The AI Governance Gap

What happens when the business moves faster than the framework

Views shared are personal and do not represent the views or positions of Howden. Practical examples are anonymised and drawn from general professional and networking experience.

INDEX

OPENING

Who am I, why this talk, and what I want you to leave with

THE PROBLEM

What governance can look like and why technical approval and privacy risk aren’t the same question

LESSONS FROM THE FIELD

Three lessons in designing governance & where it can get difficult

GOVERNANCE THEORY VS REALITY

What the data says governance is worth

CLOSING REMARKS

Embrace building the plane while flying it

BEFORE WE START

What I'd like you to leave with

1

Inventory is a control, not paperwork

You cannot assess what you cannot see.
Visibility: who owns what, what status it's in

This should come before any risk review.

2

A technical approval ≠ regulatory approval

There may be several gates as internal stakeholders try to move forward fast. Important to understand and communicate expectations and scopes of review as most likely they answer different questions.

Treating an IT sign-off as a governance pass is the single most common failure mode.

3

People bypass slow doors

If the safe route is harder to find than the unsafe one, the governance design may turn out to be the control failure.

4

Ownership and roles &responsibilities must be clear

Implementation depends less on an elegant framework than on clear decision rights, evidence, relationships and awareness: guardrails must be easy to read, to access and to digest.

5

Don't wait for perfect

Build the governance you can evidence now, say plainly what you still can't see, and record accepted risk explicitly.

Most governance talks may begin with a framework:

This one begins with a “scientific experiment”

