

BEYOND THE DEMO

Your AI Strategy Is Probably an Unfinished Cloud Strategy

Why the generative-AI race is exposing unfinished enterprise cloud transformations

Giovanni Leonardi

October 2023



Your AI Strategy Is Probably an Unfinished Cloud Strategy

Transformation · No. 296

Originally written October 2023 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“An AI strategy that cannot name its data, identity and operating constraints is not a strategy; it is a catalogue of demonstrations.”

The boardroom question has changed

By the autumn of 2023, a familiar scene was repeating in large enterprises. A leadership team would ask for an “AI strategy” and receive, within weeks, a polished inventory of opportunities: summarising service conversations, drafting technical documents, searching policy libraries, assisting developers, accelerating research. The list looked new. The obstacles beneath it did not.

In one representative programme, eighteen generative-AI use cases were scored in a two-hour workshop. Three prototypes were built in six weeks. Only one could be placed in front of employees without extensive manual preparation. The other two exposed the same unfinished business: documents spread across twenty-seven repositories, access rights that could not be inherited reliably, no agreed route for sensitive prompts, and interfaces that had never been designed for machine-to-machine use.

The programme had not discovered an AI problem. It had discovered the consequences of a cloud and data transformation that had stopped at migration.

The strategy–recycling pattern

The pattern is easy to recognise. A new technology creates executive urgency; the organisation gives the urgency a new strategy label; and the new strategy quietly inherits unresolved decisions from the previous transformation. In 2023, generative AI is the label. Underneath sit questions that cloud programmes have been debating for years:

- Where does authoritative data live, and who may use it?
- Can identity and entitlement travel with the information?
- Can systems expose reliable services rather than depend on manual extraction?
- Can usage, cost and performance be observed at workload level?
- Who decides when an experiment becomes an operated service?

These are not secondary implementation details. They determine whether a language-model demonstration can become a dependable business capability.

Generative AI is acting as an audit of enterprise architecture. It is revealing, with unusual speed, which cloud transformations changed the operating model and which merely changed the hosting location.

The distinction matters because the wrong diagnosis produces the wrong investment. If leaders believe the constraint is model selection, they fund more trials. If the constraint is data access, integration and control, more trials merely create a larger queue of attractive demonstrations waiting for foundations.

What the AI label conceals

The new vocabulary can make old dependencies harder to see. A retrieval-assisted answer may feel like an AI feature, but its reliability depends on document quality, indexing, permissions and change control. A conversational interface may look self-contained, but it still requires secure connections to systems of record. A model may be consumed as a service, but somebody must still manage capacity, latency, supplier terms, logging and cost.

What leaders call it	What is often unresolved underneath
“Enterprise knowledge assistant”	Document ownership, search quality, retention and access rights
“Intelligent service agent”	Process APIs, case data quality, escalation rules and audit trails
“Developer assistant”	Source-code boundaries, approved environments and assurance responsibilities
“AI platform”	Cloud landing zones, identity, workload observability and commercial controls

An AI strategy that cannot name its data, identity and operating constraints is not a strategy; it is a catalogue of demonstrations.

This is why the most revealing question is not, “Which use cases shall we pursue?” It is, “What must be true for the second and tenth use cases to reach operation more safely and cheaply than the first?” The answer exposes whether the organisation is building reusable capability or repeatedly staging bespoke experiments.

The serious objection

There is a strong case against turning every AI initiative into a broad infrastructure programme. Models, suppliers and techniques are moving quickly. Large foundational investments can harden assumptions that may be obsolete before procurement is

complete. Waiting for perfect data or a universal platform is another familiar way to avoid learning.

That objection is right. The answer is not to complete an immaculate cloud transformation before doing anything with AI. It is to connect experimentation to the smallest reusable foundations that real use cases prove necessary.

A sound sequence in late 2023 is therefore deliberately narrow:

1. Choose two or three workflows with different data and risk characteristics.
2. Build prototypes quickly enough to test value, not merely technical possibility.
3. Record every production barrier exposed by the prototypes: access, integration, quality, assurance, cost and ownership.
4. Fund only the shared capabilities that remove barriers across more than one workflow.
5. Require the next use case to reuse those capabilities and show whether unit cost, approval time or deployment effort falls.

This avoids both extremes: the disconnected pilot factory and the multi-year platform programme justified by hypothetical demand.

Why this matters now

The generative-AI surge has compressed expectations. Senior teams that tolerated slow progress in data catalogues, service integration or identity modernisation are now asking why an apparently impressive prototype cannot be deployed next quarter. That impatience can be useful, but only if it is directed at the real constraint.

The practical consequence is a shift in ownership. AI cannot remain a specialist innovation agenda while cloud, data, security and operations continue on separate roadmaps. Nor should it simply be absorbed into an infrastructure programme and stripped of business urgency. The work needs one decision forum that can trade value against readiness: business owners bringing consequential workflows; technology leaders exposing architectural dependencies; data owners resolving authority and quality; security and legal functions setting usable boundaries; finance testing whether consumption economics hold beyond the prototype.

The organisations likely to move fastest will not be those with the longest list of AI ideas. They will be those that use each experiment to finish a small piece of the operating foundation, then reuse it. Their advantage will look like model sophistication from the outside. Internally, it will be disciplined cloud, data and governance work finally connected to a business deadline.

The AI strategy is real only when it makes those inherited decisions visible. Until then, the organisation is not choosing an AI future. It is renaming its cloud past.