BURAK GALİBA
Thoughts on AI Transformation

ESSAY · ENTERPRISE AI

5 Mistakes Enterprises Make With AI — the Same Five, Year After Year

Five mistakes that show up in enterprise AI portfolios again and again — unowned pilots, vendor-written strategy, governance procrastination — and the fix for each.

AI STRATEGY · PILOTS · LEADERSHIP

Burak Galiba · May 12, 2026 · 4 min read
https://burakgaliba.com/blog/5-mistakes-enterprises-make-with-ai

SCAN TO READ ONLINE

I've watched enterprise AI portfolios up close for years. The industries rotate — financial services one quarter, energy the next — but the mistakes barely change.

Here are the five I see most often, roughly ordered by how much they cost. None of them are technology mistakes. All of them are fixable this quarter.

The five mistakes

1. Running pilots nobody owns

Ask who owns a given pilot and you'll usually get a committee, an "AI task force," or an enthusiastic director three levels below budget authority. That's not ownership — that's adult supervision.

A pilot without a named P&L owner has no one whose number improves if it works and no one embarrassed if it doesn't. It will run forever, politely, producing decks. The MIT NANDA finding that 95% of GenAI pilots deliver no measurable P&L return is, at root, an ownership statistic.

The fix: no pilot gets funded without a named executive owner, one business metric, and a baseline value for that metric. Three fields on a form. Brutal, cheap, effective.

2. Funding the demo, not the return

Budgets flow to the use cases that present well: sales copilots, marketing content engines, a chatbot for the annual report. Meanwhile that same report found the strongest returns sit in back-office automation — claims, invoices, document intake, reconciliations — where the work is repetitive and the cost line already has an owner.

The fix: rank your use-case backlog by measurability and volume, not by demo appeal. If your top-funded AI initiative can't state its cost-per-transaction target, it's theater.

3. Letting vendors write the strategy

Your model vendor, your cloud provider, and your systems integrator have all offered to help you build your AI roadmap. Generous of them. Each will discover — rigorously, with frameworks — that the answer is more of their platform.

This isn't cynicism; it's structure. Model vendors now run their own consulting arms, hyperscalers embed "free" engineers who steer architecture for years, and the Big 4 rent out pre-built agent portfolios. Advice and inventory have merged.

The fix: separate the referee from the players. Whoever helps you choose the stack should earn nothing from what you choose.

4. Filing governance under "later"

Texas's TRAIGA and two California AI laws took effect on January 1. Colorado's ADMT law lands January 2027, and the EU AI Act clock is running. Meanwhile Gartner predicts over 40% of agentic AI projects will be canceled by end-2027 — inadequate risk controls being a leading cause.

Governance procrastination now fails you twice: once with regulators, once with your own projects, which get killed mid-flight when risk questions finally surface. The teams that answer those questions on day one ship faster, not slower.

The fix: stand up the basics before the next pilot starts — an AI inventory, risk tiers, and an accountable owner. TRAIGA even gives you a template: implementing the NIST AI RMF earns a statutory safe harbor.

5. Never killing anything

Every enterprise portfolio I've looked at has more pilots than it did a year ago, and roughly the same number in production: near zero for most. Nothing dies, because killing a pilot feels like admitting failure, and every pilot has a sponsor who'd rather it stayed quietly alive.

But an unkilled zombie pilot isn't neutral. It consumes budget, engineering attention, and executive patience — the exact resources your two viable use cases need to reach production.

The fix: a standing quarterly triage with three verdicts — ship, fix, or kill — and a rule that "fix" requires naming the specific blocker and its owner. If a portfolio review doesn't kill anything, it wasn't a review.

The common thread

Read the list again and notice that not one mistake involves choosing the wrong model. They're all failures of basic management applied to a new object:

- No owner (mistake 1)
- No measurable target (mistake 2)
- No independent referee (mistake 3)
- No risk accountability (mistake 4)
- No kill switch (mistake 5)

That's actually good news. Management problems don't require waiting for the next model release. They require a decision, usually yours, usually this week.

Where to start

Pick the mistake that made you wince and fix that one first — for most enterprises it's mistake 1, and the three-field funding rule takes an afternoon to institute.

If you'd rather diagnose before prescribing, I built a free AI Readiness Score that covers all five failure patterns across 20 questions — pilots, data, talent, and governance — and takes about ten minutes. You'll see exactly which of these mistakes your organization is currently making, and your first three moves to fix them.

Read the essay online at https://burakgaliba.com/blog/5-mistakes-enterprises-make-with-ai

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