

# MERIDIAN

THE SYNARCHY QUARTERLY



ISSUE #01 · Q3 2026

## The Judgement Layer

Agentic AI and the decisions  
no institution can rent

A government in the region has set itself two years to run half of its government sectors, services and operations through AI agents. The private sector tends to follow; it has before. The question almost no board has answered: when an agent acts for your institution, whose judgement is it carrying out?

# IN THIS ISSUE

|     |                                                         |    |
|-----|---------------------------------------------------------|----|
|     | From the Principal                                      |    |
|     | Why Meridian exists                                     | 03 |
|     | In Brief                                                |    |
|     | The shift, and the question beneath it                  | 04 |
| I   | The Countdown                                           |    |
|     | A two-year clock, and what follows it                   | 05 |
| II  | What Changed                                            |    |
|     | From advice to action                                   | 06 |
| III | Own or Rent                                             |    |
|     | A way to sort every decision an agent touches           | 06 |
|     | • <i>The Grid in Practice - Three decisions, sorted</i> | 09 |
| IV  | The Stakes                                              |    |
|     | What mis-sorting costs                                  | 10 |
| V   | The Regional Picture                                    |    |
|     | Where the Gulf builds, and where it borrows             | 11 |
|     | What to Author Next                                     |    |
|     | The board's agenda for the quarter ahead                | 13 |
|     | Methodology & Sources                                   |    |
|     | How this issue was built                                | 14 |

## FROM THE PRINCIPAL

A government in our region recently set itself a deadline. Within two years, half of its sectors, services and operations should run through AI agents, software that does not advise a civil servant but acts in their place.

I read that announcement the way I suspect many of you did: first as ambition, then as a question. Not whether the region will move quickly; it will, faster than almost anyone expects. The question is quieter, and it sits underneath the speed. When an agent acts for an institution, it is exercising that institution's judgement. So whose judgement, exactly, has it been given?

That is the question this first issue of Meridian sets out to make usable. Not to forecast adoption, since the forecasts all point the same way, but to give a board a way to decide, decision by decision, what it should hand to an agent and what it must keep. The answer is not the same for every task, and the institutions that thrive will be the ones that can tell the difference early.

A fair question to settle first: our firm's promise is that we deliver results, not reports, so why publish at all? Because the institutions that decide these things well tend to think in the open before they act. Meridian is where we publish the thinking. The doing stays private, and it stays yours.

---

**Ramki Jayaraman**

Managing Partner, Synarchy Consulting



When an agent  
acts for an  
institution, it is  
exercising that  
institution's  
judgement. So  
whose  
judgement,  
exactly, has it  
been given?

## In Brief

Across the region, governments are handing real decisions to AI agents, and handing them fast. Everyone is watching the contest over who adopts first, and it is almost over. The contest that matters is quieter, and almost no one is scoring it.

The obvious story is that the Gulf will adopt agentic AI quickly. It is true, and it is nearly settled: a public mandate guarantees it, and the private sector has followed the state onto each new digital standard within a few years, in case after case. Adoption is no longer the hard question.

What the obvious story misses is the shift beneath it. A generative tool suggested; a person decided. An agent decides and acts. The moment execution passes from the person to the agent, the institution is no longer choosing whether to use a technology; it is choosing how much of its own judgement to carry out through software it did not build, on a model it did not train. That choice is being made, at speed, mostly by default, and rarely at the board.

The discomfort sits here: your institution may already be acting on judgement it never authored, efficiently, at scale, with the consequences landing on your balance sheet. This issue gives you a way to take that choice back. It offers a single instrument: a way to sort every decision an agent might touch into one of four kinds, and to see at a glance which ones you can safely rent and which you cannot afford to. The question it asks you to carry into the boardroom is the one beneath the technology: of all the decisions that make your institution what it is, which are you quietly handing away?



## The Countdown


A government has set a clock. The market will follow. The question is what the agents are about to be allowed to decide.

### One fact should be setting the clock in every boardroom in the region.

A government in the Gulf has committed to transforming 50 per cent of its government sectors, services and operations to agentic AI within two years (UAE Cabinet, 2026). Taken as an operating instruction rather than a press line, that commitment means a citizen's interaction with half the functions of a state could, inside twenty-four months, be handled not by a person, nor even by a person using software, but by software acting on the state's behalf. Whatever else it is, it is a deadline.

