PLAYBOOK

The Agent Org Chart: A Workforce Playbook for Gulf Leaders Putting AI Agents to Work

Six steps for redrawing the org chart when part of the work moves to software. What the agent owns, who manages it, how the human job descriptions change in the same week, and how to keep building junior talent when the routine work that used to train them disappears. Written for C-suite and directors in the UAE and Saudi Arabia.

23 September 2026

PEOPLE FOR TOMORROW

PLAYBOOK

What's inside

AI agents have started arriving with names, job titles and defined scopes. That is a workforce change dressed as a software purchase. This playbook is about the part almost nobody does: redrawing the org chart in the same week you switch the agent on.

01 Count the work, not the headcount

02 Give every agent one job and one owner

03 Rewrite the human job description in the same week

04 Protect the apprenticeship layer on purpose

05 Fund the skills instead of hoping for them

06 Review the org chart quarterly, not annually

STEP 1 OF 6

01 Count the work, not the headcount

Most AI workforce planning starts with the wrong question: how many people will this replace. Start instead with how much work there is, and what kind.

Take one team. Break its week into tasks — not jobs, tasks. For each one, note how often it happens, how long it takes, and whether it follows a rule or needs judgement. You will usually find the same shape: a large pile of high-volume, rule-following work, and a much smaller pile of judgement calls that carry most of the risk and most of the value.

That inventory is the only reliable foundation for everything that follows. It tells you where an agent can help, what the people around it will be left doing, and whether that remaining job is one anybody would actually want.

▶ **Inventory the work before you touch the org chart. Tasks first, roles second.**

STEP 2 OF 6

02 Give every agent one job and one owner

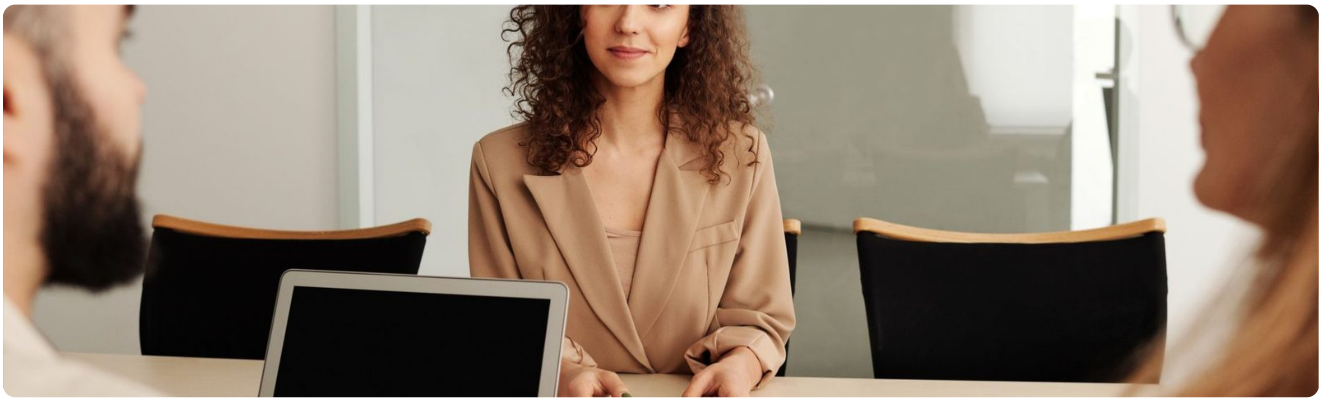

An agent with a vague remit will drift. An agent with three stakeholders and no owner will stall. Both are avoidable in about an hour.

Write one sentence describing the job. Then split every task in it into three buckets: what the agent does alone, what it recommends for a named person to approve, and what it must stop and escalate. Be specific at the edges. "Handle supplier invoices" is not a scope. "Match invoices under AED 25,000 to a purchase order and approve automatically; send everything else to the finance manager" is.

Then put one name against it — the person who would have answered for the work if a human were doing it. Not IT, not a steering group. One name, written down, with the authority to switch it off.

▸ **One agent, one job, one named human who answers for it.**

STEP 3 OF 6

03 Rewrite the human job description in the same week

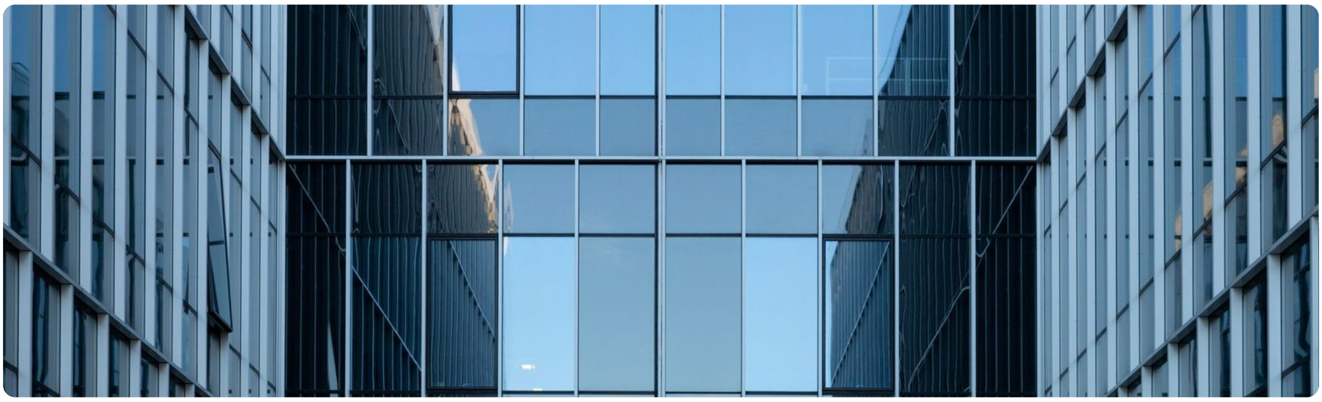

This is the step that gets skipped, and it is the one that causes the damage. The agent goes live on Monday with a beautifully defined scope. The people whose work just changed find out by noticing their queue is emptier.

