

A CXO Briefing · Getting AI to Production

# What it actually takes to ship enterprise AI

## Field notes from building and running AI systems in production

Most AI projects stall not because the model fails, but because everything around it was never built. This briefing is about that everything, drawn from systems we have put into production for Indian enterprises.

### Executive Summary

Most enterprise AI dies in the gap between a demo that impresses and a system that survives real use.

The failure numbers are real: Gartner expects more than 40% of agentic AI projects to be cancelled by the end of 2027, and a 2025 MIT study of over 300 initiatives found that roughly 95% delivered no measurable return. But the cause is rarely the model. By the time a model is famous enough for you to read about it, it is already good enough for most business work.

What actually separates the systems that ship is unglamorous, and it is where almost all the effort goes: choosing a narrow problem worth solving, taming messy real-world data, getting the compliance and data path right, and building something the client can run without the people who built it. The lessons below come from putting working systems into production for firms in real estate, financial services, and industry, not from theory.

This briefing covers where the real work sits, six lessons from shipping, the difference between a demo and a production system, what a first project actually looks like, and a short check to run before you greenlight one.

# The model is the easy part

This is the single most useful thing to understand before you fund an AI project, and the one that surprises most leaders.

In a real build, the call to the AI model is often the smallest, cheapest, simplest line in the whole system. On one of our live systems the AI processing costs a few hundred rupees a year; the platform around it is almost the entire bill, and almost the entire effort.

The work that decides success is the unglamorous part: pulling in messy inputs that never look like the demo, making data save and stay saved across many users at once, handling the edge cases real life throws up, and meeting the compliance rules your sector lives under. None of that is about which model you chose. A project that treats the model as the hard part, and everything else as a detail, is the project that stalls.

## 02 The Lessons

# Six lessons from putting AI into production

None of these are theoretical. Each one is a place we have seen projects either hold up or fall over, and each comes with the cost of ignoring it.

1

LESSON 01

### Start with one painful, measurable workflow

The systems that ship solve one specific, irritating job with a number attached: a query that takes days, work double-handled across a team, a report that swallows a morning. "AI for the company" has no owner and no finish line.

#### IGNORE IT AND

Scope sprawls, no one can say whether it worked, and the project dies of ambiguity.

2

LESSON 02

### Budget for the system, not the model

The AI call is the cheap part. The real work is ingesting messy inputs, making data persist and sync correctly, and handling the cases that never appear on stage. Most of what determines success has nothing to do with the model you picked.

#### IGNORE IT AND

You get a dazzling demo that quietly breaks the first week real data hits it.

3

LESSON 03

### Settle the data path before the model

Where data is allowed to live and be processed often rules out options before quality is even discussed, especially for regulated firms. Decide residency, retention, and access first, then pick a model that fits inside those walls.

#### IGNORE IT AND

You build on a model you are later forced to rip out for a compliance reason you could have seen on day one.

4

LESSON 04

## Never depend on a single model

Models get deprecated, repriced, or go unavailable in your region. A primary with a cheaper fallback behind it turns an outage or a price change into a configuration switch, not a crisis, and means no vendor can hold your system hostage.

### IGNORE IT AND

One vendor's bad day or pricing decision becomes your outage, with nothing to fall back to.

5

LESSON 05

## Make the output a work product, not a chat

Value shows up when the system hands a person something they can act on at once: a triaged action with the next step spelled out, a ready document, a flagged exception. Not a chat window they have to coax. Build for the user's real job.

### IGNORE IT AND

You ship a clever tool no one adopts, because using it is more work than the old way.

6

LESSON 06

## Build it so they can run it without you

A system that needs its builder for every change is not shipped, it is leased. Self-service admin, plain-English error logs, backups the team can restore, and a clean handover are what make a project something the client owns. And plan for hardening after launch: real use always surfaces what the demo could not.

### IGNORE IT AND

You leave behind an expensive dependency that stalls the moment the builder steps away.

### 03 The Gap

## What "done" looks like, versus what shipped looks like

A demo and a production system can look identical in a meeting. They are not the same thing, and the distance between them is where budgets and timelines quietly disappear. This is the gap your technology lead is paid to close.

	Looks done The demo	Actually shipped Production
<b>The inputs</b>	Clean, hand-picked sample data	Messy reality: odd date formats, forwarded email chains, duplicates, blanks
<b>The data</b>	Held in the browser or a spreadsheet	Saves, syncs across users, dedupes, and survives a refresh or a crash
<b>The model</b>	One impressive model	A primary with a fallback, chosen to fit the data path and compliance
<b>The output</b>	Plausible-looking text	A structured work product someone acts on without rechecking it
<b>When it breaks</b>	It doesn't, on stage	It fails safely, is logged in plain English, and can be recovered
<b>Who runs it</b>	The people who built it	The client's own team, with documented handover and admin controls

In our experience the bugs that surface after go-live are almost never about the AI. They are about data not saving, records duplicating on re-upload, a real-world date format the system rejected, or a screen losing the user's place under live updates. Plan for this hardening phase; it is normal, and it is where a system earns trust.

#### 04 The Shape

## What a first project actually looks like

A useful first AI project is small, fast, and narrow on purpose. The aim is one working thing in real hands, not a platform on paper. A rough shape:

## From decision to a system in real hands

Indicative, not a fixed timetable. The point is the order, not the exact weeks.

### Weeks 1–2

Pick the workflow, agree the success metric, and settle the data path. The decisions that prevent expensive rework later.

### Weeks 3–6

Build the narrow version on real data, with a model fallback and the plumbing that makes it hold up under daily use.

### Week 7 on

Put it in real hands, harden what live use exposes, then expand from a base that already works.

**The pattern:** most of the risk is retired in the first two weeks, before much is built at all. The firms that rush past those weeks pay for it in month three.

## 05 The Takeaway

# A 30-second check before you greenlight

Run a proposed project against these four. The verdict tells you whether to start now, narrow it first, fix the foundations, or wait.

### A named owner and a success metric that is a number

Someone accountable, and a clear definition of what winning looks like

GO

### Right idea, but the scope is too broad to ship in weeks

The ambition is sound; the first cut needs to be smaller

NARROW IT FIRST

### Sound use case, but the data is messy or scattered

The idea works only once the underlying data is usable

FIX DATA FIRST

### No clear owner or no agreed measure of success

An experiment with no accountability and no finish line

NOT YET



EnableWealth

AGENTIC AI & AUTOMATIONS

enablewealth.ai

This briefing is general guidance for business leaders, not investment or legal advice. The lessons here are drawn from EnableWealth's own production deployments across real estate, financial services, and industrial clients; all examples are generalised and anonymised under client confidentiality.

Third-party figures cited as published: Gartner, "Over 40% of Agentic AI Projects Will Be Canceled by End of 2027" (June 2025); MIT study of enterprise AI initiatives, as reported 2025. Figures may be revised over time; verify before relying on them. Current as of June 2026.