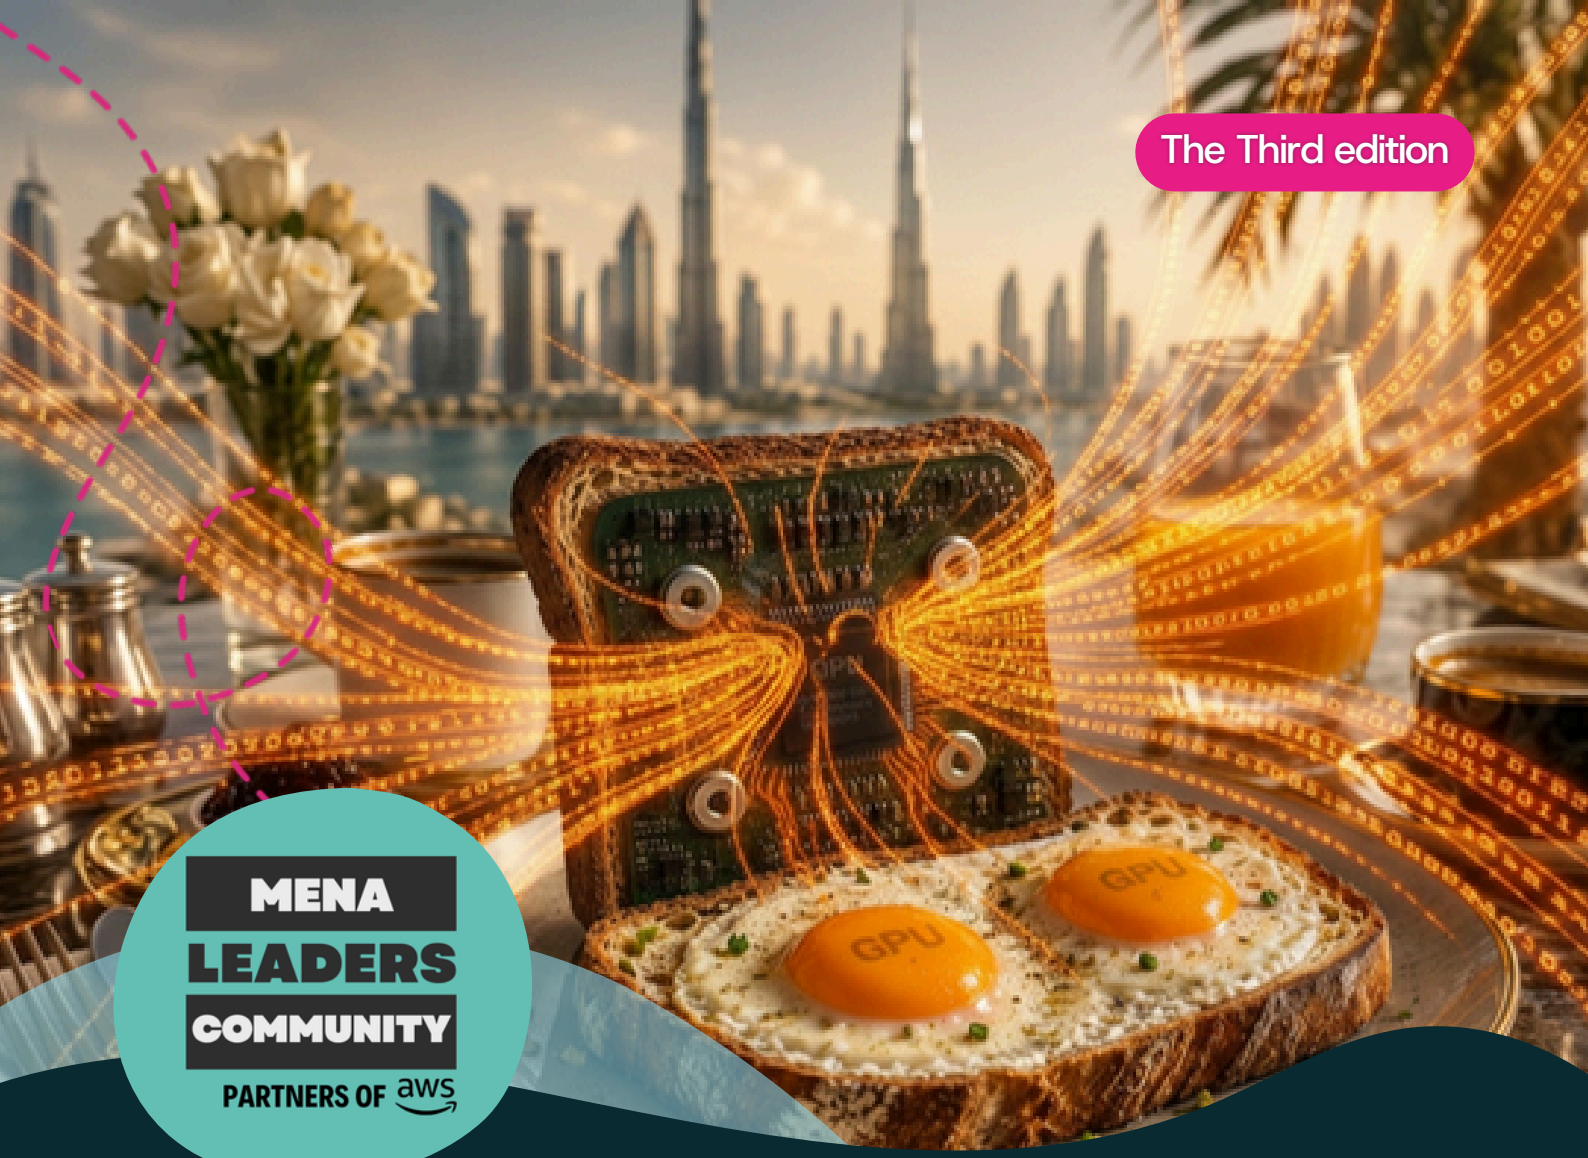


The Third edition

MENA LEADERS COMMUNITY

PARTNERS OF aws

AI. Served Fresh.

The agentic enterprise

What founders, CTOs, product and engineering leaders said when we asked: who's actually doing the work?



Co-hosted by

Tiago Henriques, Google MENA

October 2026 | AI breakfast report

salt

The one-line version

The doing is being automated. The judging isn't.

In July, the question was whether anyone trusted AI's output. By October, nobody was asking. The question was who signs it off.

Four tables started from four different questions: the work, the organisation, trust and the economics. They all ended up in the same place. Execution, from research and code to first drafts and support tickets, is moving to agents faster than anyone predicted.

What isn't moving is judgement: the taste to know what good looks like, the authority to decide what is irreversible, and the accountability to own the mistake. That is now the scarce resource in the agentic enterprise, and most organisations haven't yet redesigned themselves around it.

By the numbers

80 → 8
One startup's non-voice support team after automating its interactions. The eight who remain "review the review"

7 → 3
"Pizza team" size cited for BMW as it moves to smaller, faster AI-native teams.

48h
Creative turnarounds clients now expect, after AI collapsed production from weeks to hours.

99%