Deadlines in this region have a way of being met, and of pulling the private sector along behind them, as the record bears out. When the UAE launched its national digital identity, UAE Pass moved from a standing start to more than 140 integrated private entities in roughly five years (TDRA, 2025). When Saudi Arabia set its cashless ambition, electronic payments climbed to about 79 per cent of retail transactions (SAMA, cited in Arab News, 2025). In each case the state set the standard and the market followed within a few years, because the cost of staying off the new rails rose until it could not be ignored.

**EXHIBIT 1** The region has run this play before. The state sets the standard, the market follows



**Source:** TDRA (2025). UAE Pass private-sector integration rose roughly fourteen-fold from 2020 to 2025. The precedent shows reach; it is platform integration, not independent operational adoption.

Agentic adoption will propagate the same way, and the propagation is the point. That propagation is an inference from precedent rather than an observed fact, and we treat it as such. A board that treats the government mandate as a public-sector story is misreading the clock. The mandate is the starting gun; the race reaches your institution next, whether you entered it or not. That makes the unasked question urgent rather than philosophical: as the agents arrive, what exactly are you about to let them decide?


## II What changed

Not faster automation, but a change in who, or what, does the deciding.

**It is tempting to file agentic AI under ‘automation, but faster.’ That filing is the costliest category error of the moment.**

The shift is better understood as a change in what does the deciding. For a decade, software waited to be used: a person sat in front of it, made the call, and the tool carried out the keystrokes. Generative AI extended that arrangement: it could draft, summarise and suggest, but a human still pressed go. An agent removes the human from the middle of the loop. In practice the removal is a matter of degree: agents can run with approval gates, bounded authority and reversible actions, and the question for a board is how far along that spectrum each decision has been allowed to travel. Given a goal and access to systems, it plans, decides between options, and acts on the institution's behalf, at machine speed, across thousands of cases before anyone reviews a single one.

It redistributes who does the deciding rather than merely accelerating the same work, and it changes the question a leader has to answer. The old question was procurement: which tool do we buy, and at what price? The new question is closer to delegation: of the decisions that run this institution, which are we comfortable having made by an agent we did not build, and which are we not? Adoption will spread; the clock makes that difficult to avoid. The choice that actually separates institutions is not whether to adopt, but which decisions to keep.



Software you use, you buy. Software that decides for you, you must govern, or hand away. Agentic AI moves the institution from the first to the second.

## III Own or Rent

One instrument for sorting every decision an agent might touch.


**If the real question is which decisions to keep, an institution needs a way to sort them, quickly, and the same way twice. This is ours.**

Any decision an agent might touch can be placed along two axes. The first is how close it sits to what actually makes this institution distinctive: is it peripheral housekeeping, or is it near the core of the advantage? The second is how far it rests on judgement that lives in your people rather than in your data: is it a rule that can be written down, or a feel that cannot? Those two axes produce four quadrants, and each carries a different answer to whether the work should be owned or rented. **Synarchy's own-or-rent grid maps agentic decisions by proximity to advantage and intensity of judgement.**

In practice, placing a decision turns on three questions, whose answers locate it on both axes at once.

*continued overleaf...*

## EXHIBIT 2 The own-or-rent grid: where every agentic decision belongs



**Source:** Synarchy analysis. A decision moves toward 'author' as it nears the core of advantage and rests on judgement held in people, not data.



### HOW TO PLACE A DECISION: ASK THREE QUESTIONS

1. Does it draw on something our people know that is not in the data?
2. Would a wrong call compound silently, across many cases, before anyone noticed?
3. Is it close enough to our advantage that getting it wrong is not costly but existential?

Three yeses put a decision in the top-right corner, the work to own and govern, never to hand away; one or none, and it can be sourced on its merits.

The four quadrants divide as follows:

### Rent freely

Peripheral work that runs on rules, such as tier-one IT support, standard reporting and routine reconciliations, is the easy quadrant. It is not close to your advantage and it does not need your judgement. Automate it internally or buy it from whoever does it cheapest. Renting here is not a risk; it is good housekeeping.

### Hire expertise, episodically

Some work needs deep judgement but is not yours to own: a cyber-incident response, a tax-structuring question, an ERP migration. It is scarce, specialised, and called on rarely. An agent can give your experts more leverage, but the expertise itself is sensibly hired in from people who do nothing else. Owning it would be expensive vanity.

## Author the system, on your own data

Then there is work close to your core that nonetheless runs on rules: the data platforms, the payment rails, the identity and customer pipelines on which everything else depends. This is the quadrant the region already understands, because it is where the Gulf has invested most visibly, in national data banks, sovereign cloud and localisation regimes. Here the instruction is to author the system and run it on data you control, so that the agents reasoning on top of it are reasoning on your reality, not a vendor's approximation of it.

