

The AI POC **Death** Checklist

10 Warning Signs Your AI Pilot Is Already in Trouble

Most failed AI projects don't die suddenly. They die slowly, in plain sight, while everyone pretends it's going well. Use this checklist before you've spent another six months on a POC that will never ship.

THE WARNING SIGNS — CHECK EACH ONE THAT APPLIES TO YOUR CURRENT AI INITIATIVE

01 "The AI team owns it."

Red flag: There is no named business owner — only a technology lead or vendor driving the project. Nobody from the business unit is accountable for outcomes.

Why it kills projects: Without a business owner, the success criteria drift toward technical metrics (model accuracy, latency) that don't map to any P&L line. When it's time to deploy, there's no one with authority — or skin in the game — to push it through.

THE FIX

Name a specific business leader who will be evaluated on whether this succeeds. Not "the team." One person. With a target metric.

02 Integration is a "Phase 2" item.

Red flag: The POC works beautifully — in a notebook, with sample data, disconnected from any production system. Integration with SAP/Salesforce/your data warehouse is "something we'll figure out when we're ready."

Why it kills projects: When Phase 2 starts, you discover 4–6 months of integration work that wasn't budgeted. The project quietly gets deprioritised.

THE FIX

Have someone map the production integration architecture in week one. If you can't connect it to production, you don't have a POC — you have a demo.



03

"The data isn't perfect, but it's good enough."

Red flag: Nobody has actually audited the data. Quality, completeness, and accessibility have been assumed, not verified.

Why it kills projects: Winning AI programmes spend 50–70% of timeline on data readiness. Assuming your data is ready is how you end up rebuilding a pipeline three months in.

THE FIX

Run a structured data audit before writing a single line of model code. Document what you have, what's missing, and what the compliance constraints are.



04

Change management is scheduled for "after launch."

Red flag: The people who will actually use this system haven't been consulted, communicated with, or trained. That's all planned for the go-live week.

Why it kills projects: Adoption is not a launch-day problem. By the time you go live, users need to already understand why this exists, how it helps them, and what changes in their workflow. Surprises create resistance.

THE FIX

Start change management in week one, not week ten. The pilot users should be co-designers, not recipients.



05

Success is measured in F1 scores, not money.

Red flag: Your status updates report model accuracy, precision, and recall. Nobody has defined what "success" means in revenue, cost, or time saved.

Why it kills projects: Technical success is easy to achieve and easy to ignore. A model with 94% accuracy that nobody uses has zero business value. If you can't quantify the business impact, the project has no defence against the next budget cycle.

THE FIX

Define success metrics in business terms before writing code. What does "this works" look like in the CFO's language?

06 Legal and compliance are "on the list."

Red flag: Data governance, privacy, and regulatory review haven't happened yet. They're planned — just not started.

Why it kills projects: Compliance reviews routinely uncover requirements that change the system architecture. Discovering a data residency issue in month five, after the model is built, means rebuilding from scratch.

THE FIX

Get legal and compliance in the room in week two. Not as a final gate — as an ongoing design constraint.

07 IT found out about this project last month.

Red flag: The AI initiative was scoped and started without IT involved. They've been told about it, but they haven't been designing it.

Why it kills projects: IT controls the production environment, security architecture, and integration pathways. If they're not designing with you, they'll have veto power over your deployment — and they'll use it.

THE FIX

IT needs to be a co-owner, not a gatekeeper. If they're not in the room, you're building toward a wall.

08 There is no rollback plan.

Red flag: The project plan describes what happens when it works. Nobody has documented what happens when it doesn't — or what "doesn't work" even means.

Why it kills projects: Production AI fails in subtle ways. Models drift. Edge cases appear. Without a tested rollback path and clear failure thresholds, a bad deployment becomes a crisis.

THE FIX

Define rollback criteria and test the rollback procedure before go-live. "We can revert in under two hours" should be a documented, tested claim.



09 No actual user has been in the room yet.

Red flag: The people who will use this system daily have not been consulted during design. Their workflows have been described to you by their managers.

Why it kills projects: Managers describe idealised workflows. Users live in messy reality. The gap between the two is where AI systems break and get abandoned.

THE FIX

Get 3–5 end-users involved in design and testing from week one. Their friction is your product roadmap.



10 The timeline is driven by a board presentation.

Red flag: There is a date — a quarterly review, a leadership offsite, an investor call — that is driving when this needs to be "ready." The timeline was not derived from what the project actually requires.

Why it kills projects: Artificial deadlines force shortcuts on data quality, integration, and testing. What gets presented on the board date is a demo. What gets deployed to production six months later — if anything does — is compromised from the start.

THE FIX

Build the timeline from the bottom up. Present a realistic plan with clear phases. A credible 16-week plan is more impressive — and more useful — than a rushed 6-week demo.

If You Checked 3 or More...

Your AI pilot has structural problems that won't resolve themselves. They tend to compound: a missing business owner means unclear success metrics, which means no mandate for integration work, which means a demo that never ships.

The organisations that successfully move AI from POC to production don't have better technology. They have better processes — started earlier, owned more clearly, measured more honestly.

The question isn't whether to address these issues. It's whether you address them now, when it's fixable, or six months from now, when the project is already over.

"Most companies think they're compliant because they ran one training session. The law requires considerably more than that."

Ready to move your AI pilot from demo to production?

 [Connect on WhatsApp](#)