Accuracy that is still not enough in high-stakes settings like lending

Inside this report

01 The morning and what we asked
02 From doing to judging
03 The AI-native organisation
04 Who owns the mistake?
05 Cheaper isn't the same as better
06 Creative work and the authenticity premium
07 Where the room disagreed
08 What leaders are taking away
09 The question has changed

The morning and what we asked

Four questions, four tables, one room.

On Friday 9 October, founders, CTOs, CPOs and engineering and product leaders joined the MENA Leaders Community and Salt for the third AI. Served Fresh. breakfast, co-hosted by Tiago Henriques of Google MENA.

09:00
Networking breakfast

09:45
Tiago Henriques on the agentic enterprise

10:00
Four roundtables

11:00
Presenting back to the room

On the table

The work

From copilots to agents to digital workers: what are we actually letting AI do on its own

Trust

How much autonomy is enough and how much is too much?

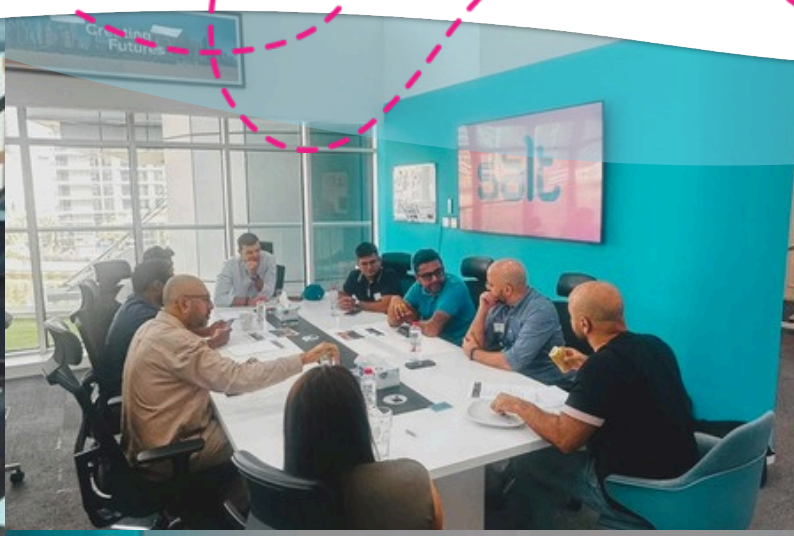
The organisation

If one person with five agents can do the work of a team, what happens to the company?

The economics

Show us one workflow where AI has genuinely changed the economics of your business







The work

From doing to judging

The conversation moved from AI as a tool to AI as a layer of the workforce. Several people compared it to Excel: something that stops being a specialist skill and becomes the default way work gets done.

