

88% of AI Agent Pilots Never Reach Production.

Here's the seven-step order that actually works — in plain English.

BUILDING YOUR FIRST AI AGENT

Why Most Agent Projects Stall

01 Why projects get cancelled

Rising costs, unclear value, and governance gaps — in that order.

02 The seven steps, in order

Skip one, and the rest of the project falls apart.

03 The one idea that matters most

Risk tiers: knowing exactly which steps need a human.

40%+

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

Rising costs, unclear value, and weak risk controls — not a bad AI model. —
Gartner, 25 June 2025

BUILDING YOUR FIRST AI AGENT

Three Reasons 88% Never Ship

64% Gaps in testing

Leaders say the agent was never tested against real, messy work.

57% Rules and approval friction

Projects stall waiting on governance sign-off that should have started earlier.

51% Not reliable enough yet

The agent itself still isn't consistent enough to trust with real work.

Forrester and Anaconda, 2026 enterprise AI agent research.

7

Steps. Same Order, Every Time.

Skip ahead — like handing an agent five tools before testing it with one — and this is exactly where projects break.

BUILDING YOUR FIRST AI AGENT

STEPS 1–4

06 / 10

Start Small, Stay Narrow

01 Pick one small task

Boring beats exciting. Small tasks are easier to test, fix, and trust.

02 Map the human steps

Write down exactly how a person does the task today, step by step.

03 Give it only a few tools

Fewer doors in means fewer things that can go wrong.

04 Add a human check on risky steps

Money, promises to customers, or private data always get a review first.

BUILDING YOUR FIRST AI AGENT

STEPS 5–7

07 / 10

Prove It, Then Scale It

05 Test with real work, not a demo

Messy, real questions beat one friendly question in a meeting room.

06 Track everything it does

Every action, every reason, every result — written down, every time.

07 Only then, add one more task

Repeat the same seven steps again. This is what the 12% do differently.

BUILDING YOUR FIRST AI AGENT

THE ONE IDEA THAT MATTERS

08 / 10

Match the Check to the Risk

RISK LEVEL	EXAMPLE TASK	HUMAN CHECK
Low	Drafting a first version of an email or summary	Reviewed before it's sent; agent works freely
Medium	Updating a customer record or scheduling a meeting	Every action logged and spot-checked

High

Moving money, changing a contract, deleting data

A person approves the action before it happens

What Good Programmes Share

01 They start small on purpose

A small agent that works beats a big one nobody trusts.

02 They set rules early, not late

Only 21% have strong governance today — the winners close that gap first.

03 They treat mistakes as data

Every error updates the task, the tools, and the human check — not just one fix.

04 They own the outcome

One named person is responsible for accuracy, cost, and safety — not just "on."

BUILDING YOUR FIRST AI AGENT

AGENTIC AI AT WORK

10 / 10

Fix the Order. Not the Ambition.

One small task, mapped properly, with the right person checking the risky parts — that's most of the work.



Terence Kok · AI & Innovation Strategist

[SAVE](#) · [COMMENT](#) · [FOLLOW](#)

BUILDING YOUR FIRST AI AGENT