Cocoon Technology

WHITE PAPER · AI GOVERNANCE

The Pilot That Never Went Live

Why AI pilots stall at the governance gate, and why the way out starts with how your work actually gets done.

cocoon.technology

November 2026

> MIT research published in 2025 found that when generative AI pilots stall, the quality of the model is rarely the cause. Many stop at the meeting where somebody asks who owns the process the AI will sit in, and nobody can answer. Governance gets the blame, and the gap it exposes is older: the organisation never wrote down how the work is done, so there is nothing for governance to attach to.

The demo worked

MIT's NANDA initiative looked at 300 public AI deployments, interviewed 150 leaders and surveyed 350 employees for its 2025 report, The GenAI Divide. It found that around 95% of generative AI pilots produced no measurable effect on profit and loss. The authors put that down to a learning gap in the organisations using the tools, and said the quality of the models was rarely the cause.

Gartner had already predicted, in July 2024, that about 30% of generative AI projects would be abandoned after proof of concept by the end of 2025. It blamed poor data quality, inadequate risk controls, escalating costs and unclear business value. In February 2025 it added that 63% of organisations lacked confidence in their data management for AI. It expected 60% of AI projects without AI-ready data to be abandoned through 2026.

Every one of those reasons is a question the business could not answer about itself.

We watched a pilot stall this way at a client that had built a generative AI assistant to help service agents answer customer queries. Its answers were accurate. Response times fell and the agents liked it. When the business tried to move it into production, risk, compliance, legal and IT each arrived with their own questions: what customer data it used, where that data went, which decisions it could make, who signed off an answer and who owned the risk. None of the answers had been designed into the pilot, and it stalled.

The questions that stop a pilot

The go-live gate is where a pilot meets the people responsible for risk, data, finance and operations. Each of them arrives with a short list of questions, and most of those questions are about the process the AI will change.

Who owns this process today, by name? Which steps will the AI carry out, and which will a person still sign off? What data does it read, and who is allowed to see it? What does the process cost now, in hours and errors, so we can tell whether the pilot made it better? What happens on the day the AI gets something wrong, and who finds out first?

In a typical case, a team that built a pilot in six weeks can answer these in an afternoon if the process is documented and owned. Where the process lives in the heads of three people and a spreadsheet called Tracker FINAL v4, the answers take months, and the pilot waits while somebody goes looking for them.

That is the stall. The governance board is asking reasonable questions, and the organisation is discovering it never knew the answers.

You cannot govern what you have never mapped

There is plenty of AI governance material. One infographic widely shared on LinkedIn sets out eight layers of enterprise AI governance, from principles and regulation through to runtime guardrails and audit tooling. All of it is useful, and some of it will be necessary for any organisation running AI in production.

Every layer assumes something most organisations lack: an accurate picture of where AI sits in the work, who owns that work and what it touches. The NIST AI Risk Management Framework, published by the US standards body, organises its guidance into four functions, Govern, Map, Measure and Manage. The second of those is mapping the context the AI runs in. ISO/IEC 42001, the international standard for an AI management system, asks for the same thing in its own language. A risk register for an AI use case is only as good as the description of the process underneath it.

A structural engineer will not sign off an extension without drawings of the wall it joins. AI governance is in the same position. The controls and the human checkpoints hang from the process, and where the process has never been drawn, there is nowhere to fix them.

Start with how the work actually runs

The fix costs less than the stall.

Pick the process the pilot is meant to change and map it as it runs on a busy Monday. Include the workarounds: the email that triggers the next step, the spreadsheet somebody updates by hand. The documented procedure from the last audit is a useful starting point and rarely the whole story.

Then mark four things on that map. Where the AI would act. Where a person must still make or check a decision. What data moves at each step and from which system. Who owns each step, by name. Add the baseline numbers, so everyone agrees what better looks like before the pilot is judged against it.

That map answers most of the go-live questions before anyone asks them. Risk gets a place to put controls, data owners get a list of what the AI reads, and finance gets a before figure. Because the map lives in a process repository, a single shared library of how the business works, it stays true after the pilot team has moved on.

That is how the service agents' assistant got through. We modelled the whole process around it: where data entered, which systems it touched, where a person had to make the call, which controls existed and who owned each decision. We then laid the risks, controls, roles and evidence over the model. The assistant went into production as a governed process, with named owners and evidence that the controls work.

