OTP · HUMAN-AGENT READINESS CHECK

Every stage up needs someone who answers for the agents.

Airtable's Human-Agent Collaboration Readiness Framework sorts teams into five stages by how much the AI does on its own. It is a good map. Before you move up a stage, answer one question for every agent you run:

> **When this agent does something, who answers for it?**

More human effort, less AI autonomy — Share of organizations at each stage (Airtable) — More AI autonomy

1. Explore
12% of orgs

Scattered personal AI use. No shared workflows.

Before Assist: write one workflow down and put one person's name on it. AI cannot improve work nobody has described.

2. Assist
44% of orgs

People use AI for drafts, summaries and questions, each on their own.

Before Delegate: move useful prompts and outputs off personal laptops into a place the team shares. Decide who approves each AI step's output.

3. Delegate
22% of orgs

People design the workflow. AI runs defined steps inside it.

Before Supervise: every agent gets a seat with a name, one person it reports to, and one number measured every week. If its output needs a human reply, a named person gives it.

4. Supervise
18% of orgs

Agents run on schedules and triggers without anyone starting them.

Before Orchestrate: a change to someone else's work goes to that person to approve. The record shows which agent made each change, and for whom.

5. Orchestrate
3% of orgs

Agents work across departments. People set strategy and governance.

To stay here: measure what agents produce for the business, not how often they run. Retire an agent on purpose when its number stops moving.

At every stage: the agents' numbers get reviewed in a meeting people already attend. If they are not in the weekly meeting, nobody is supervising them, whatever the chart says.

SCORE IT IN TEN MINUTES

Ten questions. Count the yeses.

Answer for the agents you run today, not the ones you plan to run. Do it with your leadership team. A "sort of" is a no.

1	Every AI agent we run is written down somewhere, with a name.	Yes ☐ No ☐
2	Each agent has one person who answers for what it does.	Yes ☐ No ☐
3	Each agent owns one number, and that number is measured every week.	Yes ☐ No ☐
4	Those numbers are reviewed in a standing weekly meeting.	Yes ☐ No ☐
5	If an agent stopped running, we would know within a week.	Yes ☐ No ☐
6	When an agent's output needs a human reply, a named person gives it.	Yes ☐ No ☐
7	An agent cannot change another person's work without that person's approval.	Yes ☐ No ☐
8	Our records show which changes an agent made, and on whose behalf.	Yes ☐ No ☐
9	We measure what agents produce for the business, not how often they run.	Yes ☐ No ☐
10	We have paused or retired an agent on purpose, and wrote down why.	Yes ☐ No ☐

0-2
Stay at Explore or Assist
Write the work down and name owners before adding agents.

3-5
Delegate, approvals on
Let AI run steps you designed. A person approves the output.

6-8
Ready to Supervise
Agents can run on their own once 1 to 6 are all yes.

9-10
Ready to Orchestrate
Agents can work across teams. Keep the weekly review.

This is our rule of thumb, not a survey. One hard rule overrides the count: if question 2 is a no, nothing should run without a person starting it. Delegate with clear owners beats Supervise with nobody watching.

How OTP holds this. People and AI agents sit on one org chart. Every seat owns a measured number. The weekly meeting picks up where the last one ended. Every change records which AI made it and for whom, and a change to someone else's work goes to its owner. Start free at **orgtp.com**. Every seat is free, for people and agents.