

Thirty leaders. Ten days. Five investable pilots — and one finding nobody expected.

An applied AI programme at an industrial holding, and what it revealed about why organisations stall after training.

Delivered	Format	Participants	Design & delivery
August 2026	20 sessions · 10 days	30 executives, 5 teams	Dr. Mitt Nowshade Kabir

The engagement

In August 2026, CCAIE delivered a ten-day applied AI programme to thirty executives across an industrial holding and three of its subsidiaries — general directors, finance and economics leads, production managers, and IT.

The brief was not to teach a technology. It was to test a hypothesis: **can process owners who are not programmers take an idea to the point where a board can decide whether to fund it?**

The design made evasion impossible. Each of five cross-functional teams had to name a specific step in a specific process, cost it from their own data, define the condition under which the project would be stopped, and name the person answerable for it.

What was produced

A 67-page portfolio: five processes documented step by step, each with economics, a risk matrix, safe-use rules, and a 90-day roadmap.

Five pilots reached investment-decision quality. The two strongest showed payback of **11.2 and 5.9 months** against an industry benchmark of 2.4 years across 95 comparable projects at a larger competitor — figures as published by the company concerned, not independently audited. Together the portfolio identified **160 hours** of recoverable effort per month.

The measurement caveat — read this before the numbers above. Not one of the five teams measured the actual duration of their own process step. They estimated it from process description. The arithmetic in every business case is sound; the input quantity is unverified. Until measurement is done, every payback figure in this portfolio is an illustration rather than a result.

The part that is harder to value and impossible to buy. Five of the holding's processes are now described as they actually run rather than as the regulations say they run. That description did not exist before August, and no supplier sells it.

The finding

Here is what we published to the client's owners, and what most consultancies would have buried.

Not one of the five teams measured the actual duration of their own process step. One team reported eight hours where the number turned out to describe how long a document takes to travel, not how long a person works on it.

This is not a criticism of the participants. The holding has no practice of operational measurement: process durations are known by impression, not by protocol. And the implication runs well past AI. Where a process baseline is not measured before work begins, the effect of *any* project becomes impossible to confirm or refute, and every argument about results collapses into an argument about opinions.

The larger conclusion

The programme confirmed the hypothesis it was built to test — process owners can design fundable AI projects — and then ran into a wall that had nothing to do with what anyone knew. Four things were missing, and none of them is a lesson:

- 1 **No approved environment.** No corporate contour in which an employee can assemble an assistant through a standard procedure. Without one, every initiative becomes an IT ticket and the queue stops at around the twentieth.
- 2 **No portfolio mechanism.** No intake, no selection criteria, no funding route, and — most missing of all — no procedure for stopping a project. Five pilots can be run by hand. Fifty cannot.
- 3 **No data competence.** Not programmers: someone able to judge whether accounting-system data is fit for an assistant to work with, and prove it by measurement.
- 4 **No measurement practice.** The baseline problem above, generalised.

Training does not close this gap. We wrote that sentence in the report to the owners, about our own programme, because it is true: a programme of this design takes projects to the point of an investment decision and by construction cannot take them to a working system. Saying so is not a caveat. It is the reason the next engagement is a different engagement.

What we recommended

Three options were set out — build internally, install the capability with an external team, or appoint an internal portfolio owner immediately and buy external expertise for three bounded tasks. We recommended the third, and stated our own interest in the recommendation in the report itself.

The reasoning: a pilot without an internal owner does not survive, whatever the quality of external support. The programme demonstrated this directly — the teams that advanced furthest were those where the process owner was named and attended every session.

What this engagement is evidence of

- Executives without technical backgrounds **can** produce investment-grade AI project documentation in ten days, given a structure that forces specificity.

- Cohort composition determines the outcome more than curriculum does. Two IT participants were provided against five recommended; three of five teams therefore had no technical member.
- The constraint on industrial AI adoption is organisational infrastructure, not knowledge — and it is diagnosable in ten days.
- **We publish the limits of what we measured.** In a market where efficiency claims are routinely quoted without a source, that is the point, not a weakness.

Programme design and delivery: Dr. Mitt Nowshade Kabir. The client is anonymised. Peer-company comparison figures are as published by the companies concerned and are not independently audited.

Canadian Center for AI Entrepreneurship · King, Ontario, Canada · info@ccaie.ca · <https://ccaie.ca/case-studies/executive-programme>