

WHITE PAPER · SEPTEMBER 2026

The AI-native wealth institution

Where banks can build advantage

Alexis Calla
Chief AI Officer, Alpheya

Sandeep Mukherjee
Senior AI Advisor, Alpheya

AI in wealth management
Third paper in the series

ABOUT THIS SERIES

This article is part of Alphey's thought leadership series on AI in wealth management. The series began with *Beyond the Hype: Is GCC Wealth Management Ready for AI?*, which argued that AI had crossed a practical readiness threshold for wealth management, especially in the GCC. It continued with *AI, wealth management, and trust: Could machines ever replace human advisors?*, which explored the trust equation and the case for hybrid human-AI advisory models. This piece advances that arc: how banks and wealth managers become AI-native across the client relationship, distribution, and operations. Our upcoming whitepapers will go on to cover the legal and regulatory operating model for AI-powered wealth and the emerging economics of running an AI-native operating model, as agentic reasoning changes marginal economics and adoption matures.

CONTENTS

Key takeaways	3
01 The strategic question for wealth executives	4
02 Investor adoption is already here	5
03 Distribution is restructuring beyond the front door	7
04 Where incumbent banks can build a moat	9
05 Three AI-native business models	11
06 A path to secure the moat	13
07 The opportunity for GCC wealth managers	14
08 Key questions for banking and wealth management executives	17
What Alphey is building toward	18
Sources	19

Key takeaways

01

AI in wealth management has entered its fourth wave: from systems that encode rules, recognize patterns, and generate content to agents that can take bounded action. The winning architecture combines AI reasoning with a deterministic core for calculations, thresholds, policy, and execution.

02

Investor adoption is already mainstream: 73% of affluent and high-net-worth investors use AI for finance and investment. But AI broadens exploration more than it determines final decisions: it generated 32% of recent investment ideas but was most influential in only 12% of final decisions, while professionals and institutions remained more influential at the point of decision.

03

Human advice therefore becomes more concentrated, likely not obsolete. AI can handle routine preparation and guided self-direction; advisors provide judgment, reassurance, and accountability when decisions are complex or consequential.

04

Distribution is fragmenting across owned AI experiences, agent-mediated channels, and AI-native operations. Frontier-lab AI superapps are one emerging distribution mode, not the whole strategy.

05

As model intelligence becomes widely available, the moat moves to governed context, embedded institutional judgment, trusted execution, and the accountable relationship. An institution-specific harness makes these assets operational through tools, permissions, deterministic controls, human checkpoints, evaluation, and audit.

06

For GCC institutions, the opportunity is to bridge the gap between self-directed brokerage and high-touch DPM through guided self-direction, while applying a regional lens across three business models. The path is deliberate: start bounded, govern before deployment, design human checkpoints, and make every action auditable from day one. Alpheya helps institutions build that governed operating model on an AI-ready core.

The strategic question for wealth executives

AI-native does not mean putting AI everywhere.

AI and its application to wealth management have moved through distinct waves. Symbolic systems encoded rules. Machine learning recognized patterns. Generative AI produced content: could it summarize a portfolio, draft an RM note, prepare a meeting brief, explain a market event, or help a client understand risk? That question has largely been answered. The technology is useful, and leading firms are bringing it into daily advisory workflows.

The wave now underway is different in kind. Agentic AI does not just respond; it acts: gathering the data it needs, using tools, making decisions, and continuing until the work is complete or a human checkpoint is reached. The practical frontier has shifted from “can the model reason?” to “can the agent reason, use tools, maintain context, recover from errors, and act safely inside a bounded environment?” Increasingly, with the right architecture and the right controls, the answer is yes.

The next question is broader than where the client starts or which AI interface wins. It is how a bank or wealth manager becomes AI-native across the client relationship, the advisor operating model, distribution, and operations, without surrendering trust, accountability, or economic relevance.

AI-native does not mean putting AI everywhere. It means redesigning how intelligence, judgment, and action come together, then choosing deliberately where the institution should own the experience, where it should participate through partners or external agents, and where it should automate work behind the scenes. Increasingly, it also means pairing that intelligence with a deterministic core (the systems that handle calculations, enforce thresholds, and apply policy) so reasoning handles the variable, judgment-heavy tail rather than carrying the whole task.

