ESSAY · ENTERPRISE AI

You Can't Write a Five-Year AI Plan. Write a Five-Year Commitment.

AI capability turns over in months while strategy plans run five years. The fix isn't a shorter plan — it's sorting every line by how long it stays true, and committing only to the half that survives.

AI STRATEGY · LEADERSHIP · TRANSFORMATION

Burak Galiba · August 13, 2026 · 6 min read
https://burakgaliba.com/blog/five-year-ai-plan

SCAN TO READ ONLINE

The planning template goes out, and somewhere in it there is now a section for AI. Five years of it: the investment profile, the capability roadmap, the benefit case, the target state. It gets filled in, because the template asks and a blank section is not a survivable answer in a strategy review.

Then the plan is approved, and within about two quarters the AI section describes a world that no longer exists. The vendor named in it has been repriced. The architecture it assumes has been made unnecessary by a capability that now ships in the platform you already own. The flagship use case has been quietly overtaken. Nobody reopens the document, because reopening it is a governance event — so the plan stays on the shelf, formally in force and privately ignored, until the next cycle produces an equally perishable replacement.

Michael Yaziji at IMD names the mismatch precisely: AI technology evolves in three- to nine-month cycles, and organizations plan in three- to five-year ones. That gap is not going to close, and it is not a discipline problem you can fix by planning harder.

The obvious conclusion is the wrong one

The fashionable response is to declare the five-year plan dead. Plan in quarters, stay adaptive, respond to signals.

That option isn't actually available to you, and it's worth being honest about why. Capital allocation runs on multi-year commitments. So does hiring, org design, property, and every regulatory obligation with a fixed date attached. Your board's own calendar is built around a long-range view. An enterprise that stops making five-year commitments hasn't become adaptive — it has just stopped funding anything that takes more than a year to build, which in enterprise AI is most of what matters.

So the question was never whether to plan five years out. It's **which sentences survive five years** — and the reason AI sections age so badly is that almost nobody sorts them.

The half-life test

Take the draft AI section and go line by line with one question: *if this sentence is still in the document in three years, will it be embarrassing?*

The perishable half fails immediately. Named models and vendors. The reference architecture. The use case inventory — five years of use cases is fiction, and by year two a good share of them will be capabilities you buy rather than things you build. The maturity level you're targeting. The unit economics. All of it is real, all of it matters, and none of it belongs in a document that can only be amended by committee.

The durable half is quieter and much harder to write, because it isn't about technology at all:

- **Decision rights.** Who may approve an AI use case, what a business unit is allowed to do without asking, who owns the benefit when it lands. Every model you will ever deploy passes through this, and it is still true in 2031.

- **Data access as a commitment, not an aspiration.** The access layer nobody builds is a multi-year programme with contractual, architectural and ownership dependencies. It is the single most durable thing in the section, and usually the least specific.
- **The governance clock.** Compliance dates are already fixed and knowable years out. This is the one part of your AI section that can genuinely be planned across the full horizon — and the deadlines now have teeth.
- **Workforce shape.** Not headcount numbers, but the standing answer to what the capacity is for, and what you intend to employ versus rent.
- **The financial mechanism.** That a growing share of this spend is priced by consumption and scales with adoption, so the plan needs a governing mechanism rather than a fixed line. Get this wrong and you end up rationing tokens and calling it a cost strategy.

Durable in the five-year plan. Perishable in the operating cycle, replanned every two quarters, reported to the board rather than committed to it.

The objection deserves its due

There is a serious problem with what I've just described, and I'd rather name it than let a reader find it in the review.

A plan made entirely of enablers is a plan of nothing. Decision rights, access layers and governance clocks are exactly the language an AI programme uses when it wants budget without accountability — and a CFO asked to fund five years of enabling work with no euro attached will decline, correctly.

That objection lands harder than most people running these programmes realize, because the CFO is already the least convinced person in the room. In The Conference Board's 2026 C-Suite Outlook, 38% of CFOs named AI an investment priority against 59% of COOs and chief strategy officers — and the CFO is the one who owns the long-range financial plan you are trying to get into.

So the durable commitments have to be written as falsifiably as the perishable ones. That's the whole discipline:

- Not *"improve data access."* By the end of year two, the top twenty systems expose governed, documented, permissioned read access, and a named data owner answers an access request within ten working days.
- Not *"build AI capability."* This many roles in the operating plan are AI-capable roles, with the reskilling and hiring profile named and the budget attached.
- Not *"deliver X million in AI benefit."* Each business unit's operating targets contain a named, baselined AI contribution, signed by the person who owns the process — because a central AI value line belongs to nobody who can make it true.

Every one of those is testable at a date, by someone other than the person who wrote it. That is what makes them commitments rather than vocabulary.

And it's the missing half almost everywhere. EY's 2026 CEO Outlook — 1,200 CEOs across 21 countries, fielded in March and April — found around 80% increasing AI investment, and just **11%** linking AI impact to financial reporting with regular senior management review. Four in five are funding it. One in nine can see it in the numbers. A five-year plan that doesn't close that distance is a spending plan wearing a strategy's clothes.

Three clocks, one document

The output of all this is a single AI section that visibly runs on three clocks and says which is which.

Five years: decision rights, data access, the governance calendar, workforce shape, the financial mechanism. Reviewed annually, changed rarely, and each item testable at a date.

Annually: the budget envelope, the portfolio themes, and the business-unit benefit commitments — signed by the units, inside their own targets.

Every two quarters: models, vendors, build-versus-buy, unit costs, the live use case list. This layer is reported, never committed. It is expected to be wrong, which is why it must never be load-bearing for anything above it.

Most AI sections I read have all three clocks tangled into one, which is why they can be simultaneously over-committed and under-specified: five-year certainty about a vendor, five-year vagueness about who decides.

Where to start

Before the next draft goes into the template, mark every sentence in the AI section P or D. Anything perishable that carries a commitment — a number, a date, a name — moves out of the five-year plan and into the operating cycle. It doesn't get deleted. It gets demoted to the clock it actually runs on.

Then read what's left, and take the result seriously. If the durable half is too thin to fund, you've found the plan's real problem while it's still a draft: you were planning to buy technology, when the thing that needed a five-year commitment was changing what your organization is allowed to decide, and who is allowed to decide it. That is the change work that is the project rather than the postscript, and it is the only part of this that a better model will never do for you.

If you want a read on whether the durable half is there at all, I built a free AI Readiness Score — 20 questions, about ten minutes, across pilots, data, talent and governance. It scores decision rights and ownership explicitly, because those are the lines in a five-year plan that are still true in year five.

Read the essay online at https://burakgaliba.com/blog/five-year-ai-plan