

BEYOND THE NATIVE LABEL

AI-Native Is Not a Destination You Can Declare

Why the fashionable divide between native and augmented organisations obscured the only 2024 choice that changed outcomes — who was allowed to move a decision boundary

Giovanni Leonardi

August 2024



AI-Native Is Not a Destination You Can Declare

Transformation · No. 324

Originally written August 2024 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

"Augmentation is what you get when the technology changes and the decision rights do not."

The distinction that flatters us

The slide said AI-native by design. It was the third bullet on the operating-model page of a strategy refresh, and it carried the settled confidence of a decision already taken. In the same fortnight, a service team two floors down had quietly stopped opening the assistant that had been rolled out to them, with some fanfare, in the spring. The tool could draft a courteous reply to a customer; it could not see the system where that customer's history actually lived, so every draft had to be reconciled against a screen the model could not read. The team went back to typing by hand. Both things were true in the same building in the same week: the organisation was declaring itself native to a technology that, at the coalface, it had not yet learned to hold.

This is the year of the distinction between the AI-native organisation and the AI-augmented one. It is on conference stages and in board papers, and it flatters everyone who uses it. To be augmented sounds transitional, a little apologetic — the new engine bolted onto the old chassis. To be native sounds like arrival. The trouble is that, as a guide to what an incumbent should actually do, the distinction is close to useless, and in more than a few places it is quietly doing harm.

I want to argue something against the grain of the moment: for almost every organisation that is not a start-up founded in the last eighteen months, AI-native is not a destination that can be reached by declaration, and chasing it as an identity distracts from the change that would actually matter. The useful question is not whether the enterprise is native or augmented. It is whether AI has been allowed to touch the one thing transformations habitually leave alone: the distribution of decision rights.

What "native" turned out to mean

Strip away the branding and look at what organisations calling themselves AI-native this year have actually built. In the honest cases, it is a genuinely reworked process or two: an underwriting triage, a first-line support flow, a document-review step where the model does the first pass and a human adjudicates. Real, valuable, worth having. But native

implies something total — that the organisation's default way of working now assumes the machine. Against that meaning, the label rarely survives contact with the org chart.

The pattern that recurs is an organisation that has changed its tooling and its vocabulary while leaving its structure, its incentives, and its accountability lines exactly where they were. The copilot is licensed to four thousand people; the daily-active number is nearer three hundred; and of those three hundred, most are using it to smooth an email rather than to do the work it was bought to change. Adoption is reported at the licence line, because that is the number that flatters. The operating model underneath has not moved an inch. This is augmentation wearing a native costume, and everyone in the delivery chain can see it even when the strategy deck cannot.

Declaring an organisation AI-native changes the words on the slide. Whether the model is allowed to change who decides what is the only thing that changes the organisation.

The fault line that actually mattered

Here is the distinction I would put in place of native-versus-augmented. In every deployment worth studying this year, the outcome turned on a single question: did the introduction of the model come with a change to decision rights, or did it not?

Consider two organisations that bought, in effect, the same capability: a model that reads an inbound case, assembles the relevant history, and proposes a resolution. In the first, the deployment stopped at proposal. The model drafts; the human does everything they did before, plus a review of the draft. Handling time per case went up by a measurable margin in the first quarter — from roughly nine minutes to eleven — because reading and correcting someone else's plausible but unaccountable answer is slower than composing your own. The tool was blamed. It was not the tool's fault. Nobody had been given permission to let it decide anything.

In the second, the same capability arrived with a redrawn boundary: cases below a defined confidence and value threshold were resolved by the model and sampled after the fact, while everything above the line went to a human with the draft attached. That is a change to decision rights — a real one, with a named owner for the errors it would inevitably produce. Handling time fell, but that was almost beside the point; what changed was that a class of decisions had a new default author, and the organisation had decided, explicitly, who carried the consequences when the default was wrong.

The difference between those two was not budget, model quality, or AI maturity. It was whether anyone with the authority to move a decision boundary had actually done so. Augmentation is what you get when the technology changes and the decision rights do not. And by that definition, most of what is being sold as native this year is simply augmentation that has been allowed, in one or two places, to touch the organisation's real nervous system.

The strongest case for going native

The honest objection to all of this is worth stating at full strength, because there is truth in it. The advocates of AI-native design are not confused; they are pointing at something real. Their case is that incremental augmentation is a trap: that if you only ever bolt the model onto processes designed for humans, you inherit all the coordination overhead those processes carry, and you will never see the order-of-magnitude gains that come from redesigning the work around what the machine is good at. On this view, the greenfield player who assumes the model from the first line of the process will eventually make the augmented incumbent uncompetitive — the way businesses built for the web outran those that treated it as a new sales channel.

That argument is correct — for the greenfield player, and for the specific processes an incumbent is willing to genuinely rebuild. Where it misleads is in the leap from this process should be redesigned around the model to the organisation should become AI-native. The first is a design decision you can make, fund, and own. The second is an identity, and identities are exactly the things organisations adopt rhetorically while changing nothing underneath. The useful half of the native argument is not the label; it is the instruction to redesign the work rather than decorate it. Keep that half. Drop the costume.

What the augmented organisation should actually do

So the advice that follows from a year of watching this is unfashionably modest, and I would stand behind it. Stop trying to become AI-native. Instead, pick the small number of processes where you are genuinely willing to move a decision boundary, and move it — with a named owner for the errors, a threshold that is written down, and a sampling regime that assumes the model will sometimes be confidently wrong. Then leave everything else honestly augmented, and say so, rather than dressing licence counts up as transformation.

The reason this matters is not tidiness. It is that the native-versus-augmented framing lets leaders feel they have made a strategic choice when they have made a linguistic one, and the gap between the two is where credibility quietly bleeds away — with the board, and, more corrosively, with the service team on the second floor who stopped opening the assistant and noticed that no one asked them why. The organisations that will look native in a few years will not be the ones that declared it. They will be the ones that, unglamorously, kept moving decision boundaries one defensible step at a time.