## Own the judgement

The last quadrant is the one that decides institutional fate. It holds work that is both close to the core and dependent on judgement that cannot be fully written down: the lending call on a borderline borrower, the underwriting exception, the sense of whether a long-standing client is about to leave. These decisions are the institution. They cannot be safely handed over, because to lease them out is to lease out the thing that makes you distinctive; and yet they are precisely the decisions agents are quietly being handed, because each one looks, in isolation, like a task to be sped up. Owning the judgement does not mean building every technical layer in-house; an institution may rent the model while authoring the decision rules, the accountability and the override that govern it.

The grid does not tell you where your institution sits; it gives you the axes to find out. Plotting your own decisions on it, sorting the housekeeping from the core calls and the codifiable from the tacit, is the first useful hour of work this issue asks of you. The next page shows the method in motion.



### THE INSTRUMENT

Four quadrants, one instruction: source the first three on their merits, and own the fourth, always. The only decision that can quietly sink an institution is a top-right call it mistook for a task.

## The case for renting it anyway

The grid invites an obvious objection, and it deserves a fair hearing. In work that is highly standardised and very high in volume, a rented decision layer may come to out-judge an institutional one. A provider serving hundreds of lenders sees more borrowers, more outcomes and more edge cases than any single lender ever will, and that breadth of feedback can sharpen a model in ways an in-house system, starved of comparable volume, cannot. On this reading, owning the judgement is not prudence but vanity: a refusal to let scale do what scale does best.


The objection is correct, and worth conceding plainly: across much of the grid, renting is not merely acceptable but superior, which is why three of its four quadrants point outward. It breaks down only in the top-right corner, and it breaks down for a structural reason. The decisions that sit there are valuable precisely because they are not standardised; they turn on a reading of a particular borrower, a particular client, a particular exception that the institution understands better than any aggregator of averages. Scale learns the general case well and the distinctive case poorly, and the top-right decision is the distinctive case by definition. A layer optimised across everyone converges on the median, and the median is the one thing a distinctive institution cannot afford to become.

*continued overleaf...*

## The Grid in Practice

The grid earns its place only if it discriminates, if it sorts decisions that look alike into different answers. Three illustrative cases show the method. They are examples, not prescriptions; where any of them belongs in your institution is yours to judge.

**EXHIBIT 3** The same instrument, three decisions sorted




**Source:** Synarchy analysis. Illustrative decisions, plotted to show the method, not a verdict on any institution. Each board plots its own.

**A routine invoice reconciliation** runs the three questions and answers no to each: it draws on no knowledge outside the data, a wrong line does not compound into a crisis, and it sits nowhere near the institution's advantage. This is work to rent freely or automate internally, whichever is cheaper; nothing about owning it would repay the cost.

**Building the customer-data platform** differs on one axis only. It is close to the core, since almost everything else will run on it, yet it is largely rule-based: schemas, pipelines and access controls, all written down and testable. Close to the core but codifiable, it is a system to author and run on data you control, so that whatever reasons on top of it reasons on your reality. It is also, not coincidentally, the quadrant the region has invested in most heavily.

**A borderline credit decision** answers yes three times. It leans on what a seasoned lender senses and cannot fully write down; a wrong call, repeated quietly across a segment, compounds before anyone sees the aggregate; and it sits at the heart of what the institution is for. This is a decision to own and to govern. An agent may prepare it, surface the anomalies and draft the recommendation, but a named human must remain accountable for the call. Hand it over by default, and the institution has quietly delegated a piece of itself.



Three decisions that look like 'tasks to automate' resolve to three different answers. That is the whole value of sorting before deploying: the grid keeps you from renting the one decision in the three that you could not afford to.

# IV The Stakes

What it costs to put a core decision in the wrong place.

Sorting decisions correctly sounds like an administrative nicety, and it is anything but. Putting a top-right decision in the wrong quadrant is how an institution hollows out: imperceptibly over a long stretch, and then in plain view.

The most-cited cautionary case is worth weighing precisely because it is small. A grieving passenger asked an airline’s chatbot about bereavement fares; the agent confidently described a refund policy that did not exist; a tribunal held the airline liable for what its agent had said (Moffatt v. Air Canada, 2024). The sum involved was trivial; the principle it established was not. The institution had let a system acting in its name decide what it would and would not promise a customer, and discovered that accountability does not transfer to the agent; it stays with the institution. Multiplied across an entire function, that same failure stops being trivial at all.


