

The Founder Teardown

A 20-minute self-audit for founders building agentic AI

by Sheshakanth Munagapati — Fractional CPO · ex-ServiceNow Principal PM (NowLLM, Now Assist) · 2x founder

You've got engineers and a working demo. What keeps agentic-AI products from becoming real businesses usually isn't the code — it's eight product decisions. Score yourself honestly on each. It takes about 20 minutes, and no one has to see your answers.

How to score: for each dimension, circle ● **Solid** ● **Shaky** ● **Missing**. Tally at the end.

1 Problem & fit

Are you solving a real, painful, frequent job — and is AI actually the right tool for it?

- Can you name the specific user and the exact job in one sentence, without saying "AI"?
- Does the problem clear the 4 filters: **Costly** · **Frequent** · **Persistent** · **Tolerant** (can the workflow absorb an occasional wrong answer)?
- Would a rule, a form, or a better UI solve it more cheaply? (If yes, that's a ●.)

SCORE ● **Solid** ● **Shaky** ● **Missing**

Watch-out: an impressive demo aimed at a problem nobody pays to solve is the #1 way agentic products die.

2 Data reality

Can the AI actually get what it needs?

- Is the data it depends on **accessible at the point of the workflow**, not stuck in a silo?
- Is it **current** and **owned** (someone answers for its quality)?
- Is there **enough**, and is it representative?

SCORE ● **Solid** ● **Shaky** ● **Missing**

Watch-out: MIT found weak data & integration, not model quality, is the top reason pilots fail.

3 Autonomy & trust

How much does it act on its own — and is that the right amount?

- Have you set each key action deliberately on the ladder: **assists** → **acts-on-approval** → **acts-then-notifies** → **fully autonomous**?
- Is a human in the loop everywhere an action is costly or irreversible?
- Is "what it *can do*" separated from "what it *may access*" (least privilege)?

SCORE  Solid  Shaky  Missing

Watch-out: one bad autonomous action can end user trust permanently. Most good products live at "acts-on-approval."

4 Evaluation & quality

How do you know it works when it's different every run?

- Do you have an **eval set** (real + hard cases with a way to score outputs)?
- Do you re-run it every time you change the prompt, model, or data?
- Do you have **guardrails** for when it's unsure or off-topic?

SCORE  Solid  Shaky  Missing

Watch-out: "it looked good in testing" is not a quality process. This is the single most common gap I see.

5 Roadmap & focus

Are you building the few right things, or the many tempting ones?

- Is your roadmap a set of **outcomes**, not a feature wish-list?
- Can you name the **one wedge** you're winning first?
- What did you say **no** to this month?

SCORE  Solid  Shaky  Missing

Watch-out: an agentic product that "does everything" usually does nothing well enough to pay for.

6 Go-to-market

Do you know who buys, and why they switch?

- One-sentence positioning: who it's for, what it does, why it's better?
- Do you have your first **reference logos** or design partners?
- Is your **pricing** tied to value — and does it survive your cost structure?

SCORE  Solid  Shaky  Missing

Watch-out: great agentic tech with no clear buyer is a science project, not a company.

7 Unit & token economics

Does each customer make you money?

- Do you know your **cost per use at real volume** (agents loop and multiply calls)?
- Have you modeled it at scale, not just at demo volume (production often lands 5–10× the pilot)?
- Is cost-to-serve comfortably below price?

SCORE  Solid  Shaky  Missing

Watch-out: unlike SaaS, your costs rise with usage. Growth can scale your losses.

8 Product function & team

Who owns the product decisions — and what's the plan?

- Is there one clear owner of the "what and why," or is it defaulting to the loudest voice?
- Do you have a plan (and profile) for your **first full-time PM hire**?
- Are product, design, and engineering aligned on the goal, not just the tasks?

SCORE  Solid  Shaky  Missing

Watch-out: founders scaling agentic products stall when product leadership stays a part-time afterthought.

Your tally

6+ Solid

You're in strong shape. The gains now are in sharpening focus and evals. A short advisory check-in is probably all you need.

2–4 Shaky/Missing

Normal for your stage — and exactly the set of decisions a Fractional CPO is built to fix in a few focused weeks.

5+ Shaky/Missing

The product risk is real. Address the reds before you raise or scale, not after.

What to do with this

If two or more dimensions came back Shaky or Missing, those are the calls that decide whether your product ships and scales — and they're precisely what I do as a **Fractional CPO**: I embed with your team, own the roadmap, the first PM hire, and go-to-market, and build toward a clean handoff so you're never dependent on me.

Every engagement starts with one honest call. If I'm not the right fit, I'll say so and point you somewhere better.

Book a 30-minute call → aiautomic.com/fractional-cpo

hello@aiautomic.com · aiautomic.com

*Sheshakanth Munagapati — Fractional CPO for founders building agentic AI. 21 years in product · ex-ServiceNow
Principal PM for Generative AI · took an AI product to Leader in The Forrester Wave · 2x founder.*