

BEYOND THE PLATFORM DECISION

Enterprise AI Platforms — Build, Buy, or Orchestrate: Why the Question Itself Is the Problem

A Practitioner's Perspective on Platform Strategy in the Age of Enterprise AI

Giovanni Leonardi

July 2024



Enterprise AI Platforms — Build, Buy, or Orchestrate: Why the Question Itself Is the Problem

Transformation · No. 386

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

"The build-versus-buy decision assumes a stable destination. Enterprise AI in 2024 has no stable destination."

The Familiar Fork in the Road

Every enterprise technology cycle produces its version of the same strategic question: should we build our own, buy a commercial solution, or assemble something from parts? The artificial intelligence wave is no different in asking the question — but it is fundamentally different in why the usual answers fail.

In my experience, the organisations now wrestling with their AI platform strategy are applying a decision framework designed for an earlier era of technology procurement. They are treating AI platforms as they once treated ERP systems or cloud migration: as a procurement decision with a clear evaluation matrix, a shortlist, and a winner. The pattern I have observed across sectors in 2024 is that this framing consistently produces the wrong outcome — not because the analysis is flawed, but because the question itself misunderstands what an AI platform actually is.

The Build Trap and the Buy Illusion

The "build" camp typically argues from a position of strategic differentiation. The logic runs: AI is so central to competitive advantage that ceding platform control to a vendor is an existential risk. In practice, what I have seen is that organisations pursuing a full build rapidly discover that the infrastructure layer — the compute orchestration, the model serving, the vector storage, the evaluation pipelines — consumes the vast majority of engineering effort while producing zero differentiated capability. The team hired to build intelligent applications spends eighteen months building plumbing.

The "buy" camp makes the opposite bet: that a commercial platform will abstract the complexity and let internal teams focus on business logic. The difficulty, which has become painfully visible over the past twelve months, is that commercial AI platforms in their current state are not platforms in the way that mature enterprise software is a platform. They are moving targets. The underlying models change quarterly. The

capabilities available today may be deprecated next quarter. The integration patterns that worked with one model generation may not transfer to the next. Organisations that committed to a single vendor's AI stack in early 2023 have, in many cases, already had to substantially rearchitect — not because they chose badly, but because the technology shifted beneath them.

This is the core of the problem: the build-versus-buy decision assumes a stable destination. Enterprise AI in 2024 has no stable destination.

The Orchestration Instinct

The more sophisticated organisations — and they are still a minority — have arrived at what I would call the orchestration instinct. Rather than choosing a single platform, they are designing a thin integration layer that allows them to swap models, combine capabilities from multiple providers, and maintain optionality as the landscape shifts.

This is closer to the right answer, but it carries its own risks that few are acknowledging openly. An orchestration approach requires a level of architectural discipline that most enterprise technology functions do not possess. It demands clear abstraction boundaries, robust evaluation frameworks to compare model performance across tasks, and — critically — a governance model that can handle the complexity of multiple AI providers with different data handling commitments, different performance characteristics, and different pricing structures.

The organisations that are succeeding with AI platform strategy are not the ones that made the best choice between build, buy, or orchestrate. They are the ones that designed for the ability to change their choice.

What I observe in practice is that the orchestration approach often degenerates into accidental complexity. Teams integrate three or four model providers without a coherent abstraction layer, and within months the "orchestration" has become a tangle of point-to-point integrations that is harder to manage than a single-vendor lock-in would have been.

What Actually Works

The pattern that recurs across the more successful implementations I have observed is not a platform choice at all — it is a design principle. These organisations have separated three concerns that the build-buy-orchestrate framing unhelpfully conflates:

- The capability layer — which models and services are used for which tasks. This is expected to change frequently, and the architecture is designed to make switching low-cost.
- The integration layer — how AI capabilities connect to enterprise systems, data sources, and business processes. This changes slowly and is where genuine

architectural investment belongs.

- The governance layer — how the organisation manages risk, monitors performance, controls cost, and maintains accountability across whatever AI capabilities are in use. This is permanent infrastructure, regardless of which models sit beneath it.

The organisations that have separated these three layers are navigating the current turbulence well. They can adopt new model capabilities quickly because the integration and governance layers do not need to change. They can switch providers when the economics or performance shift because the capability layer is deliberately loosely coupled.

The organisations that have not made this separation — which is the majority — are trapped in a cycle of strategic re-evaluation every time the landscape shifts, which in 2024 means roughly every quarter.

The Uncomfortable Implication

The uncomfortable implication of this analysis is that the most important AI platform decisions are not about AI at all. They are about integration architecture and governance design — disciplines that most organisations have systematically under-invested in for years. The glamour has gone to the model selection, the proof-of-concept demonstrations, the executive briefings on the latest capability announcements. The unglamorous work of building a robust integration layer and a governance framework capable of handling a multi-model, multi-provider reality has been treated as a second-order concern.

In my experience, the organisations that will look back on 2024 as the year they got their AI platform strategy right will not be the ones that chose the best model or the best vendor. They will be the ones that invested in the architecture of adaptability — the boring, essential infrastructure that allows them to treat platform decisions as reversible rather than existential.

The question is not build, buy, or orchestrate. The question is: have you designed your enterprise to change its mind?