

CASE-06 · TECH / AI — BUILD OR BUY

Building our AI platform in-house: €40 million and 18 months

Developing internally — for «control, differentiation and intellectual property»
— instead of buying the technology or partnering.

THE JUDGMENT

in-house → buy + defend

CALIBRATED CONFIDENCE

35%

SECTOR

Tech/AI



EXECUTIVE VERDICT

In eighteen months the model is a **commodity**: the moat is data and distribution.

THE CALL

in-house → buy + defend

CONFIDENCE

35%

IN ONE LINE

Recovered **12+ months** of market lead and avoided spending on a part destined to become a commodity.

§ THE PROTAGONIST & THE CONTEXT

The CTO of a B2B software company with an established client base, tasked with deciding whether to build AI in-house or rely on third-party providers. The board has already put «native AI» in the pitch deck for the next funding round.

A vertical software company with ten thousand enterprise clients, strong on process data in a specific domain. It has mid-level engineering resources and an already crowded product roadmap. The «AI project» started as an internal initiative by a small research group and has taken on a life of its own.

§ BACKGROUND

The field of large language models and general-purpose AI systems has transformed from a barrier to entry into shared infrastructure in the space of twenty-four months. The main API providers (models, embeddings, retrieval) have reduced costs by an order of magnitude in two years and continue to iterate every three to six months. Companies that invested in proprietary model development in 2022–2023 found themselves competing against freely available foundations in 2024. In this context, real competitive advantage has shifted to those who accumulate exclusive data and build product integrations that are hard to replicate — not those who train the largest model.

§ COUNTER-INTUITIVE SCENARIO

In a field that changes every three months, 18 months of in-house development means launching something **already old**. The model is becoming a commodity, something everyone has: the real defensible advantage (the «moat») is **proprietary data and sales channels**. Building the model burns your precious time on the wrong part.

IMPLICIT ASSUMPTIONS

- ◆ The advantage is in the model, not in the data or the ability to sell.
- ◆ The team delivers in 18 months.
- ◆ The technology won't shift under our feet in the meantime.

FALSIFICATION TESTS

- ◆ Will the same capability be buyable «off the shelf» (via API) within a few months?
- ◆ Timing compared: buy in 3 months versus build in 18.
- ◆ What do clients pay for: the model, or the product and the data around it?

QUESTIONS THAT RAISE THE BAR

- ◆ If in 12 months that technology belongs to everyone, what justifies 18 months building it?
- ◆ How much does the market we lose by waiting cost us?

CALIBRATED CONFIDENCE & PROVENANCE

that «build everything» beats «buy and focus on data and sales»

Provenance: vendors' public roadmaps · client interviews · time-to-launch comparisons · red-team base.

Time to market: build vs buy

ILLUSTRATIVE DATA

Build in-house (proprietary model)

18 months

Buy + product integration

3 months

Estimated time to get an AI capability into production under each strategy — the gap is the window in which competitors can consolidate (illustrative data).

Where clients recognise value

ILLUSTRATIVE DATA

54%

DOMAIN DATA
AND INSIGHTS


31%

WORKFLOW
INTEGRATION

9%

AI MODEL
QUALITY

6%



OTHER

Share of enterprise clients citing each factor as their «main reason to renew» — the AI model itself is nearly irrelevant compared with data and integration (illustrative data).

OUTCOME

Reversed: **buy what is now commodity, build what truly defends**. Model purchased, resources shifted to proprietary data and channels.

VALUE

Recovered **12+ months** of market lead and avoided spending on a part destined to become a commodity.

METHODOLOGICAL NOTE — READ FIRST

Composite cases, in the method of the Harvard Business Review: reconstructions based on real, recurring situations in each sector, merged and anonymized to protect confidentiality. The decision dynamics are authentic; names, figures and details are altered and not traceable to any single client or case. The «provenance» notes describe the type of evidence the engine cites with traceability in production. The Δ -CSI values illustrate the intensity of the pressure the contradiction put on the assumptions.