



PLAYBOOK

The Banking Leader's AI Playbook

Almost every bank and insurer now uses AI — but only a few turn it into real value, and the Gulf is short of the people to scale it safely. This six-step playbook walks bank and insurance leaders in the UAE and KSA from AI pilots to governed, scaled, value-generating AI: where to start, how to clear the data and legacy hurdles, how to meet the UAE Central Bank's governance bar, and how to win the talent race without out-bidding everyone. Written for the C-suite, Directors and VPs running BFSI in the region.

1 July 2026

PEOPLE FOR TOMORROW

What's inside

AI is now in 81% of financial firms, but only 40% have scaled it — and McKinsey says about \$170 billion of global banking profit is in play by 2030 for those who get it right. In the Gulf the banks lead the region yet employ a fraction of the AI staff of global peers. The technology isn't the hard part anymore; deploying it with governance, on real systems, with the right people is. These six steps take a bank or insurer from pilots to scaled, governed AI — built for UAE and KSA institutions moving fast under a tightening rulebook.

01 Start where the value is

02 Fix the data and legacy foundation

03 Build governance in from day one

04 Keep a human on every consequential decision

05 Win the talent race differently

06 Measure quality and economics, then scale

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 — fraud detection, customer service, document processing, claims triage. Prove value there, then move up.

This is also where most banks already are. The Cambridge 2026 report found AI deployment is concentrated in internal operations, because that's where the data is

cleanest and the risk is lowest. Use that to your advantage: bank a few clear wins before you take AI near credit decisions or customer-facing advice.

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 looks good in a demo.

- ▶ **Prove AI on high-volume, lower-risk work first. Only 40% of finance firms have scaled AI — clear wins earn the mandate for the rest.**

02

STEP 2 OF 6

Fix the data and legacy foundation

The most common reason AI projects stall in banking isn't the model — it's the plumbing. The Cambridge report names poor data quality and legacy systems among the top barriers to scaling AI, alongside the talent shortage.

You don't need to replace your core systems before you start. But you do need clean, accessible, well-governed data for the use cases you've chosen, and a realistic plan for

how AI connects to the systems you already run. Treat data quality as a deliverable, not an assumption.

This is unglamorous work, and it's why fintechs — built on modern systems — are three times more likely than incumbents to be 'transforming' with AI. The banks that close that gap are the ones that fund the foundation, not just the pilot.

- ▶ **Poor data and legacy systems stall more AI projects than weak models. Fund the foundation, not just the pilot.**

03

STEP 3 OF 6

Build governance in from day one

In the UAE, governance is no longer optional. The Central Bank's guidance requires every licensed financial institution to run a documented AI governance framework, with the board and senior management accountable for AI outcomes — and it's part of a global wave of similar rules.

Build the controls as you build the use case, not after. That means a named owner for AI risk at board level, audit trails of which model made which decision, protection

against discriminatory or manipulative outcomes, and bias testing before a model goes live.

The sharpest requirement is consumer-facing: a customer can ask for a human review or explanation of an AI decision, and an alternative if they don't want an AI decision at all. Design that path now. Governance built in is far cheaper than governance bolted on after a complaint.

- ▶ **Governance is the UAE regulator's baseline, not a nice-to-have. Name an owner, log decisions, and give customers a human to reach.**

04

STEP 4 OF 6

Keep a human on every consequential decision

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

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 value

comes from redesigning the work around the agent, with clear hand-offs, not bolting an agent onto an unchanged process.

Be deliberate about where the human sits. Routine, low-risk steps can run with light oversight; credit, fraud, claims and advice need a person who can see what the agent did and override it. That design is what keeps you fast and safe at the same time.

- ▶ **Let agents do the work; keep a person accountable for the call. About 70% of AI value is people and process, not tech (BCG).**

05

STEP 5 OF 6

Win the talent race differently

The binding constraint in Gulf banking isn't technology – it's people. Banks in the region employ an average of 316 AI staff against 1,763 at global peers, and the scarcest of all are the people who bridge technology, risk, regulation and the business.

You won't win that race on salary alone against New York and London. So split the problem. Reskill the risk, compliance and operations people you already have to supervise and govern agents. Hire directly

only for the few roles you must own outright. And for everything else – builds, launches, peaks – bring in specialists on demand, compliantly, then release them when the work is done.

This blend lets you move at the speed of your AI roadmap instead of the speed of your hiring pipeline. It's also a hedge: the talent you reskill and the specialists you can call on are capability you don't have to win in a bidding war every time.

- ▶ **You can't hire your way out alone. Reskill your people, own a few key roles, and access specialists flexibly for the rest.**

06

STEP 6 OF 6

Measure quality and economics, then scale

It's easy to celebrate speed and quietly lose quality. Track outcome measures alongside efficiency — decision quality next to time saved, error and complaint rates next to throughput — and review what the agent gets wrong, not only what it gets right.

Keep the economics in view too. McKinsey's warning is that about \$170 billion of global banking profit is at risk by 2030 for institutions that don't adapt, while early

movers could lift returns by up to four points. Scaling the wrong thing fast is how you end up on the wrong side of that line.

Set a clear bar: expand a use case only once it matches or beats human quality at lower cost. Then scale it the way you started — as a workforce and governance question as much as a technology one. That discipline is what turns AI from a pile of pilots into a durable advantage.

- ▶ **Scale only what beats human quality at lower cost. \$170bn of bank profit is in play — don't scale a problem.**



PEOPLE FOR TOMORROW

AI advantage in BFSI won't go to whoever buys the most technology. It will go to whoever turns it into governed, measurable value first — and that is a people and process problem as much as a technology one. Pair every AI plan with a workforce plan: reskill the staff you have, secure flexible access to the specialists the whole market is chasing, and keep a human accountable for every consequential decision. Move early, govern

**from the top, and the window
Evident says closes around 2030
is yours to take.**

Read the playbook → [www.tascoutsourcing.com/insights/\[REPLACE-WITH-PLAYBOOK-SLUG\]](http://www.tascoutsourcing.com/insights/[REPLACE-WITH-PLAYBOOK-SLUG])