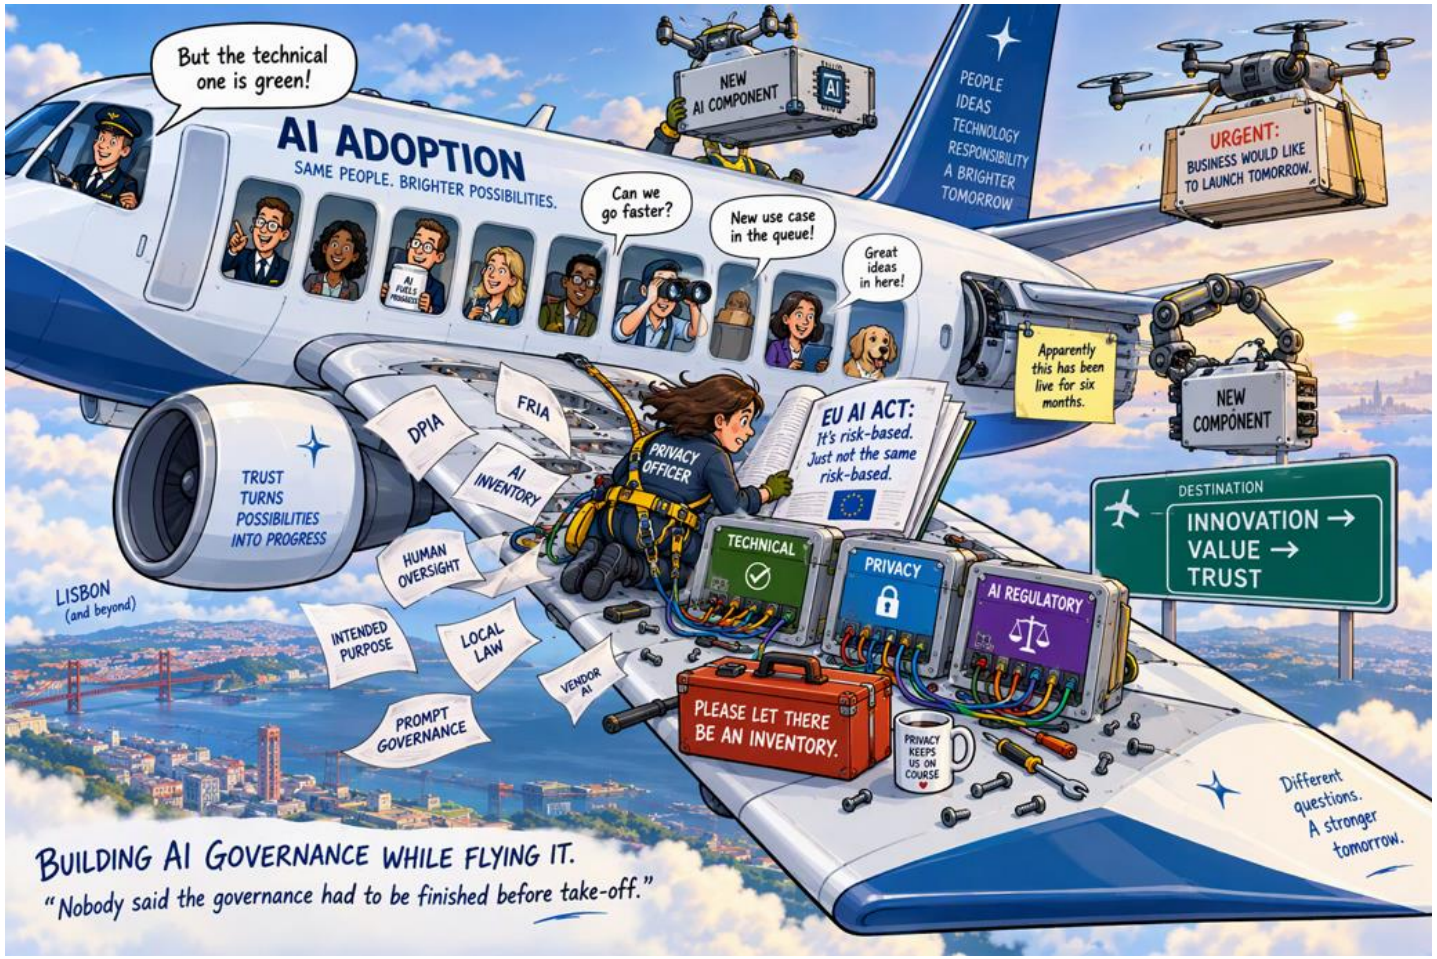

Prompt:

You're a Privacy professional. You've already got GDPR, DPIAs, incidents, vendors and local laws for 30+ countries in a multinational company in a heavily regulated sector. Now someone has given you AI governance. Add the AI Act, inventories, human oversight, prompts, three different types of risk, and obviously everything needs approving by tomorrow. All in motion at the same time.

Show me a what a completely normal working day looks like, designing a funny cartoon.

Most governance talks begin with a framework:

And also begins with a queue

What may be queuing?

- A growing list of use cases on your desk, with the words “already approved by…” added to the use cases on your desk so that your assessment is faster. “Approval” can mean different things.
- A prompt library moving toward launch.
- Local experiments already in production. Several functions using

