

Start with the room, not the pilot

Why the cheapest useful first move in AI adoption is a shared understanding and a shortlist, not a proof of concept

A C-Insight whitepaper | August 2026 | Written for boards and executive teams of medium and large organisations

The organisations that struggle with AI are rarely the ones that chose badly between technologies. They are the ones that started building before anyone in the room could describe, in plain language, what the work actually is. This paper sets out why that happens, what the evidence says about it, and why a structured day or two spent on understanding and selection is the highest-return money in the whole adoption sequence.

1. Where Australian organisations actually are

Adoption has moved quickly on paper. The Australian Bureau of Statistics reported in its Characteristics of Australian Business release for 2024-25, published on 25/06/2026 from a survey of nearly 7,000 businesses, that just over one in ten businesses (12 per cent) reported using artificial intelligence, against 1 per cent in 2022-23. Among large businesses of 200 or more employees the figure was around 35 per cent, up from 9 per cent in 2021-22. Medium businesses reached 22 per cent, against 3 per cent previously.

Those numbers describe a starting line, and it is important to be precise about what they do not describe. The ABS states plainly that the question "is not designed to measure intensity or extent of use within the business". It counts organisations that have begun. It does not count organisations where AI is working, at scale, inside a process that matters. There is no authoritative Australian statistic for that second thing, which is worth pausing on, because it is the number every board actually wants.

The closest available proxy is global. McKinsey's State of AI survey, published 05/11/2025 from 1,993 respondents across 105 countries, found that nearly two thirds of respondents said their organisations had not yet begun scaling AI across the enterprise, and that just 39 per cent reported any EBIT impact at the enterprise level, with most of those attributing under 5 per cent of EBIT to it.

Put the two together and the shape of the problem is clear enough. Starting is now close to universal among larger organisations. Landing is not.

2. The failure is not where people look for it

The instinct, when an initiative stalls, is to look at the technology. The best available evidence says look somewhere else.

RAND Corporation published The Root Causes of Failure for Artificial Intelligence Projects and How They Can Succeed on 13/08/2024, drawn from 65 semi-structured interviews with experienced practitioners conducted between August and December 2023. Its central finding is unambiguous. Failures driven by the decisions and expectations of business leadership were, in RAND's words, "far and away the most frequent causes of project failure", and 84 per cent of interviewees cited one or more of those leadership-driven causes as the primary reason AI projects fail.

RAND names five root causes: leadership misunderstanding or miscommunicating what problem is to be solved; insufficient quality data; a focus on the technology rather than the business problem; inadequate infrastructure for managing data and deploying models; and applying AI to problems that are beyond current technological capability. Three of those five are decisions. One is a judgment about scope. Only one is an engineering constraint in the ordinary sense.

One correction is worth making here, because it is a claim this paper deliberately does not lean on. The widely repeated line that "more than 80 per cent of AI projects fail" is not a RAND finding. RAND attributes it to "some estimates" made elsewhere and does not itself test it. A reader who checks will find that out, and an argument that depends on it deserves to lose.

Two further findings point the same direction. McKinsey found that organisations getting the most from AI were three times more likely to have senior leaders visibly owning adoption, and nearly three times as likely to have fundamentally redesigned the workflows AI sits inside. Boston Consulting Group, in From Potential to Profit published on 14/01/2025 from a survey of more than 1,800 C-suite executives, reported that leading companies prioritise 3.5 use cases on average against 6.1 for everyone else, and anticipate 2.1 times greater return. The same paper sets out a 10-20-70 allocation observed among top performers: 10 per cent of effort on algorithms, 20 per cent on data and technology, and 70 per cent on people, process and cultural change.

The consistent finding across four independent bodies of evidence is that the determining factors sit upstream of the technology, in problem definition, process, sponsorship and selection discipline. Those are all decisions an organisation makes about itself, before it buys anything.

