# Mind the Gap:

# 88% Run AI. 8% Have a Program.

## Implementing AI governance that actually runs

EU AI Act · NIST AI RMF · ISO 42001 · The AI Adoption Lifecycle

**Anwarul Hasan (MBA, PhD)**

Founder & CEO · Risk Vizon LLC · Zurich, Switzerland

25+ years · UBS · KPMG · Swiss Re | anwarul.hasan@riskvizon.com

**75%**
wrote an AI policy —
only 36% a framework

**74%**
plan agentic AI —
just 21% can govern it

**2027**
EU high-risk deadline
(deferred, not gone)

**CA$812**
a chatbot cost
Air Canada in court

Sources: Pacific AI 2025 · Deloitte 2025 ·
Moffatt v. Air Canada (2024) · EU Digital Omnibus (2026)

**KEY TAKEAWAYS**

- A policy is a promise; a program is a control
- Classify every system before it ships
- Govern for change, not one deadline

# Mind the Gap: 88% Run AI. 8% Have a Program.

## Implementing AI governance that actually runs

**By Anwarul Hasan (MBA, PhD)**
Founder & CEO, Risk Vizon LLC · Zurich · July 2026

In 2024, a grieving traveller asked Air Canada's chatbot about bereavement fares. The bot told him he could claim the discount retroactively. That policy did not exist. When the airline refused the refund, a tribunal sided with the customer — and rejected Air Canada's remarkable defence that its chatbot was "a separate legal entity" responsible for its own words. The airline paid.

The ruling gave us the one-line summary of this entire era: **you own what your AI says — and does.** Now scale that from an CA$800 refund to a bank's credit decisions, an insurer's pricing or an HR platform's shortlisting. Then sit with this pairing from 2025–26 surveys: the vast majority of enterprises run AI in production, yet only a single-digit share call their governance mature. One survey found **75%** of organisations had an AI *policy* — but only **36%** had a working governance *framework*. That gap is where the risk lives.

## A Policy Is a Promise. A Program Is a Control.

Any board can approve an "AI Responsible Use Policy" before lunch. The morning after, the questions get harder. Can you answer these four?

- **Do you know every AI system you run** — including AI baked into software you bought, and the ChatGPT and Copilot tabs your staff already have open?
- **Who is the one accountable human** for each system?
- **Which systems can materially harm a customer**, and what stops them?
- **Could you prove all of the above to a regulator next week?**

The distance between a signed policy and a confident "yes" is the implementation problem. It is not solved by better words. It is solved by an operating program: an **inventory**, clear **accountability**, real **controls** and retained **evidence**.

## Why AI Breaks the Old Governance Model

For decades, model risk rested on a simple assumption: a human runs the model, and another human governs that human. AI dissolves that. The same technology now both **executes** a decision *and* is pitched as the **guard-rail** that checks the decision for compliance. When the thing doing the work also grades its own homework, independent oversight quietly evaporates — and today's AI cannot yet be trusted to know right from wrong, or what "compliant" even means in your context.

> *"When the executor governs itself, independent oversight is defeated. The whole job of a governance program is to put an accountable human back in charge of the machine."*

## The Five Things a Real Program Installs

### A front door nobody can walk around.

Every candidate system — built, bought or borrowed — hits a single **intake gate** and is *classified* before any build or deployment: its EU AI Act category (prohibited, high-risk, limited, minimal) and its internal materiality tier. This one control is what kills shadow AI.

### A lifecycle with gates, not vibes.

Systems move through eight phases — from intake and design to **independent validation**, deployment, monitoring and retirement — each needing a named approver before the next begins. A high-risk credit model gets validated by someone who did not build it; a minimal-risk tool gets a light touch.

### A risk appetite with numbers and a kill-switch.

“Be responsible” is not a control. “If accuracy drops more than 10% versus baseline, or the fairness ratio falls outside 0.8–1.25, suspend and escalate” — that is a control. Amber warns; red stops. A hard “no” list (per EU AI Act Article 5) rules out social scoring, manipulation, black-box decisions on material outcomes and deceptive deepfakes.

### Special handling for the sharp edges.

- **Generative AI** — an approved-tools list, no confidential data in unapproved tools, and human review of anything customer-facing.
- **Agentic AI** — the 2026 frontier: ~74% plan to deploy autonomous agents within two years, but only ~21% can govern them. Agents that *act* need scoped permissions, spend limits, human checkpoints for irreversible steps and full logging.
- **Third-party AI** — the vendor’s model is *your* risk. Contract for audit rights, liability and exit; monitor it like your own.