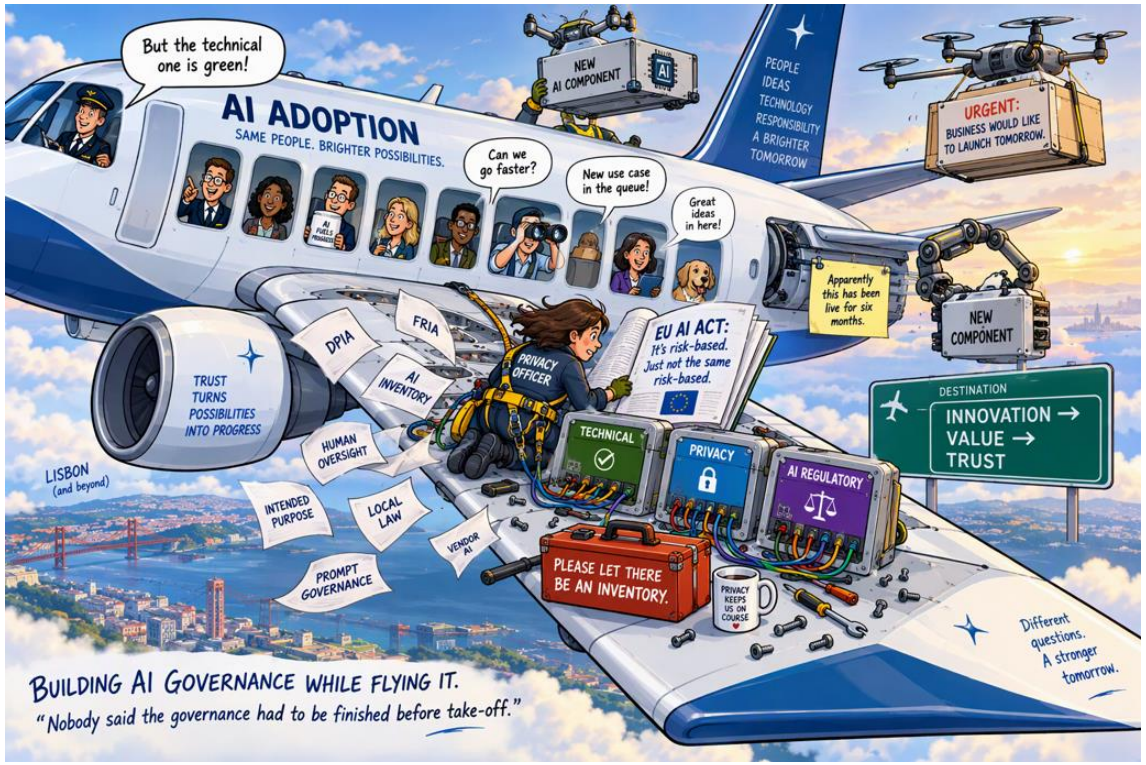

Privacy professionals are already governance professionals

WE ALREADY KNOW WHAT QUESTIONS TO ASK:

- **Purpose limitation**
- **Proportionality**
- **Data mapping**
- **Rights analysis**
- **Human oversight**
- **Auditability and influence**

WHAT DOES AI ADD? 3 COMPLICATIONS

The technology changes faster
than the review cycle built to assess it.

The same tool supports many purposes
a calendar summary and a recruitment score may share one tool.

Access is treated as permission
business teams often assume access to a tool licenses every use of it.



The job often changes before the title does

Fast Forward

Where I started

1 Why inventory comes first

Inventory is the first governance control.

2 Different questions

Technical approval and regulatory assessment answer different questions.

3 A faster safe route

Designing a route that is quicker than bypassing it.

4 What implementation depends on

Implementation depends less on the elegance of the framework than on clarity, communication and evidence.

“Don’t think Orwell; think Kafka.”

(Daniel Solove, Professor of Law, George Washington University, Chronicle of Higher Education, 2011)