## Control

This is the first thing at stake: not whether agents can do the work, but how much of the deciding the institution is quietly ceding. When an agent acts on a rented model, the call is shaped, at speed and at scale, by assumptions the institution cannot inspect. The first board question is therefore not how capable the agents are, but how much of its own judgement it is prepared to let them carry.

## Capability

Authoring anything requires being able to author it, and on this the region is exposed before it is ready. Employers expect 39 per cent of core skills to change by 2030, and 63 per cent already name skills gaps as the main barrier to transformation (WEF, 2025). In the Arab States, the ILO estimates AI could augment 14.6 per cent of jobs against 2.2 per cent it could automate, a ratio that rewards only the institutions able to do the augmenting (ILO, 2025). The exposure, moreover, concentrates where it matters most for the region: in Qatar, 93 per cent of jobs held by nationals are exposed, against 33 per cent of those held by non-nationals (IMF, 2025).

EXHIBIT 4 High exposure, high disruption, and a workforce not yet built for it



Source: World Economic Forum, Future of Jobs Report (2025), country profiles. Regional digital-transformation readiness; GCC 76.3, world 62.6, Arab States 56.8 (ILO, 2025).

## Risk

Mis-sorted decisions accumulate exposure where no one is looking. The danger of a top-right decision handed to an agent is not a single dramatic error; it is the silent compounding of small ones, across a whole client segment, before any human sees the aggregate. The public record already holds the scaled-up version: in Australia, an automated debt-recovery scheme wrongly pursued hundreds of thousands of people and ended in a royal commission that called it a 'costly failure' (Royal Commission into the Robodebt Scheme, 2023); in Kenya, a rapid roll-out of automated digital credit produced average annual percentage rates of 280.5 per cent among unregulated providers before regulation caught up (Competition Authority of Kenya, 2021; Central Bank of Kenya, 2022). Neither was an AI forecast. Both were what it looks like when an institution automates a judgement-heavy decision without owning the judgement.

## Value

Finally, the work itself does not vanish; it moves. As execution shifts to agents, value collects with whoever owns the judgement and the data the agents run on. One European bank cut support-resolution times sharply even as its headcount rose (Revolut, 2024); Synarchy reads this as work migrating towards the layer that exercises and governs judgement rather than the layer that merely processes. The institution that authors that layer compounds its advantage; the one that rents it watches its margin, and its distinctiveness, drift to the owner of the model.



Renting a top-right decision reads as prudence on the income statement, flexible, low in capital and fast. It is in fact a slow transfer of the institution itself: the judgements that define you, made elsewhere, on terms you no longer set.

## V The Regional Picture

Where the Gulf builds, and where it borrows.

**Held against the grid, the region's position comes into focus, and it is more precarious than the adoption headlines suggest.**


On the quadrant concerned with authoring the systems that run on owned data, the Gulf is genuinely strong. National data banks, sovereign-cloud commitments and data-localisation regimes mean the region is building the vaults: it increasingly controls the data its agents will reason on (SDAIA, 2024). This is real, and it is an advantage most of the world cannot match. It should not be undersold.

Yet the layer above the data, the models and the compute that turn data into decisions, is overwhelmingly leased, and from a concentrated few. A small number of laboratories account for most of the world's frontier models (Epoch AI, 2025), and the supercomputing capacity to train them is similarly concentrated. The region is buying access to that capacity at remarkable scale, which is rational; but access is not authorship, and a decision layer leased from three or four providers is a decision layer whose assumptions you do not set.

There are real counter-moves, and they matter: sovereign and regional models such as Falcon (TII, Abu Dhabi, open-weight), JAIS (G42, Abu Dhabi, open-weight), Fanar (QCRI, Doha) and ALLaM (SDAIA, Riyadh, launched on Azure) are deliberate attempts to author the layer rather than lease it (TII, 2024; QCRI, 2024). They are the exception that proves the rule, and the right exception to be building. The honest reading of the region's grid, though, is a split: it is authoring its data and borrowing the judgement that runs on it. That is a strong position on one axis and a borrowed one on the other, and agentic AI is the technology that makes the borrowed axis the one that decides outcomes.


The data is increasingly owned; the models that reason on it are leased from a concentrated few.

**EXHIBIT 5A** Where models are built



**Source:** Stanford HAI (2025), model counts;

**EXHIBIT 5B** How few hands hold it



**Source:** Epoch AI (2025), top-developer share of models above  $10^{25}$  FLOP; Trends in AI Supercomputers (2025), industry compute share.