For most institutions, that deterministic core does not need to be built from scratch: it already exists as the core banking or portfolio management platform. But it typically must be modernized, with real-time data access, exposed APIs, and governed controls, before it can safely carry this role.

We see a role for using AI to understand and orchestrate and using deterministic systems to calculate and control. This division supports repeatable financial outcomes, clearer auditability, bounded autonomy, and the ability to change model providers without rebuilding the institution's core controls.

Investor adoption is already here

Investor behavior has moved faster than many institutional strategies. HSBC and Ipsos surveyed 9,993 affluent and high-net-worth investors across ten markets in early 2026. Seventy-three percent said they use AI for finance and investment. The most common uses were analysis and research, cited by 66%; strategy support, cited by 50%; and a second opinion on their own ideas, cited by 31%.

AI is already influencing the beginning and middle of the decision journey. It helps investors discover ideas, compare options, interpret information, and approach an advisor with a more developed point of view. But it is not yet the primary source of conviction. AI generated the last investment idea for 32% of respondents and was the most influential factor in the final decision for only 12%. Financial professionals and institutions generated ideas for 62% and were most influential for 37%.


AI adoption is also not behaviorally neutral. Forty-nine percent said AI made them more willing to take calculated risks. Fifty-one percent said it made them feel more in control, while 26% said it made them feel less in control. The same tool can increase confidence, activity, and uncertainty at once. Wealth institutions therefore need to govern not only what AI can do, but how speed and convenience may change investor behavior.

The advisor becomes the assurance layer

The division of labor is changing. Agents can expand the search space, synthesize information, and prepare choices. Advisors become most valuable when a choice carries consequence and someone must interpret context, challenge assumptions, and stand behind the outcome.

The HSBC findings make that role concrete. Eighty percent cited reassurance as a reason to turn to a professional advisor at the point of decision, and 72% cited strategic expertise. The most valued advisor contributions were applying judgment and validation, spotting mistakes in AI-generated data, and providing personalized interpretation of complex information.

EXHIBIT 1 Many investors now use AI to inform their financial decisions, yet advisors retain a key role for now



Source: HSBC and Ipsos, The Trust Threshold, June 2026. Survey of 9,993 affluent and high-net-worth investors across ten markets.

That reassurance does not require an advisor embedded in every case. For a routine decision, it can come from knowing a person is a message away; for a complex or consequential one, it means the advisor stays involved throughout. That distinction (routine versus consequential, core versus edge) is what later determines which business model fits which client.

This is not a defense of the traditional advisor role unchanged. Advisors will be less valuable as information gatekeepers and more valuable as accountable interpreters, behavioral coaches, and decision partners. The winning model should automate preparation and routine coordination while concentrating human attention where trust, suitability, exceptions, or life context matter.

Generative AI raised questions about the quality of advice. Agentic AI raises a further question: accountability for actions. When a system sends a communication, triggers a workflow, or updates a record on its own initiative, the question of who is responsible becomes significantly more acute, and it is one reason execution-oriented agentic capabilities will need stronger governance than purely advisory ones.

Human checkpoints should therefore be designed, not scattered everywhere. Too little oversight creates conduct and trust risk. Too much creates review fatigue and gives back the productivity gains. The goal is precise human involvement at the moments that change the client's financial path.

Distribution is restructuring beyond the front door

Frontier-lab AI products such as ChatGPT and Claude are evolving into general-purpose work environments that combine conversation with files, tools, applications, and multi-step task execution. We use “AI superapp” as shorthand for this emerging category; the term is not yet common in GCC wealth management. These environments matter, but they are only one part of a broader shift: distribution itself is fragmenting across three complementary modes:

- 1 **Owned AI experiences:** AI embedded in investor apps, RM workbenches, advisor portals, onboarding, servicing, portfolio analytics, and product journeys. These remain the places where institutions express their brand and relationship model.
- 2 **Agent-mediated distribution:** approved external agents, partner platforms, or frontier AI workspaces call selected financial capabilities on behalf of a client or advisor. The institution may not own the first screen, but it can still own the trusted capability.
- 3 **AI-native operations:** agents orchestrate internal workflows across research, advice preparation, onboarding, servicing, compliance, portfolio operations, and exception handling. Much of the economic value may emerge here, away from the visible interface.

These are not competing futures. Institutions may use all three. The strategic task is to decide which journeys belong in each mode, how they connect, and which context, controls, and economics the institution must retain.