Across coding, research, marketing and customer operations, the group broadly agreed that execution should be close to 100% automated: research, first drafts, initial code, prototyping and market testing. Speed and cost make the case on their own.

What changes is the human job. People move from doing the work to giving it strategic direction and taste, and to judging what comes back. In one customer support operation, the people who remain spend their time assessing AI-handled transcripts and outcomes rather than answering tickets.

The constraint is no longer how much work a team can do. It's how much good judgement it can apply to what the agents produce.

The debate it sparked

Taste first

Humans stay the guiding owners

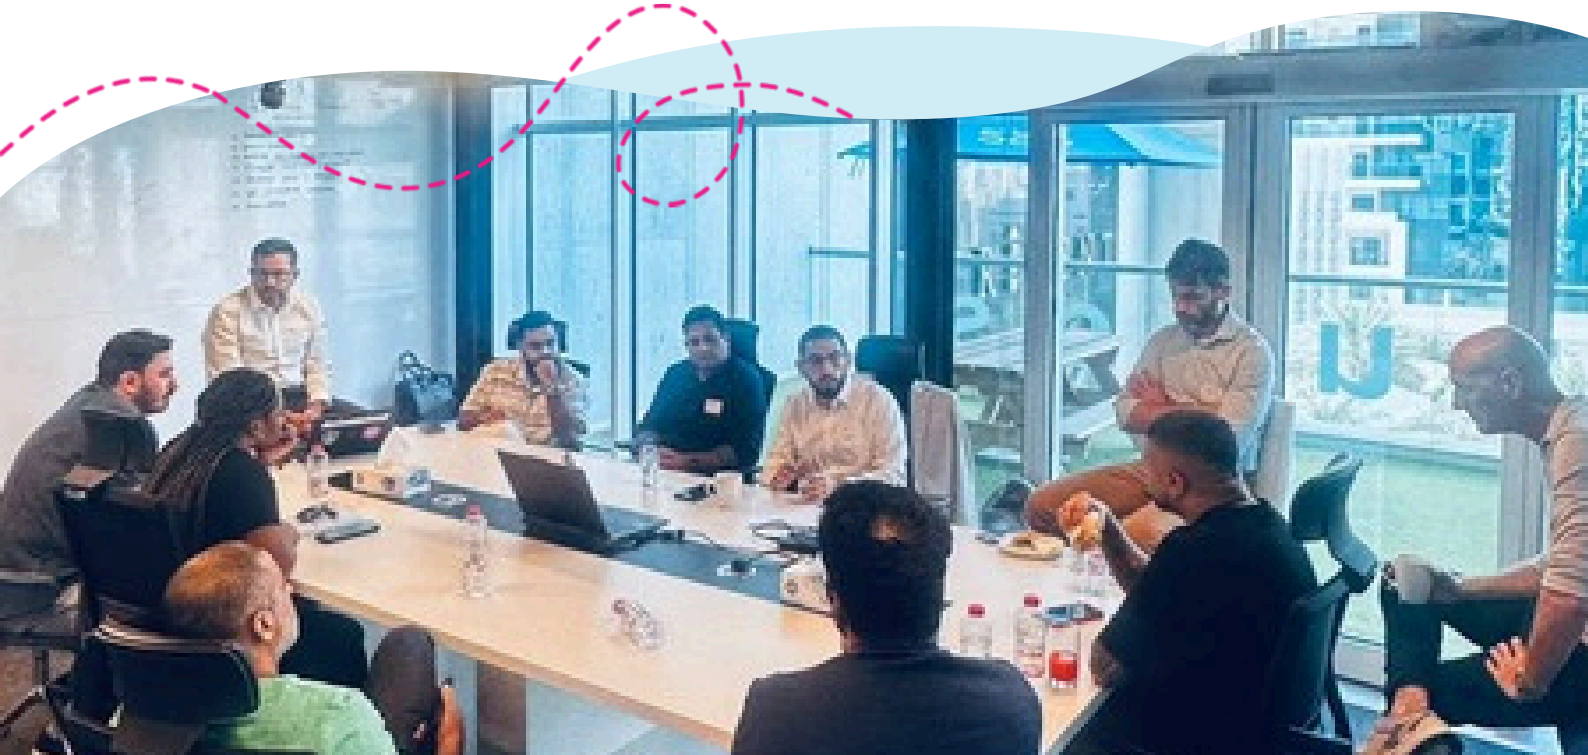
Agentic models lack taste and architectural foresight. Without a human owner, outputs drift away from what is scalable and high quality. Someone has to know what good looks like.

Done over perfect

Speed beats polish

Messy code is a human staple too. If an agent reaches the objective faster, internal sloppiness matters less, and the speed of newer models lets teams iterate past the human bottleneck.

Key takeaway **Automate execution by default, and invest in the people who can judge it. Taste is becoming a core job skill, not a nice-to-have.**



The organisation

The AI-native organisation is smaller, flatter and blurrier

If one person with five agents can do the work of a team, the team changes shape. The room had live examples.

