BURAK GALİBA
Thoughts on AI Transformation

ESSAY · ENTERPRISE AI

# Kill More AI Pilots. Seriously.

Gartner says 40% of agentic AI projects will be canceled by 2027. Good. Killing pilots early and on purpose is how winners fund what ships.

PILOTS ROI LEADERSHIP


**Burak Galiba** · July 14, 2026 · 4 min read

https://burakgaliba.com/blog/kill-your-ai-pilots

SCAN TO READ ONLINE

Gartner predicts that over 40% of agentic AI projects will be canceled by the end of 2027. The prediction made headlines again just last week, and it's universally reported as bad news.

Here's the contrarian read: the 40% isn't the problem. The *timing* is. Those projects deserved to die — the tragedy is that they'll die in 2027, after burning two years of budget, instead of in month three, after burning a memo.

Your AI portfolio doesn't have a killing problem because you kill too much. It has a killing problem because you kill too late, too expensively, and never on purpose.

## Cancellation and killing are not the same thing

A **cancellation** is what Gartner is forecasting: a project that scaled on hope, met reality — costs, unclear value, failed risk review — and was terminated by exhaustion. Maximum spend, zero learning, and a lingering organizational flinch about the next AI proposal.

A **kill** is a decision. Early, cheap, documented: this pilot doesn't meet the criteria we set, here's what we learned, the budget moves to the candidate that does. Same verdict, a tenth of the cost, and the organization gets *faster* instead of gun-shy.

The enterprises stuck in the famous 95% — the pilots that never produce measurable P&L impact — are mostly not failing to pick winners. They're failing to unpick losers. Every zombie pilot consumes the budget, engineering attention, and executive patience that the two viable use cases needed.

## Why nobody kills anything

The forces protecting zombie pilots are predictable and human:

- **Sunk cost with a badge.** "We've already invested $400K" is treated as a reason to continue, when it's precisely the reason to decide.
- **Career attachment.** Every pilot has a sponsor, and killing the pilot reads as grading the sponsor. So it lives.
- **Demo charisma.** The pilot still demos beautifully. Demos are how zombies pass their annual physical.
- **No criteria to fail.** The deepest cause. The pilot was launched with no definition of success, so nothing it does can technically be failure.

Notice that all four evaporate if kill criteria are written down *before* the pilot starts. A pilot that begins life with "this dies unless it hits X by day 90" can be killed without a fight — the decision was pre-made, and the sponsor was never on trial.

## Portfolio thinking: kills are how you fund ships

Stop evaluating pilots one at a time and the logic flips. A portfolio of ten pilots doesn't need ten survivors; it needs one or two production systems that move a P&L line. On that math, kills aren't losses — they're the funding mechanism.

I frame this as a **Ship / Fix / Kill** triage, and the definitions matter:

- **Ship** — a named owner, a baselined metric, and a credible 90-day path to production. This is rare, which is the point. Concentrate resources here.
- **Fix** — real measured value, blocked by *one nameable thing* with an owner and a date. "Fix" without a named blocker is "kill" wearing a disguise.
- **Kill** — everything else. No owner, no baseline, no user pull, no path. In most first-time triages this is honestly half the portfolio.

A healthy triage of ten pilots typically ships two, fixes two, and kills six — and the killed six are what pay for the shipped two. If your last portfolio review killed nothing, it wasn't a review. It was a status meeting with snacks.

## How to run a kill review

The mechanics take one meeting, if the preparation is honest:

1. **Inventory on one page per pilot:** owner, metric, baseline, spend to date, blocker. No slides — slides are where charisma hides.
2. **Score against pre-agreed criteria,** not against the room's enthusiasm. Measurability, volume, owner commitment, path to production.
3. **Decide in the meeting.** Every pilot leaves with a verdict. "Let's revisit next quarter" is a kill you're paying to postpone.
4. **Make kills cheap for people.** Publish what each killed pilot taught you, credit the team for the learning, and announce where the reclaimed budget went. The first celebrated kill buys you honesty forever.
5. **Total the reclaimed number.** "We freed $1.1M and 14 engineers for the two use cases that ship" is the sentence that turns kill discipline from heresy into habit.

Then put the triage on a quarterly cadence, and write kill criteria into every new pilot's funding approval. That single rule quietly prevents the next generation of zombies from ever being born.

## Where to start

This week, ask one question of your portfolio: *which pilot would we not start today, knowing what we know?* That pilot is your first kill, and everyone in the room already knows its name.

If you want to ask it with more rigor, I built a free Ship / Fix / Kill triage that runs one pilot through 15 questions — value evidence, ownership, path to production, unit economics, reliability — and hands down a verdict. The scoring can't be charmed: your verdict is your weakest dimension, not your average, and "fix" is only granted if the blocker has a name, an owner, and a date. Fair warning: it was designed to recommend funerals.

---

Read the essay online at https://burakgaliba.com/blog/kill-your-ai-pilots

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