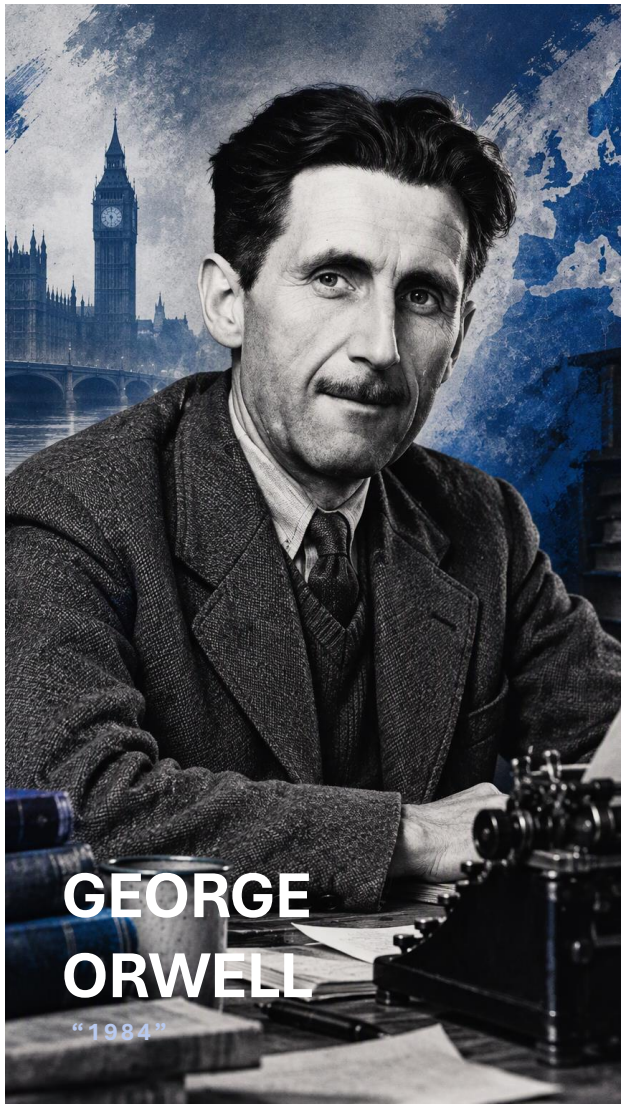



INVENTORY BEFORE ASSESSMENT

What “discovery” may turn out to mean

1 Shadow AI inside the official channel

Prompt batches reach review without prior legal, compliance or privacy sign-off and even on corporate tools, routed through an official-looking process, built in good faith.

2 The approved-tool fallacy

An enterprise licence improves security controls. It does not answer whether customer, claims, health or employee data should be used for that purpose. Tool approval is not use-case approval.

3 The supplier capability nobody had named

Automation descriptions hid real AI capability. Inventories must ask what is actually enabled and used.

4 The missing people in the room

A function can be named in the governance model and still be absent from the moments where priorities and expectations are set.

You cannot assess what you cannot see, and you cannot inventory AI by asking only which AI tools the organisation has bought.

THE TWO (or more) GATE-PROBLEM

The gates probably ask different questions

What gates can look like

GATE 1 · IT TECHNICAL / ARCHITECTURE

"How is this built?"

Environment, connectors, permissions, integrations, resilience, security, deployment route. An IT pattern can mean an easy build.

Example 1: a dev-only AI agent scores uploaded CVs from an approved folder: a routine build.

Example 2: a pilot compares recorded calls to stored voiceprints to identify callers: standalone, no downstream action.

GATE 2 · REGULATORY / PRIVACY

"What does it do, to whom?"

Purpose, data, affected people, autonomy, decision impact. The same "blue" build can touch HR health records or client risk profiling : an easy build can become high-risk.

Example 1: ranking applicants is GDPR profiling , and the EU AI Act flags CV ranking as high-risk.

Example 2: matching voiceprints is biometric identification: GDPR/DPIA risk, and a likely high-risk AI Act use.

Be careful with blanket approvals or fast tracks: an AI use case can pass a technical check and still be high-risk from a regulatory perspective.

WHERE IT CAN GET DIFFICULT

Mandate without full authority: everybody is an expert in AI

Expected to drive governance without roles and responsibilities clearly assigned and rolled out

Leadership impatience vs. incomplete evidence

Governance cannot classify “improve efficiency.” It needs data, people, impact and review detail.

Same words, different meanings

“Sandbox,” “pilot,” “approved,” “low risk” each may conceal different realities across stakeholders.

The inventory is always going to get out-of-date

Models, connectors, prompts and users change. Initial approval isn’t enough: reassessment triggers are.

Capacity consumed by preventable ambiguity

prompts reviewed line by line is sustainable only if low-risk items are routed lightly. Guidelines and guardrails need to be clear and available to users. Training and awareness and communication is key

Governance left out of the adoption story

When workshops celebrate speed but skip the assessment route, the business hears governance as a later hurdle.

LESSONS FROM THE FIELD

3 key lessons from governance designing

1

See and name it, before you assess it

Build the inventory: live tools, pilots and planned use cases. Policy comes with it or after it.

Separate the tool and technical build label from the use case purpose and data used assessment and give every submission a decision vocabulary: who can approve, when to reassess, what approval means and where it comes from. In multinational organisations liability lies locally, so should decision-making.