80 → 8

One startup automated 100% of its non-voice support interactions. The remaining 8 handle exceptions and "review the review".

7 → 3

BMW was cited as moving from 7-person "pizza teams" to 3-person teams to gain agility

Role boundaries are collapsing. In AI-native teams, engineers now pick up design, infrastructure and product research, because AI handles the baseline technical execution. Smaller, multi-disciplinary teams are producing what used to need several specialists.

Middle management is most exposed. Its core functions, emailing, synthesising data and creating content, were seen as the most vulnerable to near-term replacement. Several leaders said they are already planning for smaller, AI-native leadership layers.

Juniors still matter, but the role must change. This was the closest the room came to consensus: nobody wanted to eradicate junior roles. The task is to redesign them so they offer purpose and a path to leadership, and so the next generation learns to manage AI systems rather than compete with them.

Three pillars of an AI-native team

01

Efficiency

Do more with less. Spend less time on items and more time on context.

02

Expanded boundaries

With the doing solved, everyone needs the fundamentals, like an engineer who understands customer pain.

03

Process adoption

The bottleneck is no longer the amount of work but how fast workflows and culture can adapt.

Key takeaway

The hard part isn't the technology. It's redesigning roles, layers and career paths around it.

Trust

Who owns the mistake?

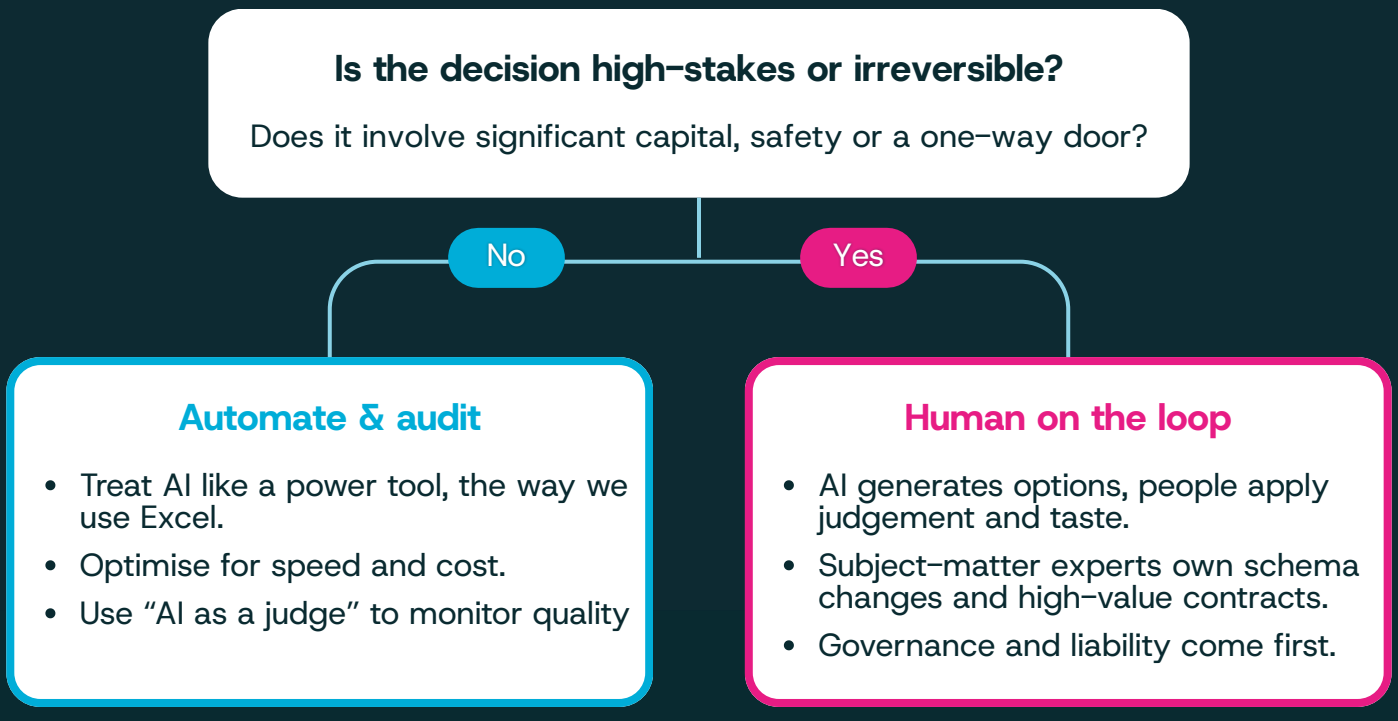
Every table found the same split: “bulls” ready to automate almost everything, and a more cautious group worried about liability.

Nobody knows who owns the mistake. This was the biggest open question of the morning. With no legal precedent yet for AI-driven failures, most organisations are defaulting to strict governance and risk management to avoid data breaches and unforced errors. Security risks in agent-to-agent interactions added another layer nobody has fully solved.