3. What this looks like from inside the organisation

I spent 25 years as a C-level technology operator before founding C-Insight, including CIO roles at BHP Billiton, DHL EMEA and Maersk's Damco, and further executive roles at Rexel Australia and Sonepar. I led transformation programs that worked and several that did not. The patterns below are not drawn from AI. They are drawn from ERP, from supply chain replatforming, from analytics, and from a decade of automation initiatives. They are recognisable in AI work now because they were never technology problems in the first place.

The problem gets named by the loudest person, not by the person who does the work

A candidate initiative arrives already described as a solution. By the time it reaches a steering committee it has a name, a vendor category and a sponsor, and nobody has gone back to ask what the work actually is, who performs it, how often, and what they look at while performing it. The description of the solution has quietly replaced the description of the problem, and every subsequent decision inherits that substitution.

The business case rewards the saving you can already cost

Efficiency and capability are not the same thing. Efficiency does existing work more cheaply, the same invoices with fewer people. Capability makes possible what was not, every contract reviewed rather than a sample. A conventional business case can price the first and struggles to price the second, so efficiency wins by default. Most portfolios end up almost entirely efficiency without anyone having decided that on purpose, and the larger prize is never put in front of the people who could have funded it.

Automating a poor process gives you a faster poor process

Most operational processes carry two or three steps that exist only because a system cannot do something, and the people who perform them stopped noticing years ago. Automate around those steps and you have paid to make a workaround permanent. Most of the available value sits in the redesign rather than in the technology, and the redesign is the part nobody sells you.

Nobody tells the board about the dip

Productivity commonly falls before it rises, because the process is being reworked while the old job still has to be done. The Reserve Bank of Australia, in its Bulletin article of 13/11/2025 drawn from interviews with 105 medium and large firms, noted early evidence that AI adoption among older established manufacturing firms may follow a J-curve, with short-term productivity losses preceding longer-term gains as organisations and production processes adjust. The same firms expected AI effects to materialise over three to five years rather than one to three. The RBA states this cautiously and so should anyone citing it, but the operational point stands on its own: the dip is normal, and the sponsorship you need is not for the launch, it is for the dip. Say that in advance and you keep support at the moment it matters. Say it afterwards and it sounds like an excuse.

One flagship project carries all the risk

A single large initiative concentrates both the risk and the visibility, which makes it politically difficult to stop and therefore difficult to learn from. Several smaller pieces at the same total cost, a few near-certain and a few deliberate bets with an agreed point to stop, produce something more valuable than any one of them: the organisation learns how to specify this kind of work, how to check it, and who needs to be in the room.

4. Why the first move should be a room rather than a build

Read those patterns together and the case for starting with education and selection is not really a training argument. It is an argument about sequence.

The decisions that determine whether AI adoption works are made before any technology is selected, and in most organisations they are currently being made by a group of people who do not share a working vocabulary. Executives have been asked for a position they are not yet equipped to hold. Teams are already using tools nobody approved. Vendor proposals arrive that nobody in the room can compare on their merits, because each one describes itself in its own language. In that state, a pilot does not test whether AI works. It tests whether one supplier's framing survives contact with one process, which is a much smaller and less useful question.

A structured education and selection engagement does four things before a dollar of build is committed.

- It gives a leadership group and the people who would own the work day to day a common description of what these systems do, expressed as the work performed rather than as product categories. That single move is what lets two proposals be compared.
- It surfaces candidates from the organisation's own processes, described by the people who perform them, rather than from a supplier's reference list.
- It produces a list of what is not worth pursuing, which is usually the more valuable half of the output and is almost never produced by anyone with something to sell.
- It leaves the organisation able to evaluate the next proposal without help, which is a capability rather than a deliverable.

The BCG finding on use case discipline is the closest thing to direct support: fewer, better-selected candidates are associated with materially higher expected return. It is worth being honest about the limits of that support. It is a cross-sectional association rather than a controlled comparison, and the return figure is anticipated rather than realised. No study I could find experimentally tests education first against pilot first. The argument here is a reasoned inference from converging evidence and from operating experience, and it is offered as that rather than as proof.

