PLAYBOOK

The Leader's Playbook for Putting AI Agents to Work

Every leader is being told to "deploy agents." Few are told how to do it without landing in the 40% of projects that get cancelled. This six-step playbook walks C-suite, Directors and VPs in UAE and KSA enterprises from a first agent to governed, scaled, value-generating agentic AI: where to start, how to redesign the work, how to build governance and human accountability in from day one, how to meet regional expectations like sovereign and Arabic-first AI, and how to close the AI talent and supervision gap without out-bidding the whole market.

29 July 2026

PEOPLE FOR TOMORROW

PLAYBOOK

What's inside

AI agents are moving into real work fast — the UAE is running government services on them and global AI spend is up 63% this year. But Gartner expects more than 40% of agentic AI projects to be cancelled by 2027, almost always over cost, unclear value or weak controls, not the technology. These six steps take a leader from a first agent to governed, scaled, value-generating AI — built for UAE and KSA enterprises moving fast while the rules and expectations tighten.

01

STEP 1 OF 6

Start where the value is

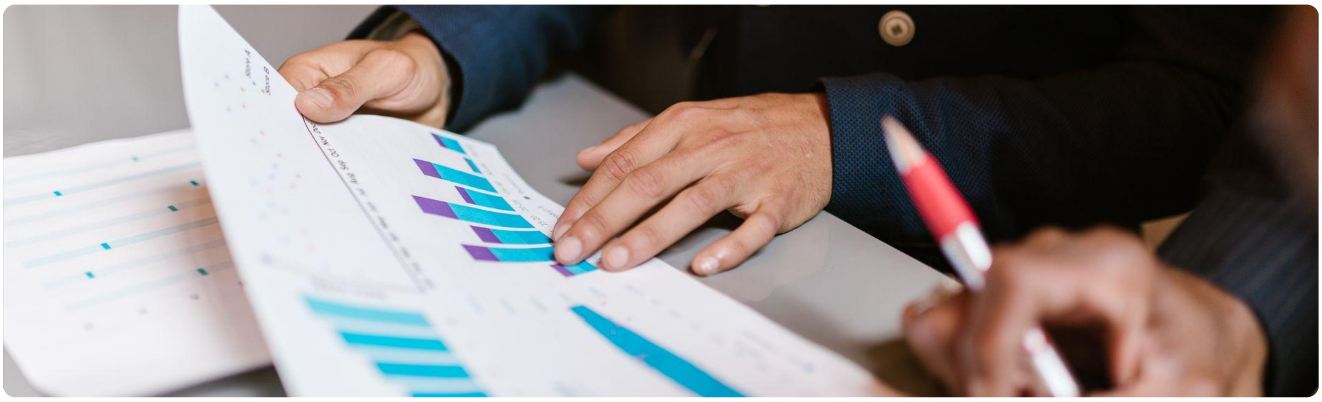

Don't start with the hardest, highest-profile use case. Start where there's high volume, lower risk and a clear measure of success — document processing, customer service triage, first-pass screening, routine approvals. Prove the value there, then move up.

That discipline matters because spending is easy and value isn't. Global AI spending is up 63% this year, yet Gartner still expects more than 40% of agentic AI projects to be scrapped by 2027 — usually because no one pinned down the value before scaling.

Pick use cases with a number attached: time saved, errors caught, cost per case. Those numbers earn you the mandate and budget for the next step, and they protect you from scaling something that only looked good in a demo.

▸ **Prove agents on high-volume, lower-risk work with a number attached. 40%+ of agent projects fail for lack of clear value — don't be one.**

STEP 2 OF 6

02 Redesign the work around the agent

The most common mistake is bolting an agent onto a process you never changed. You get a faster version of the old way, and most of the value leaks out.

BCG's 10-20-70 rule is the reminder here: only about 30% of the effort in an AI project is the technology and data; roughly 70% is people, process and change. The return comes from rethinking the workflow — what the agent does, where it hands off, and where a person steps in — not from the model itself.

Map the whole task before you deploy. Decide which steps the agent owns end to end, which it drafts for a human, and which stay human entirely. That design work is unglamorous and it is where the payoff lives.

> ▸ **About 70% of AI value is people and process, not tech (BCG). Redesign the work — don't just bolt an agent onto the old way.**

STEP 3 OF 6

03 Build governance in from day one

Governance is not a thing you add after launch. In the UAE, the direction is already set: Ajman is running a live government service on an agent "within a framework of governance, trust, accountability and data protection," and the Cabinet is training 80,000 staff to supervise agents as it moves half of government services onto them.

Build the controls as you build the use case. That means a named owner for each agent, an audit trail of what it did and why, clear limits on what it's allowed to do, and bias checks before it goes live.

The sharpest test is consumer-facing: can a customer get a human review or explanation of a decision an agent made? Design that path now. Governance built in is far cheaper than governance bolted on after a complaint.