Trust depends on the stakes. Confidence is high for low-stakes work like meeting minutes, research and itineraries. It drops sharply for high-stakes or irreversible decisions such as investing money or analysing medical records. The useful test was whether a task is a “one-way door”: once it’s done, can you undo it?

Evals are the bottleneck. Validating chains of agents is still hard. The direction of travel is “AI as a judge”, with a stronger model checking agent output against quality benchmarks. But in lending, for example, 99% accuracy isn’t enough. There, evals need to be deterministic or overseen by a subject-matter expert who directs the tools rather than doing the work by hand.

The room’s automation triage



Key takeaway **Decide who owns the mistake before an agent makes one.**

The economics

Cheaper isn't the same as better

We asked every table for one workflow where AI had genuinely changed the economics of the business. The answers challenged the idea of AI as a simple cost-saver.

Cheaper can mean worse

There was real concern that cost-cutting through AI, such as automated phone systems, is degrading customer experience. Done badly, reactive service stops being a solution and becomes a barrier between a company and its customers.

The better version is the opposite: using AI to shift customer service from reactive to proactive, and judging success by customer satisfaction, not just speed.

The savings are visible. The costs are hidden.

Short-term FTE savings are easy to point to. But they are often offset by new costs: token and infrastructure spend, and the quieter loss of institutional context when the people who held it are gone.

Nobody can isolate AI revenue

Most organisations have no metric for how much revenue AI actually generates. Much of today's adoption is reactive, doing the "same thing faster", rather than the proactive innovation that creates new value.

Faster isn't automatically better. If you only measure headcount saved, you'll miss both the hidden costs and the real upside.

What to measure instead

Full cost

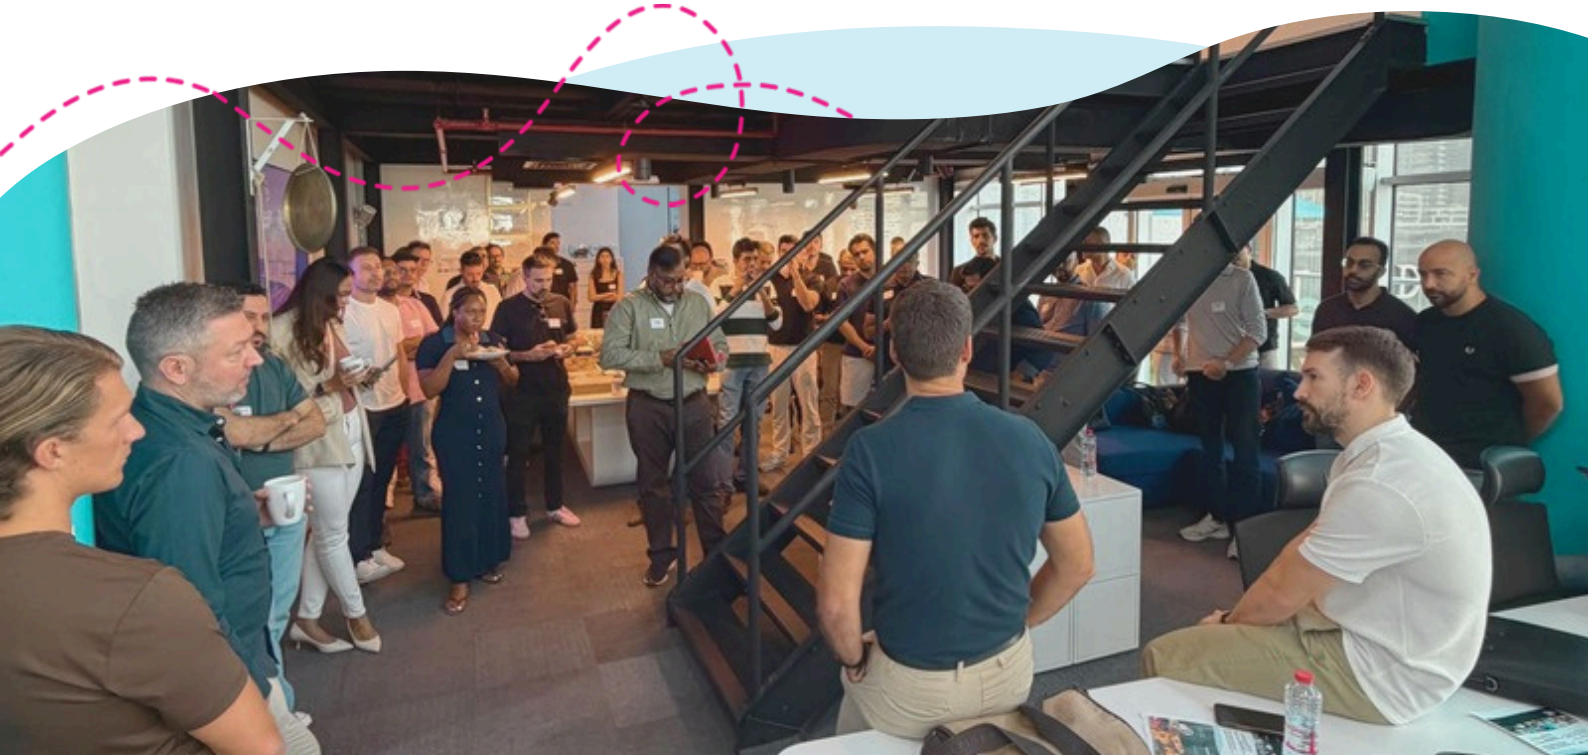
Balance short-term labour savings against long-term token, infrastructure and context costs.