### Monitoring, people and proof.

Live dashboards across technical health, model performance and business outcomes; a four-rung escalation ladder from model owner to board; role-based training so controls survive contact with real users; and an audit trail that lets Internal Audit — and a supervisor — confirm the whole thing works.

## Inside the Program: An Eight-Gate Lifecycle

This is the engine the white paper details in full. Every system — built, bought or borrowed — travels the same eight phases, and cannot move to the next until an accountable human signs the gate. Governance scales with risk: a minimal-risk tool gets a light touch; a high-risk credit model gets every gate.

| # | Phase | Gate to pass before the next phase begins |
|---|---|---|
| 0 | Strategy & demand | Strategic fit and sponsorship confirmed |
| 1 | Intake & classification | Classified; permitted to proceed — the front door that kills shadow AI |
| 2 | Design & development | Design + data and bias plan signed off |
| 3 | Independent validation | Validated by someone who did not build it |
| 4 | Approval & deployment | Go-live approved; rollback tested |
| 5 | Monitoring | Within thresholds — continuously |
| 6 | Incident & change | Incident closed / material change re-approved |
| 7 | Revalidation & retire | Retain, revalidate or decommission cleanly |

© 2026 Risk Vizon LLC

## A Phased Rollout: 12 Months to Control

You don't boil the ocean — you sequence it, so each phase delivers a defensible control even if the next slips. The white paper's roadmap, in brief:

| Phase | What you do | What you can prove at the end |
|---|---|---|
| **Months 0–3 Foundations** | Stand up a lean AI Governance Office; sweep business lines, procurement and telemetry for an inventory; publish a GenAI acceptable-use rule now. | You know what you run — the intake gate is live. |
| **Months 3–6 Control the risk** | Classify everything; gate the 5–15% that are high-risk; add independent validation; screen for prohibited uses. | Every material system has an owner, a class and a status. |
| **Months 6–9 Operationalise** | Turn on monitoring and KRIs; wire incidents into op-risk; roll out training and vendor due diligence. | The first quarterly AI risk report to the board, from live data. |
| **Months 9–12+ Assure** | First internal audit; annual attestations; optional ISO/IEC 42001 certification; continuous regulatory monitoring. | You can prove control to a supervisor. |

## The Moving Target — and Why It Doesn't Matter

Here is the twist that proves the point. The framework was built expecting the EU AI Act's high-risk obligations to bite on 2 August 2026. As of mid-2026 that date has *moved* — the EU's **Digital Omnibus** pushed most high-risk deadlines out to **December 2027**, while keeping the "tell users they're talking to AI" transparency rule on for August 2026. Deadlines will keep shifting. If your governance is pinned to one date, every change is a fire drill. If it is built around a **classified inventory and a repeatable lifecycle**, a regulatory change is just a reason to re-run classification. That is the whole game.

The institutions that win with AI won't be the most cautious or the most reckless. They'll be the ones that move **safely at speed** — approving good use cases fast because the guard-rails are trustworthy, and catching bad ones early because someone is actually watching. That never came from a policy document. It comes from an implemented program.

> *"The tribunal was right: you own what your AI says. Build the program that lets you stand behind it."*

I would welcome your perspective — connect on LinkedIn or reach out at anwarul.hasan@riskvizon.com.

**Read the complete White Paper:**

***"Implementing an AI Governance Program in a Medium-to-Large Institution"***

Available at: **www.riskvizon.com**

#AIGovernance #EUAIAct #RiskManagement #ModelRisk #ResponsibleAI #BoardGovernance #CRO #AgenticAI #ISO42001 #RiskVizon

**ABOUT THE AUTHOR**

**Anwarul Hasan (MBA, PhD)** is the Founder and CEO of Risk Vizon LLC, Zurich. With 25+ years at UBS, KPMG and Swiss Re, he is the architect of **GRIP — the Governance & Risk Intelligence Platform**, designed to empower boards, regulators and executives to govern, monitor and steer enterprise risks — including AI risk — in real time.

anwarul.hasan@riskvizon.com · www.riskvizon.com | © 2026 Risk Vizon LLC