> ▸ **The UAE is running live government services on agents — with governance built in. Match that bar: named owner, audit trail, human review.**

STEP 4 OF 6

04 Keep a human accountable for every decision that matters

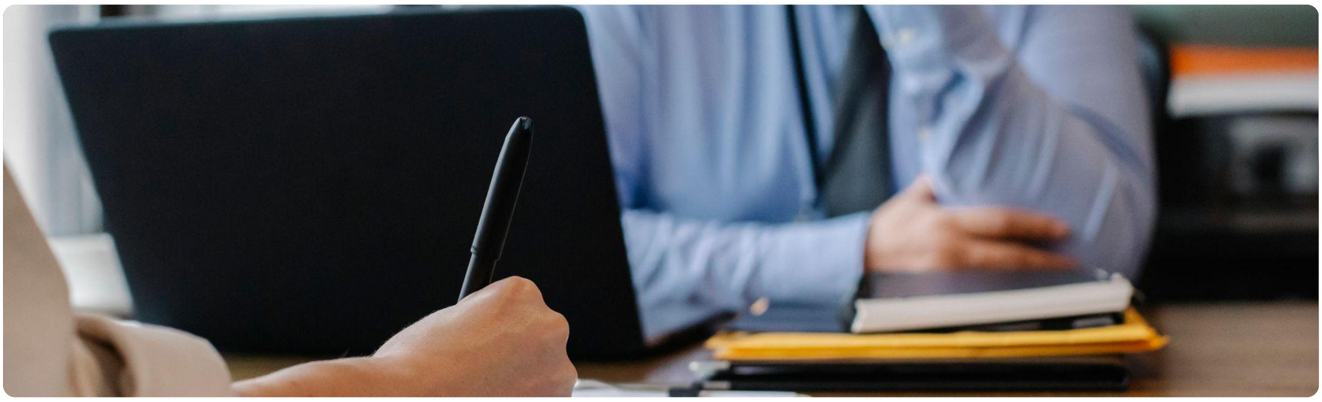

Agentic AI can draft, rank, flag and act — but someone still has to be accountable for the decisions that affect a customer's money, a person's job or the business's risk. Design the workflow so the agent does the heavy lifting and a person owns the call.

Be deliberate about where the human sits. Routine, low-risk steps can run with light oversight. Credit, fraud, hiring, claims and advice need a person who can see what the agent did and override it.

This is also how you keep trust. The teams getting value from agents aren't the ones who removed people — they're the ones who put people where their judgement counts most and let the agent clear everything else.

> ▸ **Let agents do the work; keep a person accountable for the call. Light oversight for routine steps, a human owner for anything consequential.**

05

STEP 5 OF 6

Make it regional — sovereign, Arabic, compliant

What works in a New York pilot doesn't automatically clear in Riyadh or Dubai. Saudi Arabia is building sovereign AI — models hosted and controlled in-Kingdom, tuned for Arabic — with HUMAIN and Cohere, and regulated work is heading that way by default.

For every agent you deploy in the region, ask three questions early: where does the data live, does it work properly in Arabic as well as English, and does it meet local rules on data protection and accountability? Answering those late is how good projects stall.

The upside is real: get the regional fit right and you can move at the speed the Gulf is setting, instead of retrofitting compliance and language after the fact.

▸ **Sovereign and Arabic-first AI is becoming the default for regulated Gulf work. Design for data residency and language from the start.**

06

STEP 6 OF 6

Win the talent race differently, then measure and scale

The binding constraint isn't the technology anymore — it's people who can run and govern agents. You won't win that on salary alone against the global market, so split the problem. Reskill the people you already have to supervise agents, hire directly only for the few roles you must own, and bring in specialists on demand for builds and peaks, then release them when the work is done.

That blend lets you move at the speed of your AI roadmap instead of the speed of your hiring pipeline — and it's a hedge, because the people you reskill and the specialists you can call on are capability you don't have to win in a bidding war every time.

Then scale with discipline. Expand a use case only once it matches or beats human quality at lower cost, and keep watching what the agent gets wrong, not just what it gets right. That's what turns a pile of pilots into a durable advantage.

> ▸ **You can't hire your way out alone. Reskill your people, own a few key roles, access specialists flexibly — then scale only what you can prove.**

PEOPLE FOR TOMORROW

The advantage in agentic AI won't go to whoever buys the most technology. It will go to whoever turns it into governed, measurable value first — and that's a people and process problem as much as a technology one. Pair every AI plan with a workforce plan: reskill the people you have, keep a human accountable for every decision that matters, and bring in the specialists the whole market is chasing on a flexible basis. Start small, govern from the top, and scale only what you can prove.

Read the playbook → www.tascoutsourcing.com/insights/tasc-ai-news/edition-04-playbook