AI ‘superapps’ are one potentially disruptive distribution model

Wealth management is predominantly a trust and distribution business. Frontier lab AI products are moving from answer engines toward work environments that can use files, tools, applications, and background workflows. The user no longer asks only for information; the user delegates an outcome. Financial research, document analysis, planning, comparison, and eventually bounded action will be pulled into those environments; they are evolving into trusted distribution sources for a range of content, the core of a wealth manager’s business.

Governance is not merely a technical control layer. It is how legal duties become executable system behavior.

This is not only a retail brokerage phenomenon. The major AI providers have independently converged on a similar architecture for letting agents act safely: separating the reasoning layer (the model and its instructions) from an execution layer (the tools and sandboxes that take action) and an audit layer (a durable log of everything the agent did). Anthropic's Claude Agent SDK, OpenAI's Agents SDK, and Google's Agent Development Kit each formalize some version of this pattern, and protocols are emerging to standardize it: the Model Context Protocol for how agents connect to tools and data, and the Agent2Agent standard, now under the Linux Foundation, for how agents connect to each other. The vendor differs; the pattern is the same, and for wealth managers, the pattern matters more than any single vendor's product.

Robinhood's Agentic Trading product is an early market signal of this pattern in a client-facing context. It allows a third-party AI agent to connect to a dedicated account, access defined account information, and place orders inside a bounded structure. The lesson for wealth executives is not that every bank should copy agentic trading. It is that financial capabilities are becoming callable by agents.

That creates a new distribution question. If a client's preferred AI can understand a portfolio, compare options, prepare questions, or request an action, will the bank be invisible, obstructive, or trusted enough to participate? The answer depends less on owning every interface and more on making selected capabilities safely available without giving away control of the relationship.

Execution will move more slowly than research and preparation because liability and regulation require stronger boundaries; the accountability question above applies with full force once money moves.

That caution is not theoretical. In June 2026, U.S. lawmakers pressed the SEC on how broker-dealer obligations apply when third-party AI agents access accounts, use customer data, and potentially initiate trades, covering Regulation Best Interest, supervision, best execution, and trade approval procedures. Institutions should not wait for that clarity before acting, though: the discovery, explanation, and document layers are already moving.

Agentic AI turns governance into a legal and compliance design question

Once a system can recommend, communicate, or initiate action, institutions must determine which activities fall within the regulatory perimeter, who is acting on whose behalf, and who remains accountable when an outcome is challenged. The practical answer is to treat each use case as a regulated operating model: classify whether it informs, recommends, or acts; allocate responsibilities between the licensed institution and technology provider; involve legal and compliance in defining the permitted purpose, users, data, disclosures, jurisdictions, and degree of autonomy; and enforce those decisions through deterministic policy gates, explicit client or human approval for consequential actions, complete records, monitoring, escalation, and suspension. Governance is therefore not merely a technical control layer. It is how legal duties become executable system behavior.


Where incumbent banks can build a moat

Three shifts are converging here, not one. Frontier intelligence is available to every institution and technology provider, so model access alone is not a durable advantage. Distribution itself is restructuring, as financial capabilities become callable by agents and new modes of reaching clients open up for incumbents and challengers alike. And investors still want human judgment at the point of decision, even as they lean on AI everywhere else. Together, these shifts move value toward assets that cannot be replicated overnight:

- **Governed context:** positions, transactions, liabilities, client history, goals, preferences, mandates, and interactions held with current entitlements and provenance.
- **Embedded institutional judgment:** suitability rules, product logic, investment policy, Shariah considerations, escalation thresholds, and knowledge of which exceptions matter.
- **Trusted execution:** the regulated accounts, approval paths, custody, payments, trading, servicing, and audit trails that turn an idea into a defensible action.
- **The accountable relationship:** the advisor, institution, and brand that clients trust when markets move, life changes, or a recommendation must be challenged.

Incumbency can be a liability when these assets remain fragmented across legacy systems and organizational silos. It becomes an advantage when the institution can orchestrate them as a governed system of action. The bank should not become a passive data provider to an external AI. It should become the trusted context, judgment, and execution layer that approved AI experiences must call.

EXHIBIT 2 Where incumbent banks can build a moat



Source: Alpheya.

