

40% of AI agent projects are about to be cancelled, and it isn't what you think.

Tim Daines · Synaptix Solutions · October 2026

40%

of agentic AI projects will be cancelled by the end of 2027.

GARTNER · 25 JUNE 2025

Short on time? **Skip to the six-step pre-flight checklist at the end.**

Four in ten, scrapped. Quietly shut down after the money's been spent.

When people read that number, they assume they know the reason. The tech isn't ready. The models aren't good enough. The teams are out of their depth.

None of that is the real reason. I'll tell you what is.

The problem was never that they couldn't build

Here's what everyone assumes. These projects get cancelled because the people running them are out of their depth, not clever enough, not technical enough, or behind.

That's not it.

The people running these projects are often very good. They've shipped hard things before. They know how to build.

What they don't have is a map.

A technology leader showed me his AI project last month, a smart person, with a good team. They had a real budget, and six months in. It wasn't working, though, and he couldn't tell me why.

So I asked him one simple question:

"Before you built it, how did you check it was the right thing to build?"

"I didn't know there was a way to check."

That's it. That's how a project ends up in the 40%. Right there in one sentence.

Someone sold them a lie, and the lie was "the model is everything"

The lie is this: the AI model is the whole thing. Buy the model, plug it in, and you have an AI project. That's the story the whole industry has been telling, the model is the hero, the answer, is everything.

It isn't. The model is a small part, maybe five percent of what makes an AI project work in a real business.

But if you've been told the model is everything, then when it doesn't work, you have nowhere to look. You stare at the model, swap it for a newer model, tweak the prompt, and it still doesn't work, because the thing that was broken was never the model.

It was everything around the model, and nobody gave them a picture of what "everything around it" even is.

// BY THE NUMBERS

What the data says about the other 95%

40%

of agentic AI projects
cancelled by end of 2027

GARTNER, JUN 2025

79%

of large US/UK enterprises
had AI cost overruns
last year

DOIT / SAPIO,
FEB 2026

80%

of companies miss AI
spend forecasts by 25%+

CLOUDZERO

COST — THE SHARPEST, MOST BUDGET-OWNING PAIN

89%

of the most mature
FinOps teams
still had AI cost
overruns, mean
overspend 30.9%

DOIT/SAPIO

15%

can calculate AI ROI
without significant
bottlenecks

DOIT/SAPIO

\$117,558

average monthly
agent spend on
inference and
orchestration

IDC, JUL 2026

10–20×

normal cost a
modelled agent
system hit in a bad /
extreme month

SYNAPTIX

VALUE & RISK — WHY THEY GET KILLED

70%+

of agentic use cases
will fail to deliver
expected value

GARTNER, 2026

~130

of thousands of
"agentic" vendors are
real — the rest is
"agent washing"

GARTNER

95%

of enterprises already
run at least one agent
workflow in
production

IDC, JUL 2026

24

public AI failure cases
reviewed, incl.
Starbucks retiring its
AI inventory tool nine
months after rollout

SYNAPTIX

No map, no frame of reference, no idea where you stand

Imagine someone hands you a car and says, "drive to the airport", but they don't give you a map, they don't tell you which airport, and you've never been to this city before.

You're not a bad driver. You just have no idea where you are or where you're going.

That's where most people sit with their AI project. They're not bad at it, they simply have no frame of reference for their own situation. No way to stand back and ask:

Where am I, really? What's strong here? What's about to break? Is this even the right thing to build?

I built the thing that answers this, and still watched good people miss it

I know this because I built the thing that answers those questions, and I still watched capable people fail to use it. Not because it was hard, because nobody had ever shown them that an assessment was a thing you could do. It had never occurred to them to check their ground before they built on it.

This is why we start with an assessment, not a model

Not a model. Not a tool. Not a platform. A simple check of where you are.

Before you spend anything else, you answer a few plain questions and you get a plain answer back:

"Is this strong enough to build on, does it need reshaping, or should you pause before it costs you more?"

