

Before you hire an agent: six steps

From the talk **From Tools to Teammates**, by Alexandrine Fayet (WWF International) and Samuel Harper (CURA).

We know how to hire a person: a job description, an interview, a manager, a budget line. There was no process to hire an agent, so we started writing one. These are the six steps we learned while preparing our first agents at WWF. None of them is technology.

HOW TO USE IT Go through the six steps before an agent joins a team. Tick what you have. Where you cannot tick, you have found your next piece of work, and it is good to know that early. Reuse it, adapt it, share it.

01 Write the house rules

Before anyone joins, you write the house rules. This takes longer than you expect, and it is what lets you begin.

- ☐ A risk-tiered AI policy
- ☐ A framework for explainability, accountability and transparency
- ☐ A human in the lead where decisions matter
- ☐ A full audit trail
- ☐ Education first: people feel confident to question an answer, not only allowed to use a tool

02 Ask if you should build

One front door for every request, and five questions before anything is built.

-
- ☐ What problem, for whom? No tool names allowed.
 - ☐ What does it cost today: time, money, pressure?
 - ☐ Would a better process, a clearer rule, an existing tool, or no AI at all solve it?
 - ☐ Will building, running and overseeing it cost less than doing it by hand?
 - ☐ Who owns it, and what one measure shows in 30 days if it worked?
-

In one live review, about half of our requests did not need us to build anything. They needed a licence, a setting, or a governance decision.
Sometimes the best agent is a better process.

03 Write the job description

You hire an agent the way you hire a person: with a good job description. For an agent, it is a plain markdown file you can read, test and audit.

-
- ☐ The process is written down first, step by step
 - ☐ The job, in one sentence
 - ☐ The process the agent follows
 - ☐ When it must ask a human
 - ☐ What it logs, for audit
-

The first questions will be about your organisation, not about models. The rules are real, but they often live in people's heads, and **an agent cannot read a head.**

04 Assign a line manager : Name who owns the decision

For people, shared accountability is normal. For an agent, it is a wall: it needs one name.

- ☐ One named owner for each decision in the process
- ☐ The risky decisions are flagged, and a human approves them
- ☐ A rule for when the owner does not answer: how many reminders, then who decides

If you cannot name the person, you are not ready yet. When humans lead, the human can become the bottleneck, so plan for it.

05 Book the busiest people first

The people an agent will help most are the busiest people in the building, and the ones who know the process.

- ☐ First conversation, about trust: why you want to help, and why it is safe
- ☐ Second conversation, about the work: one three-hour block to walk the whole process together
- ☐ Their time is in the plan before the build is
- ☐ A baseline is measured before anything is built

If AI mainly helps the people who already have time to experiment, we have solved the wrong problem.

06 Register every agent

Most organisations do not know how many agents they already have. The register is the org chart of your hybrid team.

- ☐ Where it runs
- ☐ Who owns it
- ☐ What it can touch
- ☐ A review date: should it still exist?

No orphan agents.

Hiring is only the start

If agents are part of your team, they need what a team needs. After hiring comes the rest of the life cycle: onboarding, performance reviews, development, and one day, retirement.

One assumption to keep throughout: AI does not automatically save time. **It moves work.** Measure where it goes, including where the saving does not appear.

Learned something we did not? Tell us.

Your dead ends will save someone else the trip.

CURA

Samuel Harper

CURA

sam@cura-im.com



Alexandrine Boggiano Fayet

WWF International

aboggiano@wwfint.org