

BEYOND THE DEMO

What AI-Ready Actually Means

Access to the technology arrived overnight. The discipline it demands did not.

Giovanni Leonardi

December 2022



What AI-Ready Actually Means

Transformation · No. 289

Originally written December 2022 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“Two black boxes stacked on top of each other do not make a decision; they make an incident waiting for an audit.”

1. The Screenshot on the Board's Agenda

Somewhere in your organisation, in the past three or four weeks, a person typed a plain-English question into a chat box and watched a coherent, confident, grammatically immaculate answer assemble itself in seconds. They showed a colleague. The colleague showed someone else. By now the screenshot has travelled far enough that a member of the board has forwarded it downward with a subject line that is some variation of a single question: what is our plan for this?

It is a reasonable question, asked in an unreasonable state of excitement, and it almost always carries a hidden assumption — that the hard part has just been solved. The tool is here. It is astonishing. It is, for the moment, more or less free. Surely readiness is now a matter of moving quickly, before competitors do.

We should slow down on precisely that assumption, because it has the logic backwards. The arrival of a machine that anyone can use has not made organisations ready for artificial intelligence. It has made visible, for the first time and rather brutally, how ready they are not.

2. Readiness Was Never a Procurement Problem

For most of the past decade, "getting ready for AI" was sold — and bought — as a purchasing exercise. Stand up the data lake. Hire the data scientists. License the platform. Run the proof of concept the vendor helpfully offered to co-fund. The barrier to entry was money and scarce talent, and so readiness looked like a budget line.

What the last few weeks have done is collapse that barrier to something near zero. The capability a well-funded team could not reliably build a few years ago now sits behind a web address a summer intern can reach. And when the expensive, scarce thing suddenly becomes cheap and abundant, it stops being the constraint. Something else becomes the constraint. That something else is the organisation itself — and it always would have been, which is why the procurement framing was a comfortable evasion all along.

When the model is free, the constraint is no longer the technology. It is the quality of your data, the clarity of your decisions, and the discipline of the people reading the output.

This is uncomfortable, because an organisation is a great deal harder to fix than a contract is to sign. But it is also clarifying. For the first time we can see what readiness actually requires, now that access has stopped standing in for it.

3. What "Ready" Actually Means

Strip away the platform and the pilot, and readiness resolves into four unglamorous conditions. Not one of them is technical in the way the word usually implies.

- Data you can stand behind. A model is only ever as good as what it is given, and most organisations sit on data they quietly do not trust — duplicated, undefined, scattered across systems that disagree with one another about basic facts. For years this was tolerable because humans in the loop silently corrected for it. Feed the same data to a machine that corrects for nothing, and the mess is not filtered; it is amplified and handed back to you in fluent prose.
- Decisions with a traceable logic. If you cannot today explain how a particular judgement gets made in your organisation — what goes in, who weighs what, where the accountability sits — then you are in no position to hand any part of it to a system that also cannot explain itself. Two black boxes stacked on top of each other do not make a decision; they make an incident waiting for an audit.
- Processes stable enough to augment. You cannot augment a process you cannot describe. Automating a chaotic workflow does not tame it — it produces the same chaos faster, and with more confidence. The processes worth pointing this technology at are the ones you already understand well enough to know exactly where a machine would help and where it would quietly do harm.
- The temperament to tell plausible from correct. This is the subtle one, and the most important. The defining property of these systems is fluency: they produce answers that are well-formed, confident, and persuasive whether or not they are right. That is precisely what makes them dangerous inside an organisation trained to treat a polished document as a reliable one. Readiness, at the human level, is having people who can read a beautifully written answer and still ask whether it is true.

Notice that none of these is bought. Each is earned, slowly, in the years before the interesting technology arrives — which is why the organisations now scrambling to look ready cannot become ready in a quarter, however large the cheque.

4. The Objection Worth Taking Seriously

There is a strong counter-argument, and it deserves better than to be waved away. It runs like this: you are overthinking a productivity tool. People are already getting real value from it — drafting, summarising, unsticking themselves on a blank page — and every week you spend building "readiness" is a week a nimbler competitor spends simply using

the thing. Readiness is a luxury framing that will leave you standing still.

Much of this is correct, and I would not slow the part that is. An individual using one of these systems to draft an email, rough out a first version, or think aloud needs no organisational readiness at all; the judgement stays with them, the stakes are contained, and the value is immediate and real. Adopt that freely, and quickly.

But there is a hard line between an individual drawing value from a tool and an organisation making decisions with one, and the whole of the readiness question lives on the far side of it. The first needs nothing from you. The second needs everything on the list above. The failure mode of this moment is not being too slow to let people experiment – it is mistaking the ease of the first thing for permission to do the second. That is how an organisation comes to scale its trust in a system faster than it has earned the right to.

“The danger is not adopting too slowly. It is trusting too quickly.”

5. A Small, Expensive Illustration

Consider a composite that is beginning to play out, in one form or another, in a great many finance functions. A shared-services team responsible for month-end reporting wires one of these language models into its variance commentary – the narrative that explains why each line moved against budget. It is a natural fit. The task is repetitive, language-heavy, and swallows the better part of two analyst-days every month. The machine does it in minutes, in clean prose, and for a full quarter the output is indistinguishable from what a competent analyst would have written.

Then comes a month where a number moves for a reason the model cannot possibly know – a one-off reclassification, agreed verbally in a meeting that left no trace in the ledger. The model does what it always does: it produces a fluent, confident, entirely plausible explanation. It is also entirely wrong. And because it reads exactly like the eleven true commentaries around it, it passes untouched into the board pack. Two analyst-days saved a month, quietly traded for one distorted decision at the top of the house.

The instructive part is where the failure actually sat. It was not in the model, which behaved exactly as designed. It was that, somewhere in the enthusiasm of automating a tedious task, the organisation stopped noticing that the tedious task had contained a judgement – and it removed the judgement without ever deciding to. Nobody owned the question the analyst had been quietly answering all along: does this explanation actually make sense? Readiness would have been knowing that the judgement was there, and who held it, before handing away the task that hid it.

6. The Boring Work Is the Real Work

Here is the part that should reframe the board's forwarded screenshot. Almost everything genuine readiness requires — trustworthy data, decisions you can trace, processes you can describe, people who can tell fluent from correct — is work you should have done anyway. It is the governance no one is promoted for, the decision rights nobody enjoys litigating, the documentation everyone defers. The technology did not create these gaps. It merely made them impossible to keep ignoring.

So the honest answer to what is our plan for this? is neither a tool selection nor a race. It is to treat this extraordinary moment as the audit it actually is. Let people experiment where the stakes are contained; they will learn faster than any strategy deck. But before you let a machine anywhere near a decision that matters, ask the four unglamorous questions — and be willing to hear that the answer, for now, is not yet.

The organisations that look prescient a few years from now will not be the ones that moved first. Everyone has the same access; access confers no advantage. They will be the ones that spent the unfashionable years building the substrate underneath — and who, when the astonishing tool finally arrived, already had somewhere solid to plug it in.