

The Build-Buy-Partner Grid

Most banks deploy artificial intelligence in-house they should have bought. This is the decision architecture for sourcing every AI capability — and for knowing the few where building in-house is the right call, not the expensive one.

"We build only where we own a data moat — buy or partner for everything else — and we almost never build our own model."

82–85%

of in-house bank AI builds fail
(reported) · vs 25–50% for normal IT

~95%

of pilots never reach production
(reported)

~56%

lower 3-year TCO when you buy vs
build (reported)

Week 6 of the AI Business Architect series · *Stop Building Your Own Banking AI*

THE PROBLEM

The most expensive instinct in banking

The build bias hides behind a flattering story: we are a technology company now, our engineers can build anything, the cost of code is falling. All three are partly true. None justify what building actually costs.

Roughly eighty percent of in-house bank AI builds fail — and not because the algorithms are wrong. They fail because the data and infrastructure underneath cannot carry the model in production. A fraud model with perfect sandbox accuracy is useless if it cannot reach real-time data behind a legacy core.

The economics, side by side

	BUILD (IN-HOUSE)	BUY (COMMERCIAL)
Upfront	\$500K–\$2M	\$50K–\$200K
Time-to-value	12–24 months	3–6 months
3-year TCO	baseline	~56% lower
Ongoing opex	30–40% of capex / yr	included

When the instinct is to build, leaders should ask one thing first: do we own a moat here — or are we about to rebuild something we could have bought this quarter?

"These aren't AI failures. They're infrastructure failures wearing an AI label."

THE FRAMEWORK

Two axes, four verdicts

X · Strategic Differentiation — a defensible moat, or a commodity competitors can buy off the same shelf?

Y · Internal Readiness — the proprietary data, talent, and architecture to build *and sustain* it.

BUILD

High differentiation · High readiness

A real moat you can sustain on proprietary data + elite talent. The only quadrant that justifies an in-house build.

PARTNER

High differentiation · Low readiness

The edge is real but you can't sustain it alone — co-develop with a specialist, keep the proprietary workflow.

BUY

Low differentiation · any readiness

A commodity function. Purchase it, take the fast time-to-value, carry none of the maintenance.

ASSEMBLE

Low edge · High readiness

Fine-tune or retrieve on a bought foundation model instead of training from scratch.

Build your own model? Training from scratch: **\$78–100M**. Fine-tuning: **\$500–5,000** — 60–90% less. For a bank, **Assemble — don't Build**.

The Capability Map

CAPABILITY	VERDICT	WHY
Fraud detection on your own ledger	BUILD	Proprietary graph + latency a vendor can't match
Credit underwriting on proprietary data	BUILD	The data moat is the model
AML network / entity resolution	PARTNER	Differentiating but niche engineering
Wealth-advisory copilot	PARTNER	Proprietary archive + guardrails on a partner model
Knowledge retrieval / search	ASSEMBLE	RAG / fine-tune on vetted internal docs
OCR · HR chatbot · code-assist	BUY	Commodity — building wastes quant talent
Your own foundation model / LLM	ASSEMBLE	Fine-tune, don't train from scratch

Three forces regulated finance adds

1 · Liability stays home

Buying a model doesn't move the regulatory responsibility. Demand vendor audit logs + explainability.

2 · Speed vs opportunity cost

3 months often beats 24 — the cost of waiting exceeds the contract.

3 · Lock-in

"Tokenflation" + outages. Architect for model agnosticism — swap providers by task.

Buying is the easy part. The value is redesigning the operating model around the tool — the purchase order is not the strategy.

THE ACTION

The Sourcing Decision Loop — 4 steps

- 1 **Score differentiation.** Moat or commodity? If a competitor can buy the identical capability, it is not a moat.
- 2 **Score readiness.** Can you sustain it — data, talent, pipelines — not just ship a pilot?
- 3 **Route the verdict.** Build · Partner · Buy · Assemble. For any Build, demand the data moat in writing.
- 4 **Govern the vendor.** Deployer controls, a concentration limit, and a model-agnostic exit before you sign.

Proof: Capital One & BofA **build** on proprietary data · HSBC & Morgan Stanley **partner** · DBS **assembled** ~SGD 1B of value by orchestrating 1,500+ models. Outsourcing the model never outsources the liability.

The regulation rewards the discipline

SR 26-2 (scales by materiality; GenAI governed) · EU AI Act deployer duties (oversight · FRIA · right to explanation) · DORA concentration risk · Vietnam Decree 356/2025 (ex-ante transfer review). And shared models on shared clouds create systemic concentration risk.

Stop building what you can buy. Build only what you alone can own.

► Watch the 8-minute briefing: youtu.be/a8z_pktfhFU