## What to Author Next

Meridian closes with an agenda rather than a playbook: the questions a board should put to itself once it has plotted its own decisions on the grid.

The board's first task is singular and clear: identify its top-right decisions, the calls close enough to advantage, and dependent enough on human judgement, that delegating them by default would change the institution itself. Everything else can be sourced on its merits: rented as commodity, hired in as expertise, or authored as system on data you own. But the core, judgement-heavy calls are the institution's to keep, and the work of the coming quarters is to find them before an agent inherits them unasked. The grid is how you find them; the questions below are how a board acts on what it finds.



### FOR THE NEXT BOARD MEETING

Which of our decisions sit in the top-right corner, close to our advantage, and resting on judgement our people hold rather than our data?

Of those, how many are already being handled, in whole or in part, by an agent, and who decided that, and when?

Where are our agents reasoning on data we own, and where on a model whose operating assumptions we cannot inspect?

Have we built the capability to govern an agent before a mandate obliges us to deploy one?

If we had to defend an agent's decision to a regulator tomorrow, could a named person account for it?

There is less time than the optimism implies. The same clock that guarantees fast adoption guarantees that mis-sorted decisions surface just as quickly. The window to decide which decisions to keep, the gap between the deployment and its consequences, is measured in quarters, not years.



For a decade the contest was access: who could reach the technology first. Agentic AI has settled it. The agents are here, they can be hired by the afternoon, and they will act on whatever judgement they are given. The only question left is whose judgement that is.

Adoption was never the advantage; everyone adopts. The advantage is the judgement you choose to own: the one thing no one can rent back to you.

## Methodology & Sources

This issue is built on primary and multilateral evidence, named independent indices, and official regulatory and corporate disclosures, assembled and tiered through Synarchy's research process. Each quantitative claim is linked to a recorded source in the references that follow; figures that could not be independently reproduced are treated as directional and identified as such.

A sourcing discipline governs the whole. The proprietary numbers of the large global strategy and analyst firms have been excluded entirely, and any statistic whose origin traced back to such a firm, even when quoted by a credible secondary source, was dropped. Independent research partners are used at theme level only. Three limitations are stated plainly. First, this issue is Gulf-led with deliberate African adjacency: the evidence base is richer for GCC institutions, where the public record is fuller, and Africa enters chiefly as a contrast rather than as an equal weight; we frame it that way rather than claim a parity the evidence would not support. Second, no Tier-1 figure exists for the size of the regional agentic-AI market, so any market sizing is treated as direction, not magnitude. Third, country-level private-sector adoption rates were not available for most of the region, so the propagation of adoption is presented as inference from precedent rather than as observed fact.

The own-or-rent grid is Synarchy's analytical instrument, not a verdict on any institution. The placement of the example decisions within it is illustrative; where a specific decision belongs depends on the institution, and each must plot its own. Referencing follows the Harvard (author–date) convention; sources are listed once, alphabetically, below. All web sources were accessed on 22 June 2026.