2

Design for speed, not just control: Building the plane while flying it

Give low-risk work a fast lane, but be clear about guardrails and evidence and reserve deep review for real risk. Set a minimum viable submission (problem, users, data, owner) so vague requests don't enter the queue.

3

Make the system defensible

Keep evidence inside the workflow, not in an inbox. Frame and communicate challenge as part of the design, not an obstacle to it. Position governance as leverage for revenue generated AI adoption.

And when leadership accepts a risk despite the gaps, document it: explicitly, by name.

GOVERNANCE THEORY VS OPERATIONAL REALITY

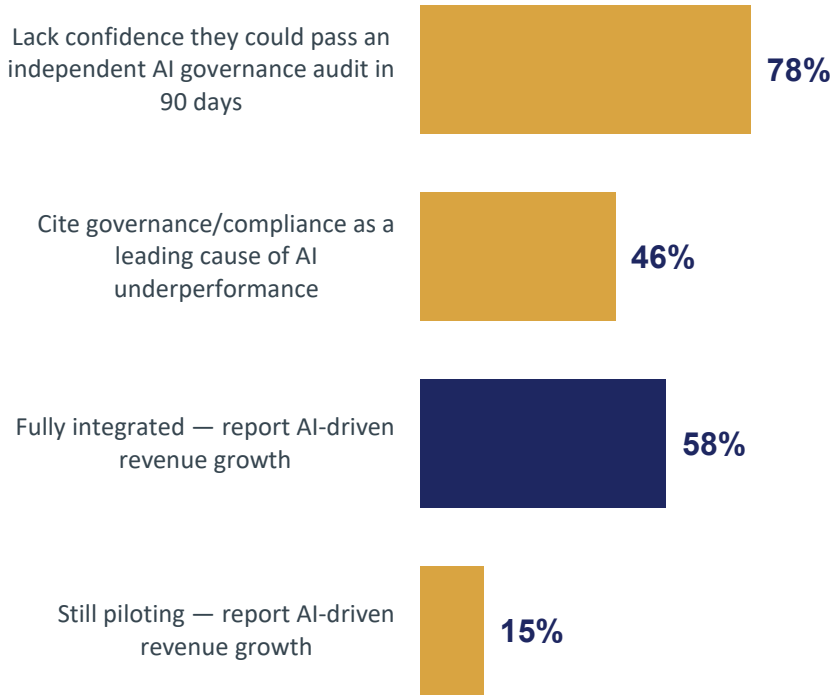

**Source: Grant Thornton 2026 AI Impact Survey of ~1,000 senior US business leaders.*

AI- driven revenue climbs with governance maturity. However, this survey also shows that yet, most senior leaders still doubt their own governance and compliance would hold up: 78% lack audit confidence.

This survey also points out that too often boards approve the AI investment without setting governance expectations. Leadership deploys AI without naming who owns the outcome. The organisation scales without building the infrastructure to prove any of it works.

The organisations closing this gap aren't waiting for the plane to land and provide time to design the governance, or for better technology, a regulatory mandate, or an incident to force the issue: they're building governance now, embracing building the plane while flying it why it pays to.

What does the "governance design" require?

Governance reality: a working model for a fundamental new way of working, ensuring you can prove which version, decision and condition applied to this use case today.