5. What this engagement is, and what it deliberately is not

C-Insight's AI Bootcamp and Use Case Workshop exists to occupy exactly this position in the sequence. It runs as eight modules covering what AI does described as work, where value has and has not been realised, process redesign, data quality as an operational question, governance under Australian obligations, workforce change, and how to evaluate a proposal. Every module ends in a diagnostic question answered against the participant's own operation. The workshop then converts that understanding into a shortlist of candidate use cases drawn from the organisation's own processes, with a named owner and an open question against each, and a documented list of what was removed and why.

Two rules govern the delivery. No product, platform or supplier is named at any point, including in worked examples. And a question the room cannot answer is treated as a finding rather than a gap to be talked past, because it usually is one.

The boundary matters as much as the content, so it is stated plainly rather than left for a client to discover afterwards.

Question	Answered here?	Where it is answered
What is AI, in terms of the work it performs	Yes	Education component
Where might it apply in our processes	Yes, as a shortlist of candidates	Use case workshop
Which candidates are clearly not worth pursuing	Yes	The not worth doing list
What is each candidate worth in dollars	No	AI Opportunity Diagnostic
In what order, with what dependencies and cost	No	AI Opportunity Diagnostic
Can this organisation actually execute any of it	No	AI Readiness Assessment
Which platform or supplier should we use	No, and no C-Insight product answers this	A separate and later decision, made by you

It is not a compressed opportunity diagnostic and it is not priced as one. Formats run from a half-day executive briefing to a full two days, from \$5,500 to \$18,500 excluding GST, for up to sixteen participants.

6. The test of a good first engagement

There is one test worth applying to any first AI engagement, whoever provides it. It must be capable of concluding that nothing here is worth doing yet.

An engagement that cannot reach that conclusion is not an assessment, it is a preamble to a sale, and the room will work that out inside an hour. The value of an honest first pass is not that it finds opportunities. It is that the shortlist it produces can be defended to people who were not in the room, including the ones holding the budget, precisely because the method was published before the work started and the same method was capable of returning nothing.

If your organisation cannot currently compare two AI proposals on their merits, that is the problem to solve first. It is cheaper than the pilot you were about to fund, and it is the only thing that makes the pilot meaningful.

Where to from here

If any of this is recognisable, a short conversation is usually enough to work out whether the first move is education, an honest look at whether the organisation could absorb AI at all, or nothing at all for now.

C-Insight offers a free 30-minute discovery call. There is a booking link at c-insight.com.au, or write to frits.devroet@c-insight.com.au. Every enquiry is answered personally.

About the author. Dr Frits de Vroet leads C-Insight, an independent, vendor-neutral AI adoption advisory based in Perth, Western Australia, working with medium and large organisations across Australia and South-East Asia. He spent 25 years as a C-level technology operator, including CIO roles at BHP Billiton, DHL EMEA and Maersk's Damco, and holds a PhD in strategic management. C-Insight names, recommends and scores no AI product or platform in any deliverable, and every engagement is fixed fee and delivered by the person who scopes it.

Sources. Australian Bureau of Statistics, Characteristics of Australian Business, 2024-25, released 25/06/2026. RAND Corporation, Ryseff, De Bruhl and Newberry, The Root Causes of Failure for Artificial Intelligence Projects and How They Can Succeed, 13/08/2024. McKinsey, The State of AI, 05/11/2025. Boston Consulting Group, From Potential to Profit: Closing the AI Impact Gap, 14/01/2025. Reserve Bank of Australia, Technology Investment and AI: What Are Firms Telling Us?, Bulletin, 13/11/2025. Every figure quoted has been checked against the publishing organisation's own document. Where a widely circulated figure could not be verified against its primary source, it has been left out.