BURAK GALİBA
thoughts on AI

ESSAY · ENTERPRISE AI

Whose Number Is It? Why IT Should Never Own the AI Value Line

When the AI benefit line in the strategy plan belongs to IT, adoption stalls and finance can't audit the claim. The value is only real when business units own the number.

ROI | AI STRATEGY | LEADERSHIP

Burak Galiba · August 6, 2026 · 5 min read
https://burakgaliba.com/blog/who-owns-the-ai-value-number

SCAN TO READ ONLINE

Strategy season does something predictable to enterprise AI programs. The planning template goes out, every function drafts its commitments, and somewhere in the technology section a slide appears: the AI program, presented by IT, with a cumulative benefit number on it. So many million over the planning horizon. The room nods.

That nod is the problem. Not because the number is wrong — it may be scrupulously built, case by case, from real projects. The problem is who's holding it. An AI benefit number presented by IT is a forecast of *other people's* P&L improvements, committed by the one function that can't deliver a single euro of it alone.

I have sympathy for how this happens, because the sequence is almost forced. Someone has to start the program, and it's usually a central team. The central team needs a value story to defend its budget, so it aggregates the portfolio into one impressive line. The business units are quietly happy to let someone else carry the commitment — enthusiasm without accountability is a comfortable position. Three reasonable moves, and you've arrived somewhere unreasonable: the strategy plan now contains a number that belongs to nobody who can make it true.

Three things then go wrong, on a schedule.

Accountability without authority

Every euro of AI benefit materializes inside a business process: an invoice cycle that shortens, a claims queue that needs fewer escalations, a forecast that wastes less inventory. Those processes belong to business units. The workflow changes, the staffing decisions, the adoption work that is the actual project — all of it sits outside IT's authority.

So when IT commits the number, it signs up to defend a result it cannot cause. Year one, this goes unnoticed — the plan is new and everyone applauds the ambition. Year two, the gap between committed and realized opens, and IT stands alone in the review explaining variances in processes it doesn't run. The program's credibility — not the business's — takes the damage. I've written about what applause is worth; a benefit line owned by IT is applause converted into a liability.

A number finance can't audit

Here's the test your CFO will eventually run, so you should run it first. The claimed benefits land in business-unit cost lines. Either the business units signed the baseline for those savings, or they didn't.

If they didn't, the claim is unauditable. Finance cannot reconcile "AI saved this much" against budgets that never assumed the saving — the money can't be found because nothing was baselined before the work started. The benefit line is a story, and finance treats unpriced stories exactly as they deserve.

If they did sign the baseline — good. But then the saving is already inside the business unit's target, and IT's separate line counts it a second time. A benefit that appears both in a BU's plan and on IT's slide isn't governance; it's double-counting with a logo.

That's the deletion test: remove the IT-owned AI value line from the strategy plan and see what happens to the total. If the total drops, the money existed nowhere else — no business unit had actually planned for it, which means no business unit was actually accountable for it. If the total doesn't move, the line was duplicative. There is no third outcome. A standalone AI value line fails the test in both directions, which tells you where the number has to live: inside the business targets, or nowhere.

Adoption incentives, inverted

The subtlest cost is behavioral. A business unit that owns its AI number *pulls*: it clears workflow blockers, staffs the process changes, pushes for the rollout, because its own target now depends on the result. A business unit watching IT own the number *hosts*: it politely accommodates the project, attends the steering meetings, and waits to see if the central team's bet pays off.

The portfolio looks identical on a slide. The outcomes are not close. The MIT NANDA finding that 95% of enterprise GenAI pilots show no measurable P&L return is, at root, an ownership statistic — and centralizing the value commitment in IT manufactures exactly the unowned condition that number is measuring. You can't buy pull with a bigger central target. You can only transfer the target to the people whose pull you need.

What IT should commit instead

None of this means IT shows up to strategy planning empty-handed. It means IT commits the things IT actually controls — the capability, not the benefit:

- **Time from approved idea to production.** The cycle time of the delivery machine itself.
- **Unit cost per use case**, trending down as platform and patterns get reused.
- **Reuse rate** — the share of new use cases built mostly from existing components, which is the honest measure of whether you have a platform or a pilot factory.
- **Governance service levels**: how fast a use case clears risk review, how completely the portfolio meets the compliance dates that now have teeth.

These are commitments IT can defend in any review without borrowing anyone's P&L. And they make the division of labor legible: IT runs a delivery capability with delivery metrics; business units commit euros against baselines they signed, with named owners. The strategy plan then shows AI value where value belongs — inside the operating targets of the units that will earn it — and shows IT's contribution as what it truly is: the capability that makes those targets reachable.

The handover doesn't have to be a cliff. In the first cycle, co-sign: joint targets, IT and the business unit both named, with attribution rules agreed before the work starts instead of litigated after. By the second cycle, the euro belongs to the business and IT's name stays only on the capability metrics. If a business unit refuses to co-sign even that first cycle, that's not a process failure — that's the single most useful signal in your portfolio. It's telling you which "committed" benefits were never believed by the people who would have to deliver them.

Where to start

Take the largest single benefit in your current AI value line and ask its receiving business unit to co-sign it — named owner, agreed baseline, the saving written into their next budget. One conversation. If it ends in a signature, you've converted your best number from a story into a plan. If it doesn't, you've learned your number was fiction while there's still time to fix it quietly.

For a wider view of whether your operating model puts AI value where it can actually be earned, I built a free AI Readiness Score — 20 questions, about ten minutes, and the ownership dimension is scored explicitly, because in my experience it predicts more than the technology ever does.

Read the essay online at https://burakgaliba.com/blog/who-owns-the-ai-value-number

I built a free AI Readiness Score — 20 questions across pilots, data, talent and governance, about ten minutes, no email required: https://burakgaliba.com/readiness

© 2026 Burak Galiba · https://burakgaliba.com · Views are my own.