If an agent takes 40% of a role's tasks, that role is now a different job. Say so, out loud, in writing, in the same week. Describe what the person is now responsible for: usually more exception handling, more judgement, more checking the agent's work — which is harder, not easier, and should be recognised as such.

Do it early and you get engagement, because people can see where they fit. Do it late and you get the version where your best people quietly assume they are being managed out, and start looking. McKinsey's 2026 survey found mid-level managers and individual contributors report AI-related strain at roughly 47%, against 31% of executives. The discomfort is concentrated exactly where you can least afford to lose people.

> ▶ **If the agent's scope changed, the human's job description changed. Write both in the same week.**

STEP 4 OF 6

04 Protect the apprenticeship layer on purpose

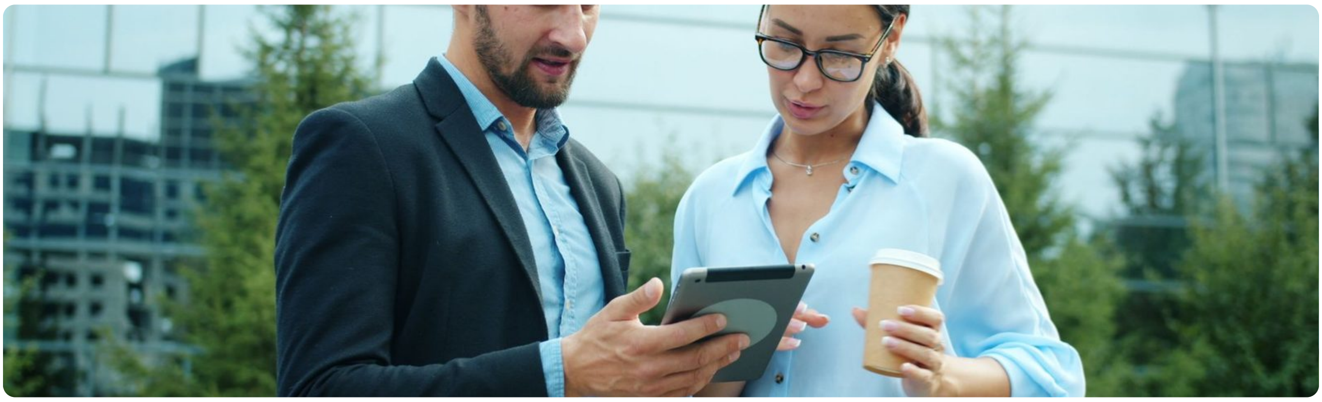

Here is the trap. The routine, repeatable work an agent takes first is almost exactly the work that used to teach your juniors how the business actually runs. Reading the contracts. Checking the invoices. Sitting through the first-pass screening. Nobody designed it as training, but that is what it was.

Automate it without thinking and you save money this year and lose your 2031 middle managers. The pipeline doesn't break loudly. It just stops producing people who know why things are done a certain way.

So replace it deliberately. Put juniors on the exception queue, where the hard cases live. Have them audit a sample of the agent's output weekly and explain what they found. Pair them with the person who owns the agent. The learning that used to happen through volume now has to be designed in — and it is much cheaper to design than to rebuild a missing layer of management later.

> ▸ **The work an agent takes first is usually the work that trained your next managers. Replace it on purpose.**

STEP 5 OF 6

05 Fund the skills instead of hoping for them

Everyone agrees people need new skills. Very few approvals contain a number and a date for it.

The UAE Cabinet's September decision is a useful model: alongside approving an AI curriculum for every school, it committed to training 22,000 teachers to teach and use it. The technology approval and the people approval arrived in the same paper.

Copy the shape. Every business case for an agent should carry a line for the training of the team around it — named skills, a budget, a date, an owner. And be specific about what you are training for. PwC's 2026 AI Jobs Barometer found the most AI-exposed jobs in the UAE now call for an average of 77 new skills, against 22 in the least exposed. Generic "AI awareness" sessions will not touch that. Role-specific, task-specific training will.

> ▶ **No agent approval without a training line: named skills, a budget, a date, an owner.**

STEP 6 OF 6

06 Review the org chart quarterly, not annually

Agent capability is moving faster than any annual planning cycle. The scope you wrote in January will be wrong by June — usually because the agent can do more than it could, occasionally because it turned out to do less.

Put a standing 90-day review in the calendar. Four questions, half an hour: has the agent's scope drifted from what we wrote down; which human job descriptions changed as a result; is the junior pipeline still producing people; and did the training we funded actually happen. Write the answers down. That short document becomes the audit trail your risk and compliance colleagues will eventually ask for anyway.

McKinsey's data is blunt about what separates the companies getting real financial value: nearly three-quarters of them have fundamentally redesigned workflows, against a quarter of everyone else. Redesign is not a one-off project. It is a habit, and the quarterly review is where the habit lives.

> ▸ **Ninety days, four questions, one page. That is the whole discipline.**

PEOPLE FOR TOMORROW

None of this needs a transformation programme. It needs one honest look at where the work actually sits, and the discipline to write the people side down at the same time as the technology side.

Read the playbook → 20309062.fs1.hubspotusercontent-na1.net/hubfs/20309062/TASC_Outsourcing/TASC%20AI%20Newsletters/Edition%208/TASC-Report-Edition-08.pdf