The harness is where domain knowledge becomes operational

That governed operating layer is increasingly described as an agent harness: the context, tools, memory, permissions, deterministic controls, human checkpoints, evaluation, and audit infrastructure that surrounds a foundation model and enables it to act reliably.

A generic harness is not a moat. The defensible asset is the institution-specific harness in which proprietary client context, investment and suitability judgment, workflow knowledge, execution controls, and learning from real outcomes compound over time. As model intelligence becomes more widely available, this layer becomes more important: models may absorb prompt tricks and basic scaffolding, but greater autonomy increases the need for domain context, verification, and accountability.

Three AI-native business models

Most cases are simple; a minority carry most of the complexity, and that minority is where human judgment earns its cost.

Not every client situation needs the same mix of agent and advisor. AI agents cannot replace advisors outright, but for the core, everyday cases (routine rebalancing, standard goal tracking, common product questions), they can be a genuinely sufficient solution on their own, with an advisor available on demand rather than involved throughout. It is the edge cases (complex family structures, unusual tax or legal situations, major life events) that still need an advisor end to end.

A comparable pattern shows up in banks' own operations. In a small vendor-led study of 85 reconciliation events at a Nordic fund administrator, 36% of cases needed no AI involvement at all, 34% needed a single model call, and only 19% and 11% needed multi-step or deep reasoning respectively. Most cases are simple; a minority carry most of the complexity, and that minority is where human judgment earns its cost.

Banks and wealth managers can use their moat across three practical models, each suited to a different point on that core-to-edge spectrum. They are strategic lenses rather than mutually exclusive product lines:

- 1 **The augmented relationship model:** equip advisors with richer preparation, portfolio explanations, product intelligence, meeting follow-up, and next-best-action support, best suited to the complex, edge-case relationships where the advisor needs to stay involved throughout. The value is greater advisor capacity and a more consistent premium relationship, not a faceless digital substitute.
- 2 **The AI-enabled self-directed brokerage model:** add an always-on layer of research, portfolio and risk explanation, alerts, education, product comparison, and goal-aware guidance to the self-directed journey, best suited to the core, everyday cases. AI helps investors frame decisions and understand trade-offs; deterministic systems enforce eligibility, risk, order limits, disclosures, and execution controls. Human advisors or specialists enter when a decision is consequential, complex, or outside the investor's confidence. This creates guided self-direction without presenting every interaction as full advisory.

- 3 **The trusted capability model:** expose selected data, analytics, explanations, suitability checks, workflow steps, and bounded actions to approved agents and partners, governed by clear identity and consent, data scopes, and action levels from read-only to fully executable, and priced for what's exposed rather than for owning the interface. The bank can participate in new distribution channels even when another platform owns the interface.

None of these three is complete without a regional lens. In the GCC specifically, differentiation also comes from Arabic experiences, Shariah-aware journeys, local product shelves, regulatory interpretation, in-region infrastructure, and the cultural context of family and intergenerational wealth. That isn't a fourth model so much as a filter: it should shape how whichever of the three an institution pursues actually gets built in this market, and that's also what makes the result hard for a horizontal global platform to reproduce.

The strongest institutions will combine the three models. Advisor augmentation covers complex edge cases. AI-enabled self-direction provides always-on guidance across core brokerage relationships. A trusted capability layer extends distribution. Run through a regional lens, all three become harder to copy. The cost-to-serve and margin implications of that combination will be the subject of a separate report.

A path to secure the moat

The governance layer is not a constraint on agentic AI; it is what makes it deployable in a regulated industry.

Knowing the moat is not the same as securing it. The transition from advisory AI to agentic, execution-capable AI is not a single leap; it is a deliberate progression, and firms that have done the foundational work on generative AI (data pipelines, governance, advisor trust in AI recommendations) are considerably better positioned to take the next step than those that haven't. Four principles govern that progression:

- 1 **Experiment: start with bounded agents.** The first agentic deployments should operate within tightly defined boundaries: a single workflow, a single data source, a specific and limited set of permitted actions. Expand scope only as reliability, and the trust that goes with it, is demonstrated.
- 2 **Govern: establish explicit governance before deployment.** Define what agents can and cannot do, when human approval is required, and how every action is logged, before deployment, not after. The governance layer is not a constraint on agentic AI; it is what makes it deployable in a regulated industry.
- 3 **Build trust: design every workflow around the human in the loop.** The question is not whether humans should be involved, but at which moments and for which decisions. For high-stakes actions (trade execution, client communications, compliance escalations), human approval should be mandatory by design, not optional by default.
- 4 **Audit first: treat every agent action as auditable from day one.** Before deploying an agentic system, ensure every action it takes can be logged, retrieved, and explained. In regulated financial services, auditability and recordkeeping are often required and are foundational to the trust that makes agentic AI useful in a client-facing context.