PAUSE

Stop before it costs you more.

RESHAPE

Rebuild the scope around what's actually true.

BUILD

Go, on ground proven strong enough to hold it.

Pause, reshape, or build. Three words, no jargon. That's the map. That's the frame of reference. That's the thing that was missing.

It's not clever, and that's the point. The whole problem was that everything got too clever, too complicated, too dressed-up, when what a person running an AI project actually needs is the simplest thing of all: to know where they stand.

The one decision nobody told them they were allowed to make

I don't think most of that 40% get cancelled because the people are behind.

I think they get cancelled because of one decision nobody told them they were allowed to make — the decision to pause and check before they build.

Give someone that one simple step, a real map, in plain words, and most of the frustration goes away. Not because the work got easier. Because, for the first time, they could see it.

If you're running an AI project right now and some part of you suspects it's heading for that 40%, that feeling is information. Don't ignore it, check your ground.

Start there. Everything else gets simpler after that.

// PRE-FLIGHT CHECKLIST

Six questions to ask before you build, or build any further

Run these in order. Each one is a plain question with a plain, honest answer. Where the honest answer is "no" or "I don't know", that's the ground you check before you spend another pound.

● PAUSE ● RESHAPE ● BUILD

1

Is the business case real?

Can you say, in one sentence, what this project changes for the business and how you will measure it?

If you can't, you don't have a project yet, you have an interest. Pause until the sentence exists.

PAUSE

2

Is the data you need actually usable?

Can your team get to it, trust it, and work with it today, without a six-month clean-up first?

No usable data is the most common quiet killer. If the answer is no, reshape around the data you have before you buy the model.

RESHAPE

3

Will it fit the systems you already run?

Is there a realistic path for AI to live inside your current platforms and workflows without ripping them up?

If AI only works by replacing the operation around it, you're not buying AI, you're buying a transformation programme. Reshape the scope.

RESHAPE

4

Does someone own the outcome?

Is there one named person accountable for the AI decisions, the oversight, and the risk, before anything ships?

If ownership is shared, ownership is missing. Pause until a name goes next to the project.

PAUSE

5

Do you know what it costs at ten times the volume?

You've seen the pilot bill. Do you know what the same agent costs on a bad month at full volume?

If you can't put a number on the bad month, you can't defend this to the CFO when it arrives. Reshape the architecture, or price it properly.

RESHAPE

6

Can you prove it worked, or didn't?

If something goes wrong six months after launch, can you reconstruct what the system did, and why?

No audit trail is no defence. If this answer is no, don't go live. If the first five answers are yes, then build.

BUILD

Six questions. If the honest answer to any of them is "no" or "I don't know", that's where to look first, before you spend another pound on the model.

Sources

1. Gartner, *Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027* (25 June 2025), gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027. Also the source for the >70%

- value-failure and "agent washing" (~130 real vendors) figures.
2. DoiT, survey by Sapio Research (Feb 2026, 500 finance leaders, US/UK, 1,000+ staff), 79% cost overruns; 89% among mature FinOps; 30.9% mean overspend; 15% can calculate ROI without bottlenecks.
 3. IDC (Jul 2026), 95% of enterprises run at least one agent workflow in production; \$117,558 average monthly agent spend among those with cost visibility.
 4. CloudZero, 80% of companies miss AI spend forecasts by 25%+.
 5. Synaptix / The AI Agent Architect, *Why Enterprise AI Projects Fail: Three Decisions That Matter*, theagentarchitect.substack.com/p/240eb7c6-88d6-43d4-93b3-241522beafa0. Source for the 5% / 95% framing, the 24 cases, the Starbucks nine-month retirement and the 10x/20x agent-cost model.

The DoiT/Sapio, IDC and CloudZero figures are from vendor-published research; re-check against the primary release before use in print.



Synaptix Solutions

Applied AI, minus the hype. From prototype to production.

info@synaptix.ae · synaptix.ae