Thank you

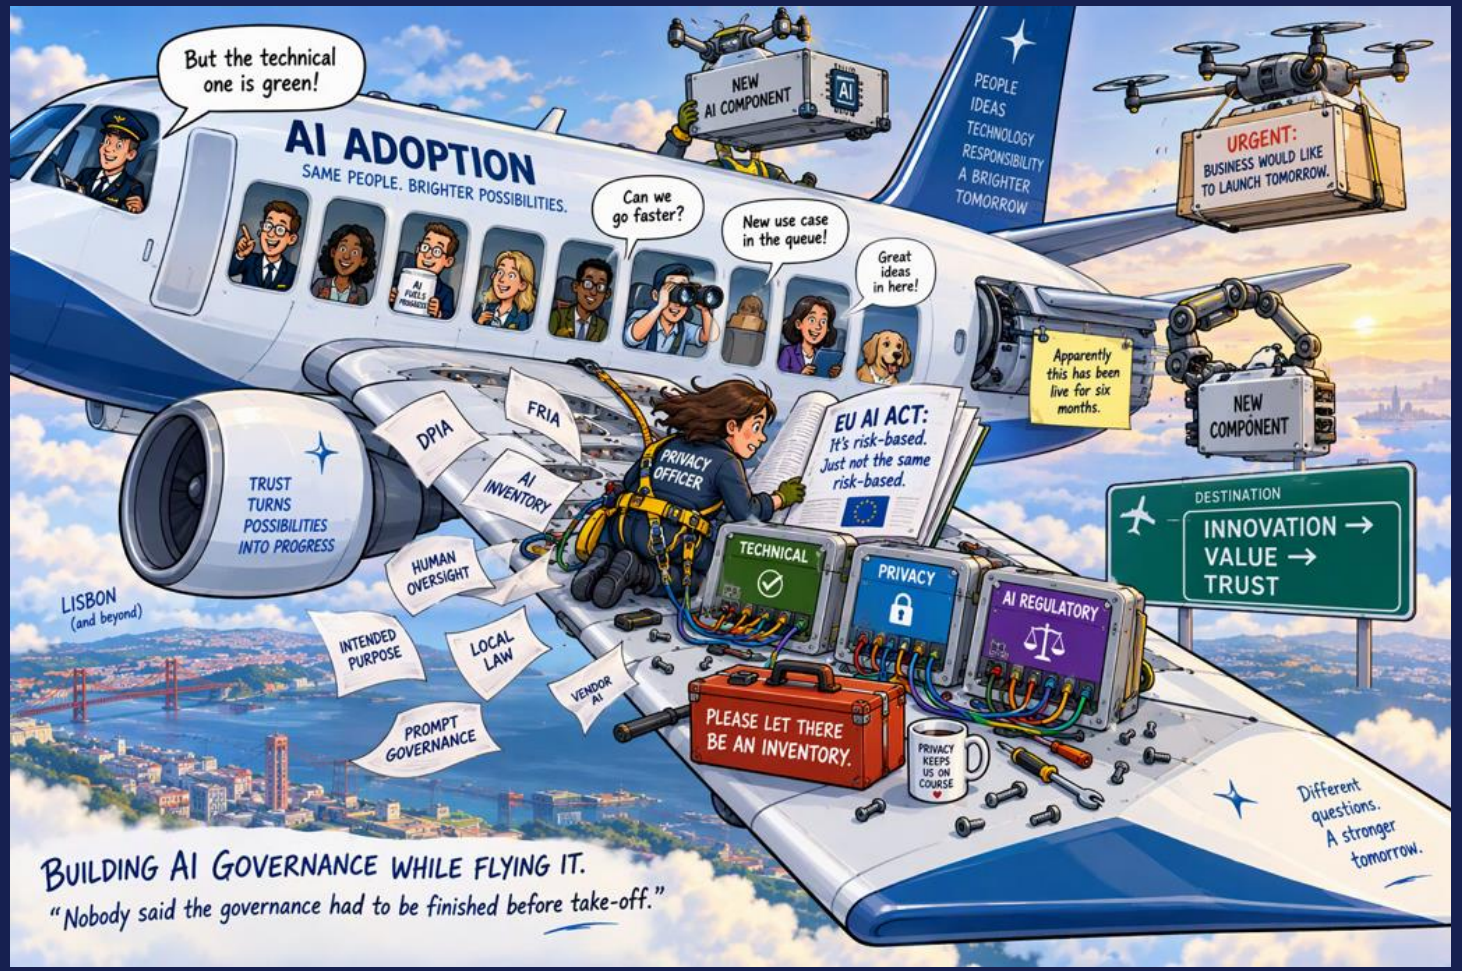

telma.candeias@gmail.com