AI-driven revenue

Build metrics that isolate what AI actually adds, not just what it replaces.

Customer outcomes

Put satisfaction gains ahead of pure operational speed.



Creative work

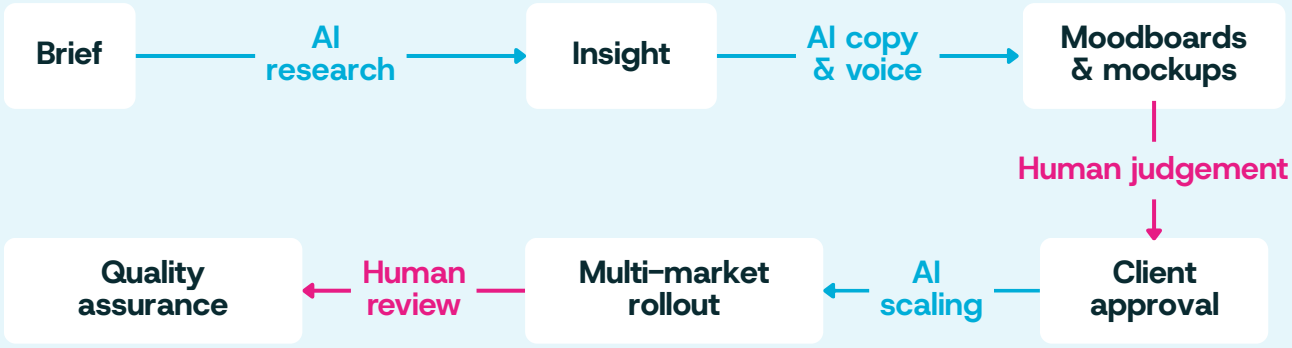
Production is cheap. Authenticity is the premium.

Advertising gave the clearest picture of what happens to a whole industry when execution gets cheap.

Timelines have collapsed. Production that took weeks now takes hours. Teams present four creative options where they used to present two, and clients now expect 48-hour turnarounds. The flip side: as production gets more efficient, clients are reducing retainers.

Everything is starting to look the same. As AI-generated content becomes ubiquitous, consumer taste is swinging back to human imperfection. Raw, unedited mobile video is outperforming polished work, and Gen Z's unfiltered aesthetic is part of the same shift. Human insight is becoming the differentiator.

How AI now runs through a creative commission



Blue hand-offs are AI-led.
Pink hand-offs stay with people, who own client approval and final quality.

Key takeaway **When anyone can produce, the value moves to the insight behind the work and the judgement applied to it.**



Fault lines

Where the room disagreed

Not everything converged. These were the live arguments nobody fully settled.

Taste or speed?