Applied to the moat rather than to agentic AI in the abstract, these four principles map onto the four assets that make up the moat itself. Bounded experimentation is how governed context gets tested without exposing the whole client relationship. Governance-before-deployment is what turns embedded institutional judgment into something an agent can act on safely. Human-in-the-loop by design is what keeps trusted execution trusted. And auditability from day one is what keeps the accountable relationship accountable, even when the action was taken by a system rather than a person.

The opportunity for GCC wealth managers

The GCC has an opportunity to shape an AI-native wealth model rather than import a generic one. Regional wealth is growing, client needs are increasingly cross-border, and institutions sit at the intersection of private banking, family wealth, investments, lending, and complex servicing. The opportunity is larger than adding a chatbot to existing channels.


Bridging the GCC's self-directed-DPM divide

In much of the GCC, the market is polarized. Self-directed brokerage serves a broad, digitally active investor base, while discretionary portfolio management and private banking serve the wealthier end. Scaled advisory between these poles is comparatively limited. The result is a missing middle: investors have access to execution, but often not to continuous, personalized guidance.

AI creates a practical way to narrow that gap without forcing every client into a traditional advisory model. At the self-directed end, it can provide always-on education, research, portfolio and risk explanations, goal-aware alerts, product comparison, and decision preparation. Deterministic systems keep eligibility, suitability where applicable, pricing, order controls, disclosures, and execution within policy. When stakes or complexity rise, the journey can escalate to a human advisor or specialist.

At the DPM and private-banking end, the same architecture can strengthen mandate monitoring, advisor preparation, personalized communication, and exception handling. The strategic opportunity is to create a continuum from self-service to guided self-direction to human-led advice, with the level of support and control matched to the client and the decision.

EXHIBIT 3 AI can bridge the GCC's missing middle



Source: Alpheya.

The UAE offers an early signal of the pace of investor change. Among 703 UAE respondents in the HSBC/Ipsos survey, 83% used AI for finance and investment, compared with 73% globally. Sixty-three percent said AI helped them take calculated risks, 56% preferred a hybrid AI-and-advisor model, and 39% preferred using AI for initial exploration followed by advisor validation.

83%	63%	56%	39%
of UAE respondents use AI for finance and investment, against 73% globally	said AI helped them take calculated risks	preferred a hybrid AI-and-advisor model	preferred AI for initial exploration followed by advisor validation

Source: HSBC and Ipsos, 703 UAE respondents, 2026.

GCC governments have invested significantly in AI infrastructure and national AI strategies, while financial regulators including SAMA, DFSA, and the CMA have generally supported responsible adoption. But the frameworks governing autonomous AI action in financial services are still being developed, and the data sovereignty requirements that apply to generative AI apply equally, and more urgently, to agentic systems that act on client data. Firms that engage proactively with regulators, document their governance frameworks carefully, and design systems that can demonstrate compliance as the rules evolve will be better positioned than those that react.

GCC institutions are not starting from a regulatory blank page. Existing obligations around governance, suitability, customer protection, outsourcing, data protection, recordkeeping, and senior-management accountability already apply. What is still developing is how those duties should be interpreted when AI agents act across systems, channels, and third parties with increasing autonomy.

These findings should not be generalized across the GCC. They do suggest that some regional institutions may need to govern highly AI-active investors earlier than peers elsewhere. Arabic and bilingual journeys, data residency, in-region infrastructure, Shariah-aware investing, market-specific regulation, and the cultural importance of trusted relationships should be designed into the model from the start.

