI is already inside.

It is not a future risk: adoption arrived bottom-up before the top-down strategy.

02

It is not indiscipline, it is design.

Where there is no good, governed, and fast path, the shadow fills the gap. Close the gap; do not punish those who fall into it.

03

Banning worsens the position.

It keeps the risk and removes observability. Ban specific categories with alternatives, not the entire tool.

04

Make visible, govern, enable.

The order matters: see first, control in the architecture, and make the official path irresistible.

Conclusion

From silent liability to *governed capability.*

Shadow AI is, above all, good news badly managed. It means your people have already found value in artificial intelligence without being told to do so. The problem is not the appetite; it is that the appetite is being satisfied outside your view, with your data and without your limits.

That is why the thesis should be repeated plainly: **the invisible risk is not that your employees use AI, but that they use it where you cannot see it, and banning it does not solve that. What is not governed is not controlled; and what is not controlled creates exposure.**

The path is not to slow adoption, but to channel it: make visible what is already happening, govern it from the architecture, and enable an official path so good that the shadow has no reason to exist. Done this way, what is now a silent liability becomes a governed capability, and security stops being the department of "no" and becomes the one that makes it possible to say "yes, with judgment."

> **Mature security does not build walls. It builds *channels*.**

Thinkia

AN EXECUTIVE WHITEPAPER · 2026

SOURCES AND METHODOLOGY NOTE

Sources.

01 AI adoption data for knowledge workers (≈75 %, frequent use without IT supervision). Pending independent verification against the primary source before publication.

02 Reference regulatory framework: EU AI Act and General Data Protection Regulation (GDPR). Verify current versions and deadlines at the time of publication.

03 Shadow IT pattern as the conceptual precedent for Shadow AI. General technology-governance reference, with no attribution to a single author.

Note: the quantitative figures in this document must be independently verified against their primary source before publication. The regulatory frameworks cited may have changed; confirm the current versions.

HOW TO CITE THIS DOCUMENT

Alejano, D. (2026). *Shadow AI*. Thinkia thought leadership series. thinkia.com

thinkia

thinkia.com