




■ ARTICLE · OCT 2026 · 5 MIN READ

Five Forces of Enterprise AI

AI programmes rarely stall because the technology doesn't work. They stall because they're pulled in several directions at once, by people who are each right about their own piece.

Daniel Mackle · CognaOne · cognaone.com



Four pressures, one decision.

Over the past few years, working inside and alongside regulated organisations, we've noticed a pattern that keeps repeating. AI programmes rarely stall because the technology doesn't work. They stall because they're pulled in several directions at once, by people who are each right about their own piece.

We call them the five forces of enterprise AI. Four of them pull, and one of them decides.

The four pulls

Data. Every AI conversation eventually becomes a data conversation. Teams are refining, cleaning, unifying and governing data so models and agents can use it. What has changed is the stakes. When AI only produced reports, poor data meant a poor report. When agents act, poor data can mean a poor action. Gartner expects organisations to abandon a large share of AI projects that are not supported by AI-ready data [3].

Agentic AI. New platforms, copilots and agents are arriving through many doors at once: central technology teams, federated business teams, vendors adding agents to software already in use, and staff simply bringing their own tools. One large survey found most AI users at work were bringing their own AI tools, and fewer than half had received training from their employer [4]. Knowing which AI systems exist is essential, but no longer enough. We also need to know what each one is allowed to do, and under whose authority.

Regulation and risk. Regulators increasingly expect active oversight, not a policy on a shelf: evidence that controls work, reporting when they don't, and clear lines of accountability. In Australian financial services, the Financial Accountability Regime (FAR) places responsibility with named executives across banking, insurance and superannuation [5], and APRA's operational risk standard, CPS 230, makes boards accountable for operational risk, including risk arising from material technology providers [6]. AI may take the action, but the organisation still owns the outcome.

People. This is the force we think is easiest to overlook. AI succeeds or fails on whether people trust it, understand it and know who owns its decisions. The University of Melbourne and KPMG found only around a third of Australians are willing to trust AI, among the lowest of the countries surveyed [8]. That doesn't read to us as resistance to technology. It reads as a reasonable request for assurance.

The one that decides

The fifth force is commercial. Boards and executives ask one question of each of the other four: is this returning on the investment, at a level of risk we can stand behind?

That question is getting harder to answer with confidence. An MIT study found most organisations it looked at had yet to see a measurable return from generative AI [1]. Gartner expects a large share of agentic AI projects to be cancelled by the end of 2027, and names escalating costs, unclear business value and inadequate risk controls as the reasons [2].


We see organisations caught between two poor outcomes: moving fast without enough control, or stalling good projects because they can't show the risk is managed. Either way, the question has shifted from "can we deploy AI?" to "which systems do we authorise, under what conditions, and can we show the decision was sound?"

Nor is the pressure easing. Australia's National AI Plan chose existing laws over new AI rules [7], which leaves organisations to show control under the regimes they already answer to.

Where the forces meet

These forces aren't separate problems. Agents make data quality an action problem. Regulators ask about data and AI together. Accountability sits with people, not systems. And all of it lands on the same budget, in front of the same board.

That is the gap we're working on with CognaOne: a governance control layer between these pressures and the decision to scale, with policy on each agent action, human oversight where it matters, and evidence a regulator can follow. It won't fix every data or workforce problem. It gives the decision to invest, scale or stop a basis that can be shown and defended, and most Australians say real assurances like these would make them more willing to trust AI [8].



Governance is how AI reaches production

We see governance less as a brake on AI and more as the thing that lets it reach production safely. We'd be glad to hear whether this matches what you're seeing.

Daniel, CognaOne.com Founder

AI use: research and drafting for this article were supported by AI tools. Multiple models were used for research and to cross-check facts, and every factual statement was verified by the author against the cited sources. The author is responsible for the content.