Cocoon Technology builds these maps in BusinessOptix, the process modelling platform we partner with, and connects them to the systems that hold the data. We have run seven AI governance engagements in the last 18 months. Our other projects met the same questions at sign-off: we have connected the systems behind an iconic British motorcycle brand, automated supplier payments across a national pub estate, rebuilt patient intake for a private hospital operator and untangled know-your-customer and anti-money-laundering compliance across several jurisdictions. Only the thing waiting at the gate has changed.

The regulation clock

UK GDPR already applies to any AI that processes personal data. The Information Commissioner's Office expects a data protection impact assessment, or DPIA, wherever that processing is high risk.

Firms that sell into the EU, or whose AI outputs are used there, also fall under the EU AI Act. The AI Omnibus, Regulation (EU) 2026/1744, came into force on 27 July 2026 and pushed back the main high-risk obligations. Stand-alone high-risk systems, including AI used in recruitment, credit decisions and access to essential services, now face their obligations from 2 December 2027. AI built into regulated products follows on 2 August 2028. The duty on organisations to support AI literacy among their staff has applied since February 2025. The transparency rules for AI that interacts with people applied from August 2026.

The delay gives organisations time to do the mapping. Every one of those obligations still assumes you know which of your processes use AI, what the AI decides and who is accountable for it.

Twelve questions to answer before your next pilot

These are the questions the go-live gate will ask, grouped the way a governance board asks them. If you can answer all twelve for the process your pilot will change, governance will be a short meeting.

Process

- [ ] 1 Is there a current map of the process as it actually runs, including the workarounds?
- [ ] 2 Do you know how long it takes and how often it goes wrong today, in numbers?

Ownership

- [ ] 3 Does the process have a named owner who will be accountable for the AI's part in it?
- [ ] 4 Has someone agreed which decisions the AI may make alone and which a person must approve?

Data

- [ ] 5 Can you list every data source the AI will read, and which system holds the master copy?
- [ ] 6 Do you know whether any of that data is personal, and has a DPIA been considered?

Risk and controls

- [ ] 7 Is this use case on a register of AI use cases, with a risk rating?
- [ ] 8 Are the controls tied to specific steps in the process, so an auditor could find them?

Regulation and policy

- [ ] 9 Do you have an AI policy staff have read, and have the people using this tool been trained on it?
- [ ] 10 Have you checked whether the use case falls under the EU AI Act high-risk categories?

Measurement

- [ ] 11 Were the success measures agreed before the pilot started?
- [ ] 12 Will someone be told, and know what to do, when the AI gets something wrong?

About Cocoon

Cocoon Technology is a UK business improvement consultancy and BusinessOptix partner. We map how organisations actually work, agree what better looks like in numbers, and then decide what technology earns its place. AI is the newest technology in that queue. It still goes last.

Score your pilot in five minutes: cocoon.technology/ai-governance-assessment

Talk it through: calendly.com/richard-tisdall-e-cocoon/ai-pilot-review

Tell us where your pilot is stuck. We will tell you honestly whether we can help.

Sources

1. MIT NANDA, The GenAI Divide: State of AI in Business 2025, as reported by Fortune, 18 August 2025. https://fortune.com/2025/08/18/mit-report-95-percent-generative-ai-pilots-at-companies-failing-cfo/
2. Gartner prediction on generative AI projects abandoned after proof of concept, as reported by Computerworld, 29 July 2024. https://www.computerworld.com/article/3478532/nearly-one-in-three-genai-projects-will-be-scrapped.html
3. Gartner survey on AI-ready data, as reported by Freevacy, 26 February 2025. https://www.freevacy.com/news/gartner/gartner-reveals-60-of-ai-projects-with-data-issues-will-fail/6155
4. Gibson Dunn, EU AI Act Omnibus Agreement: Postponed High-Risk Deadlines and Other Key Changes. https://www.gibsondunn.com/eu-ai-act-omnibus-agreement
5. Mondaq, EU Digital Omnibus on AI enters into force. https://www.mondaq.com/eu-digital-omnibus-on-ai-enters-into-force/1825722
6. NIST AI Risk Management Framework (AI RMF 1.0). https://www.nist.gov/itl/ai-risk-management-framework
7. ISO/IEC 42001, Artificial intelligence management system. https://www.iso.org/standard/42001