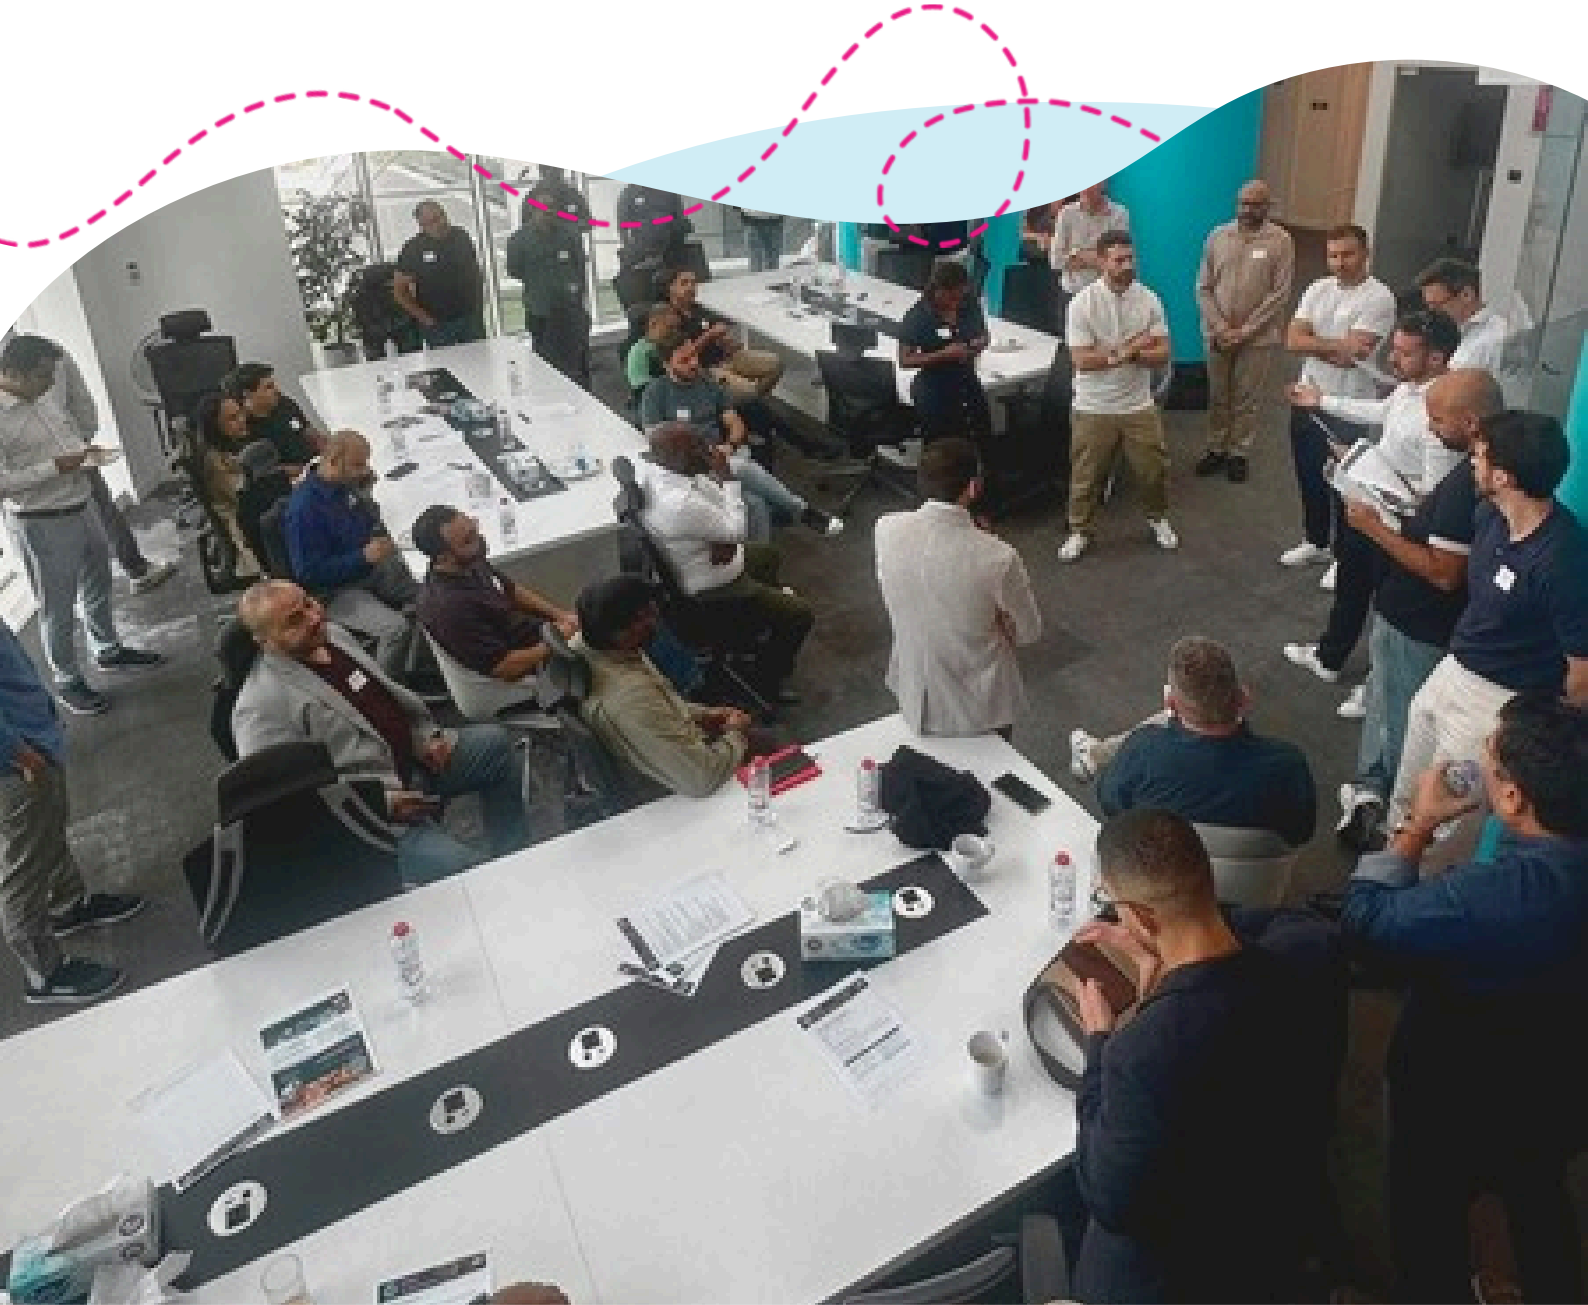
Do humans need to stay the “guiding owners” of quality and architecture, or is done-over-perfect good enough when agents iterate faster than people can review?