Key questions for banking and wealth management executives

The practical agenda is not to launch an agentic AI program in the abstract. Executive teams should make a small number of strategic choices:

- | | |
|--|--|
| <p>01 Which of the three AI-native business models best supports our growth, service, and margin priorities, and where does a regional lens sharpen it further?</p> | <p>02 Where can AI-enabled guidance close the gap between self-directed brokerage and high-touch DPM or advisory?</p> |
| <p>03 Which outcomes must be calculated or controlled deterministically, and where can AI add reasoning, synthesis, or explanation?</p> | <p>04 Where are clients and advisors already using AI outside our controlled environment?</p> |
| <p>05 Which moments require human reassurance, judgment, or accountability, and which do not?</p> | <p>06 Which client context, institutional judgment, and execution capabilities must remain proprietary?</p> |
| <p>07 Which capabilities should be usable through owned channels, approved partners, or external agents?</p> | <p>08 Where on the experiment – govern – build trust – audit path are we today, for each moat asset, and what evidence would justify moving to the next step?</p> |
| <p>09 Which GCC requirements create differentiation rather than only compliance work?</p> | <p>10 Which outcomes will demonstrate value: cost-to-serve, advisor capacity, client confidence, exception rates, conversion, or revenue per relationship?</p> |
| <p>11 Who is legally and operationally accountable at each level of AI autonomy, and what disclosure, evidence, review, and redress must accompany that accountability?</p> | |

The future of wealth management is unlikely to be fully autonomous. It is also unlikely to remain merely assistive. It will combine AI-led exploration, human assurance, governed capabilities, and increasingly autonomous operations. Banks do not need to own every interface to remain economically central. They do need to own the context, judgment, execution, and trust that make financial action credible. The interface may fragment. The trust layer and the economic relationship do not have to.



What Alpheya is building toward

Alpheya's remit is to sit on that path (experiment, govern, build trust, audit) mapped to each moat asset rather than offered as a generic capability list. Concretely:

- **AI-ready core:** we help transform the legacy core, positions, transactions, and ledgers, into a platform with real-time data access, exposed APIs, and governed controls, so the deterministic foundation agentic AI depends on is actually in place.
- **Governed context:** we connect AI to real client, portfolio, product, and suitability logic, rather than a disconnected conversational layer, so context stays governed as it's put to use.
- **Embedded judgment:** we design AI outputs to be machine-readable as well as human-readable, so suitability rules, product logic, and escalation thresholds can be encoded and acted on, not just described.
- **Trusted execution:** we build the governance layer (the permissions model, the audit trail, the human-approval checkpoints) in parallel with the capability layer, not after it.
- **The accountable relationship:** we design the human-in-the-loop role deliberately, and build GCC requirements (Arabic and bilingual journeys, Shariah-aware investing, data residency, local regulation) into the model from the start, rather than retrofitting them.

The transition to agentic AI in wealth management will not happen all at once. It will happen use case by use case, as firms develop confidence in specific, bounded workflows. The firms that move earliest and most deliberately (not the fastest, but the most thoughtfully) will accumulate the operational experience and client trust that makes broader deployment possible. We are committed to being their partner in that journey.

Sources

- 1 HSBC/Ipsos, The Trust Threshold: AI makes investors bolder, but they want human judgement to make decisions, June 24, 2026.
- 2 Emirates News Agency (WAM), HSBC report finds UAE investors lead globally in AI adoption, June 24, 2026.
- 3 Robinhood Agentic Trading overview:
<https://robinhood.com/us/en/support/articles/agentic-trading-overview/>
- 4 AdvisorHub, Democratic lawmakers press SEC over brokers' use of agentic AI, June 2026.
- 5 OpenAI Codex product page: <https://openai.com/codex/>
- 6 Anthropic Claude Cowork product page:
<https://www.anthropic.com/product/claude-cowork>
- 7 IBS Intelligence and Profile Software, Agentic Triada: The New Economics of Wealth Management, 2026.
- 8 OpenAI, Harness engineering: leveraging Codex in an agent-first world, February 2026.
- 9 Anthropic, Effective harnesses for long-running agents, November 2025.
- 10 Birgitta Böckeler, Thoughtworks, Harness engineering for coding agent users, April 2026.
- 11 McKinsey, The symbiotic enterprise: A new model for growth, July 2026.



AUTHORS

Alexis Calla

Chief AI Officer, Alpheyra

Sandeep Mukherjee

Senior AI Advisor, Alpheyra

ABOUT ALPHEYA

Alpheyra is an AI-native wealth and trading platform for banks and wealth managers, headquartered in ADGM, Abu Dhabi, and backed by BNY, Lunate and Pershing.

alpheyra.com