

BUILDER · 6 MINUTE READ

Multi-agent environments: a team of specialist AIs

Several AI agents, each with one clear job, can handle complex work better than one generalist, if a person or orchestrator coordinates and checks them.

What it is

An AI agent is an assistant that can carry out a task in steps: it reads, decides what to do next, uses tools and reports back. A multi-agent environment is a set-up where several agents work on the same piece of work, and each one has a single job.

Each agent gets instructions, the tools it may use and the information it needs. One might research, another write, a third review. Something has to coordinate them. That can be a person, or an orchestrator, which is an agent whose only job is to split the work, pass it to the right specialist and collect the results.

Think of a small newsroom. One person gathers facts, another writes, an editor checks, and a senior editor decides what goes out. A multi-agent set-up copies that pattern.

Why it matters

- A short, focused set of instructions is easier for an AI to follow than one long list that tries to cover everything.
- A separate reviewer agent looks at the work with fresh eyes, so it can catch some mistakes the writer agent missed. It will not catch them all.
- Each agent only gets the tools it needs, which limits the damage if something goes wrong.
- The costs are real: more agents means more AI usage, more time and more places where errors can creep in.

How to do it

1. List the steps of the work, as if briefing new colleagues.
2. Group the steps into jobs that need different skills, such as finding facts, drafting and checking.
3. Give each agent one job, a short set of instructions, a clear input and a clear output.
4. Limit each agent to the tools and documents its job needs, and nothing more.
5. Decide who coordinates: a person at first, an orchestrator agent only once the flow is stable.
6. Put a person's check before anything leaves the team: sent, published, paid or changed in a real system.
7. Set limits on rounds and time, so agents cannot pass work back and forth for ever.
8. Test on made-up examples, compare with what one assistant gives, and keep the simpler option if it is as good.

Example

Imagine a team that writes a monthly customer update. Before, one person asked one assistant to "research our product changes and write the update". The draft mixed up two releases and nobody spotted it.

After, the team used two agents and one person. The writer agent drafts as before. A reviewer agent checks the draft with these instructions:

```
You are a reviewer. You do not rewrite. Check this draft against the source notes below.
List every claim that is not supported by the notes, every date or number that differs,
and anything unclear for a customer. Give a list, most serious first.
Draft: [paste draft]
Source notes: [paste release notes]
```

The person reads the reviewer's list, fixes the draft and approves it before it is sent.

Common mistakes

- **Vague job descriptions that overlap.** Give each agent one job and say what it must not do.

- **Errors that pile up.** A small mistake early on gets built upon by every later agent. Check the output of early steps, not only the final result.
- **Endless loops.** Two agents can keep sending work back to each other. Set a maximum number of rounds, then hand it to a person.
- **No one watching cost.** Each agent uses AI capacity, which may be paid per use. Track usage while testing.

Try it this week

1. Pick a document you write often, such as a short report or a reply template.
2. Ask your approved assistant to write a draft, using made-up details.
3. Open a new conversation and paste the reviewer instructions from the example, with the draft and your source notes.
4. Compare the reviewer's list with what you would have spotted yourself.
5. Decide whether the reviewer step is worth keeping.

Going further

- Start with a reviewer agent: it is often a good first step, because it only reads and comments.
- Some assistants, including Claude, ChatGPT, Microsoft Copilot and Gemini, let you save custom assistants with their own instructions; one per specialist role is a simple way to begin.
- Give reviewer agents a checklist from your own quality rules, not a general "check this".
- Do not use multi-agent set-ups for simple, one-step tasks, for work that needs one consistent voice throughout, or where you cannot afford the extra cost and checking time.

Key words

- **Agent:** an AI assistant that carries out a task in steps and can use tools.
- **Specialist agent:** an agent with one job, short instructions and limited tools.
- **Orchestrator:** an agent or person that splits the work, assigns it and gathers the results.
- **Reviewer agent:** an agent that only checks work and reports problems, without changing it.
- **Loop:** agents passing work back and forth without finishing.