*continued overleaf...*

## References

- **ADGM (2024)** Big Data Analytics and Artificial Intelligence (AI): Rulebook. Abu Dhabi Global Market. Available at: [en.adgm.thomsonreuters.com](https://en.adgm.thomsonreuters.com) [Accessed 22 June 2026].
- **Abu Dhabi Media Office (2024)** G42 launches JAIS 70B and other AI models to champion Arabic natural language processing. Available at: [mediaoffice.abudhabi](https://mediaoffice.abudhabi) [Accessed 22 June 2026].
- **Arab News (2025)** Saudi electronic payments reach 79% of retail transactions (citing SAMA). Available at: [arabnews.com](https://arabnews.com) [Accessed 22 June 2026].
- **Central Bank of Kenya (2022)** Digital Credit Providers Regulations 2022. Nairobi: CBK. Available at: [centralbank.go.ke](https://centralbank.go.ke) [Accessed 22 June 2026].
- **Competition Authority of Kenya (2021)** Annual Report & Financial Statements 2021. Nairobi: CAK. Available at: [cak.go.ke](https://cak.go.ke) [Accessed 22 June 2026].
- **Epoch AI (2025)** Over 30 AI models have been trained at the scale of GPT-4. Available at: [epoch.ai](https://epoch.ai) [Accessed 22 June 2026].
- **European Commission (2024)** Regulation (EU) 2024/1689 (Artificial Intelligence Act). Official Journal of the European Union, OJ L, 2024/1689. Available at: [eur-lex.europa.eu](https://eur-lex.europa.eu) [Accessed 22 June 2026].
- **ILO (2025)** Navigating the digital and artificial intelligence revolution in Arab labour markets. Geneva: International Labour Organization. Available at: [ilo.org](https://ilo.org) [Accessed 22 June 2026].
- **IMF (2025)** Artificial intelligence in Qatar: assessing the potential economic impact. Selected Issues Paper. Washington, DC: International Monetary Fund. Available at: [imf.org](https://imf.org) [Accessed 22 June 2026].
- **ISO/IEC (2023)** ISO/IEC 42001:2023 – Information technology – Artificial intelligence – Management system. Geneva: International Organization for Standardization. Available at: [iso.org](https://iso.org) [Accessed 22 June 2026].
- **Moffatt v. Air Canada (2024)** 2024 BCCRT 149. British Columbia Civil Resolution Tribunal. Available at: [decisions.civilresolutionbc.ca](https://decisions.civilresolutionbc.ca) [Accessed 22 June 2026].
- **NIST (2023)** Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1. Gaithersburg, MD: National Institute of Standards and Technology. Available at: [nist.gov](https://nist.gov) [Accessed 22 June 2026].
- **Pilz, K. et al. (2025)** Trends in AI Supercomputers. Epoch AI. Available at: [epoch.ai](https://epoch.ai).
- **QCRI (2024)** Fanar: Arabic large language model. Doha: Qatar Computing Research Institute. Available at: [qcai.qcri.org](https://qcai.qcri.org) [Accessed 22 June 2026].
- **Revolut (2024)** Revolut Group Holdings Ltd annual report 2024. Available at: [assets.revolut.com](https://assets.revolut.com) [Accessed 22 June 2026].
- **Royal Commission into the Robodebt Scheme (2023)** Final report. Canberra: Commonwealth of Australia. Available at: [robodebt.royalcommission.gov.au](https://robodebt.royalcommission.gov.au) [Accessed 22 June 2026].
- **SDAIA (2024)** National Data Bank. Riyadh: Saudi Data and Artificial Intelligence Authority. Available at: [data.gov.sa](https://data.gov.sa) [Accessed 22 June 2026].
- **SPA (2024)** SDAIA, Microsoft launch ALLaM Arabic language model on Azure. Saudi Press Agency. Available at: [spa.gov.sa](https://spa.gov.sa) [Accessed 22 June 2026].
- **Stanford HAI (2025)** Artificial Intelligence Index Report 2025. Stanford, CA: Stanford Institute for Human-Centered AI. Available at: [hai.stanford.edu](https://hai.stanford.edu).
- **TDRA (2025)** UAE digital achievements. Telecommunications and Digital Government Regulatory Authority. Available at: [tdra.gov.ae](https://tdra.gov.ae) [Accessed 22 June 2026].
- **TII (2024)** Falcon 3: Technology Innovation Institute launches its small AI models. Abu Dhabi: Technology Innovation Institute. Available at: [tii.ae](https://tii.ae) [Accessed 22 June 2026].
- **UAE Cabinet (2026)** Framework to deploy agentic AI across 50% of government sectors, services and operations within two years. Available at: [uaecabinet.ae](https://uaecabinet.ae) [Accessed 22 June 2026].
- **WEF (2025)** Future of Jobs Report 2025. Geneva: World Economic Forum. Available at: [reports.weforum.org](https://reports.weforum.org) [Accessed 22 June 2026].
- **World Bank (2024)** Digital credit. World Bank Digital Finance. Available at: [digitalfinance.worldbank.org](https://digitalfinance.worldbank.org) [Accessed 22 June 2026].

## About Meridian

Meridian is Synarchy Consulting's quarterly institutional foresight publication on the structural shifts in strategy, capital, technology, talent and governance that reshape how institutions from the Gulf across Africa compete and endure. Where our monthly briefing reacts to signals, Meridian discerns the shifts beneath them, and asks what institutions must decide to stay ahead.

Each issue takes a different shift, but reads it through one enduring question: which foundations of advantage must an institution own, and which can it safely rent? The subject changes every quarter; the question does not.



**Synarchy Consulting** is a boutique strategy-to-execution advisory firm, headquartered in Dubai, that helps governments, public-sector institutions, enterprises and family businesses across the Middle East and Africa translate strategic ambition into measurable outcomes. Synarchy means leading together. We deliver results, not reports, and we publish the thinking behind them.

[synarchygroup.com](https://synarchygroup.com)