Bulls or sceptics?

Some are ready to automate almost everything now. Others won’t move on high-stakes work until someone can answer who owns the mistake.

Cost-saver or experience-killer?

Is AI-driven cost-cutting creating leverage, or quietly degrading customer experience in ways the numbers don’t yet show?



Moving forward

What leaders are taking away

Core positions

Work & workflow

Default to automating execution, and name the human who owns the judgement.

Sort every workflow by reversibility before deciding how much autonomy it gets.

Use AI as a judge for low-stakes chains. Keep subject-matter experts on high-stakes ones.

Trust & governance

Decide who owns the mistake before an agent makes one.

Make evals deterministic where accuracy has to be near-perfect.

Treat agent-to-agent interactions as a security surface, not an afterthought.

People & economics

Redesign junior roles as a path to leading AI systems, not a casualty of them.

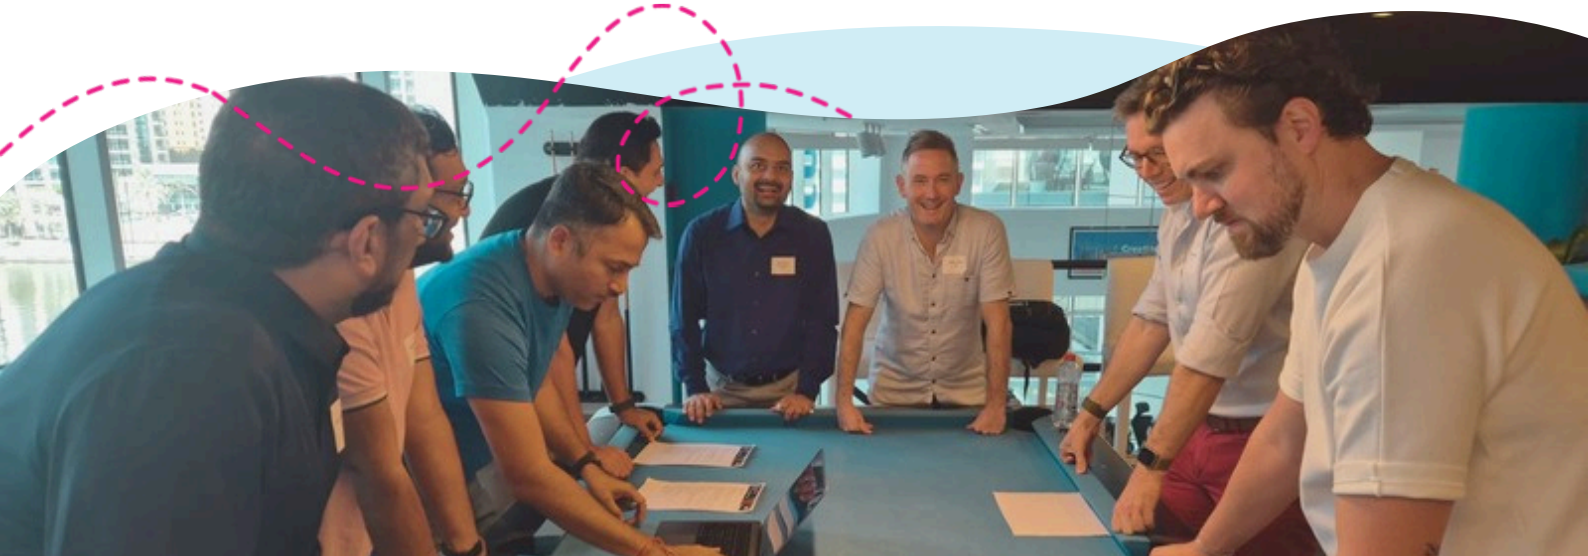
Measure full cost, AI-driven revenue and customer satisfaction, not just headcount saved.

Invest in process and culture change: that is now the real constraint.

The question has changed

In July, the room was still asking whether AI's output could be trusted. By October, the question had become what we're prepared to hand over, who stays accountable, and what the organisation looks like when the doing is no longer the hard part.

Execution is becoming a commodity.
Judgement, taste and ownership are not.





AI. Served Fresh.

See you at the fourth edition.

Each breakfast is small by design: a room of founders, CTOs, CPOs and engineering and product leaders talking candidly about what is really working. Places go fast. Follow Salt and the MENA Leaders Community to hear first when the next date is announced.

If judgement is the scarce resource, it's also the hire.

Salt specialises in placing AI and digital transformation leaders across the Middle East, Asia Pacific, Africa and Europe. We work with organisations at every stage, from first AI hire to scaling a global team.

salt
Creating Futures

MENA LEADERS COMMUNITY
PARTNERS OF aws

welovesalt.com/insights

© Salt Solutions Group.
All rights reserved.
Branding and design by Salt Marketing.