

Building Your First Agent

A strategic guide to finding and shipping your first agent, and not just piloting it

By Julian Lancaster, CISO & Agentic AI Practice Lead, Prowess Consulting

Most agentic AI projects don't ship. But surprisingly, the blocker is rarely the model, the platform, or the tools. It's methodology. Don't start with ambitious end-to-end automation. Start small. Focus on the friction between systems. A great first agent solves a well-known bottleneck, burden, or distraction by turning manual workarounds into reliable, repeatable actions. This guide shows you how to find the right use case and take the first step.

What a first agent actually looks like

A good first agent is narrow, well-scoped, and intentionally simple. It isn't autonomous or especially clever. Think of it as a reliable assistant that handles one repetitive task consistently, without supervision.

The gap between a well-written prompt and a functioning agent is smaller than most people realize. In many cases, an agent is simply a proven prompt connected to the right inputs and triggered automatically.

Where the best first agents live: The seams

The seams between systems are strong candidates for first agents: manual data entry, reconciliation spreadsheets, and the "export, clean, paste" workflows everyone relies on but nobody enjoys. These workarounds signal a real, measurable cost.

To find a good candidate, ask:

- Where is the same data entered more than once?
- Who maintains a spreadsheet just to connect two systems?

Follow those leads, and you'll often uncover the ideal opportunity for your first agent.

Find your first agent: The discovery process

You don't even need to author the scoping yourself. A great technique is to let your LLM interview you. The prompt below runs a structured discovery conversation. You answer in plain language, and the transcript becomes your scoping doc.

- **Answer in plain language.** Don't skip details. Make sure to include all the specific workarounds and tribal knowledge.
- **Save the transcript.** It becomes the scoping doc for the build and the baseline for your next discovery.

First Agent Discovery Prompt

```
I'm scoping my organization's first agentic AI use case. Act as an experienced agentic AI practitioner and walk me through a discovery conversation. Ask me one question at a time:  
1. What workflow do my teammates complain about most, and what workaround have they already built around it (spreadsheets, side tools, manual re-entry)?  
2. Which systems sit on either side of that workflow, and where does data get re-entered or reconciled by hand?  
3. What decision logic lives only in someone's head? (The exceptions, judgment calls, and "we always do it this way" rules).  
4. What would it cost the business if the agent got this wrong, and who could catch the mistake?  
5. Who would own responsibility for this agent's behavior after launch?  
After my answers, tell me: (a) whether this is a viable first agent or whether I should keep looking, (b) where the highest-risk unknown is, and (c) the three questions I should ask the people who actually do this work before scoping further.
```

Build it so you can trust it

The same principles that make prompts effective (clarity, completeness, and the ability to flag uncertainty) are also the core elements of making agents trustworthy. The difference is that you're operationalizing them.

Four principles matter most:

- **Always validate outputs.** Decide how you'll know the agent's results are right or wrong, and who responds and monitors outputs.
 - **Assign your agent permission after it can be trusted.** Outputs should be read-only recommendations (not actions) to start! Allow the agent to earn each new permission with evidence of stable success and move on to Human-In-The-Loop review before allowing it to act without approval.
 - **Assign an operational owner, not a technical one.** The best owner is someone who directly benefits from the agent's success, as they'll be invested in keeping it effective over time.
 - **Work with your technology team.** IT teams already know how to govern automation, so make them a trusted partner to help ensure agents fit securely within your organization's systems and processes.
-

Use the discovery process to scale agents across your organization

Scaling the agentic workstyle across the organization requires cultural adoption. To start the cultural change, simply run the discovery process across 5–10 different workflows with different teams, adapting it to your industry, risk tolerance, and business needs.

Then formalize the discovery process into an onboarding process for new ideas: require new agent ideas to go through discovery, creating a consistent pipeline for identifying, evaluating, and implementing opportunities quickly across the organization.

Agents scale successfully in organizations that drive cultural adoption. Using the discovery process can turn scattered enthusiasm into repeatable gains.

From the field

A tech enterprise's compliance team was running eleven spreadsheets to track exceptions across four systems. We didn't replace the systems; we put an agent in the seams to reconcile them. **Cycle time was cut by more than half, and the spreadsheets were retired.**

Where to go from here

→ **Talk to our team:** Brendon Shaw, brendon.shaw@em.prowessconsulting.com | [Book a call](#)

→ **More from Julian:** The Road to Agentic AI video series | [Watch now](#)

→ **Register for our next webinar:** Live with BDPA CEO | [Register now](#)