



Autonomous
Finance
by
Amna Usman Chaudhry

The State of Autonomous Finance 2026

An Executive Brief for Financial Institutions

RESEARCH, ANALYSIS & EDITORIAL DIRECTION BY

Amna Usman Chaudhry

Launched at the Fintech Tuesdays Business Breakfast

Dubai, UAE
17 July 2026

CONTENTS

Inside the Brief

00	Publication note	3
01	Where the Market Stands	4
02	The Sixty Second Read	6
03	Ninety Days That Moved the Market	7
04	Defining Autonomous Finance	12
05	The Map: Six Layers of the Autonomous Finance Stack	15
06	Capability domains	16
07	The Case for Waiting	21
08	The KPI Is Not the Number of AI Agents	22
09	Safe Adoption as an Operating Model	23
10	Beyond Individual Agent Governance	24
11	Seven Questions Every Financial Institution Should Ask	26
12	The UAE Advantage	28
13	What to Watch Next	29
14	Conclusion	30
--	The Author	31
App. A	Methodology and Scope	32
App. B	Notes and Sources	33

EXHIBITS

01	The position at mid-2026	4
02	Ninety days that moved the market	10
03	The Autonomous Finance stack	18
04	The unit of governance keeps moving	24

Publication note

This report is a Fintech Tuesdays Business Breakfast Brief, prepared for the business breakfast Agentic AI and the Future of Autonomous Finance (Dubai, 17 July 2026) and for the wider Fintech Tuesdays community. The research, analysis, and editorial direction are those of Amna Usman Chaudhry and were developed through her Autonomous Finance publication.

It is an executive market snapshot based on publicly available information current to 14 July 2026. It is intended to support discussion, strategy, and further research. It is not legal, regulatory, investment, or technology advice, and it does not represent the views of Fintech Tuesdays or of any organization with which the author is affiliated. Product availability, market data and regulatory positions may change.

SUGGESTED CITATION


Chaudhry, Amna Usman (2026). The State of Autonomous Finance 2026: An Executive Brief for Financial Institutions. Fintech Tuesdays Business Breakfast Brief.

SECTION 01

Where the Market Stands

Two findings frame the position of financial institutions at mid 2026.

The first is commercial. A survey of 522 finance leaders published in June found that 71 percent already factor a bank's agentic AI capability into choosing or expanding a banking relationship, rising to nearly eight in ten among mid-market companies. Sixty-nine percent said they would pay their bank for secure, compliant, auditable agentic connectivity.^[37] Agentic capability has become a selection criterion for commercial banking relationships. The competitive contest is moving from the account layer to the workflow layer, and the ERP systems, treasury platforms, and fintech applications already embedded in clients' workflows are positioned to capture that role if banks do not.



The second is operational. Half of banking executives report that governance and compliance considerations are already limiting AI performance, while 18 percent are confident they could pass an independent audit of their AI controls.^[38] Around 12 percent of financial institutions describe their AI strategy as well defined and resourced.^[40] A third of bank leaders report that they do not understand agentic AI, while 84 percent are concerned about AI-enabled fraud targeting their customers.^[39]


Those two findings describe the same institution from different angles: rewarded by the market for moving, and not yet able to evidence that it is in control while doing so.

In the UAE, both conditions arrive with a timetable attached. In April 2026 the Cabinet approved a framework to move 50 percent of government sectors, services, and operations to agentic AI within two years. The first live government agents are in service across procurement, tax auditing, customer happiness, and technical support, and 80,000 government workers are being trained.^[18] In June, Dubai approved an executive plan extending the shift to the private sector: 295,000 companies, 100 specialized AI assistants over two years, and 50 new agentic AI firms.^[33] Meanwhile the Central Bank of the UAE's Responsible AI guidance, in force since February 2026, makes boards and senior management responsible and accountable for AI systems and outcomes, states that institutions should not deploy AI they cannot control, and confirms that responsibility for outsourced AI functions remains with the institution.^[12]

The accountability principle arrived before the agents did. For institutions in this market, deployment speed and board accountability are not sequential phases. They are the same deadline.

Beneath both findings sits a shift that neither survey names. Financial agency, the capacity to interpret conditions, decide, and initiate a transaction, is ceasing to be exclusively human. That is the subject of this report.

It sets out what changed in the first half of 2026, provides a working map of the infrastructure now assembling, presents the strongest case for caution, and closes with the questions an institution can use to move at market speed while remaining able to stand behind the outcome.



One view, stated once: Autonomous Finance will not scale on capability alone. It will scale where value, trust, and accountability advance together.

SECTION 02

The Sixty Second Read

Five things that changed in the first half of 2026

- 01 **Software became a counterparty.** ClawBank and Shodai reported a demonstration in which two AI-agent entities negotiated, signed, and executed a Ricardian agreement linked to automated settlement, without human signatures.^[22]
- 02 **An AI-enabled adviser entered a fiduciary structure.** Coinbase Advisor launched through an entity registered with the SEC, CFTC, and NFA, placing an AI-enabled advisory experience inside a regulated structure subject to fiduciary obligations.^[28]
- 03 **The payment networks shipped.** x402 has processed more than 100 million machine-to-machine payments; Mastercard launched Agent Pay for Machines with more than 30 partners; Visa opened a single on-ramp across competing protocols.^{[6][7][8]} Volume remains early, which is precisely the window in which governance can still be designed in rather than retrofitted.
- 04 **The UAE cleared programmable money built for machines.** DDSC, the regulated dirham-backed stablecoin, received approval to reach VARA-regulated exchanges, with its sovereign-linked backers explicitly naming machine-to-machine transactions and trade between AI agents as design goals.^{[20][21]}
- 05 **Supervisors moved from principles to practice.** The FSB, IOSCO, Singapore's IMDA, and the CBUAE all issued operational guidance for agentic AI within weeks of each other, converging on one theme: human accountability must survive autonomy.^{[9][10][11][12]}

The infrastructure is moving faster than settled business models, customer behavior, and institutional practice. That gap is the opportunity.

SECTION 03

Ninety Days That Moved the Market

What actually happened and why each moment matters

The clearest market signal of 2026 is not any single product. It is the simultaneous movement across contracts, advice, execution, payments, money, and supervision, much of it inside a single quarter. The following moments tell the story.

18 June: Two AI companies signed a contract with each other

ClawBank and Shodai reported what they described as the first Ricardian contract negotiated and signed entirely between autonomous AI agents: two AI systems, each operating through its own registered legal entity, negotiated terms, applied legal e-signatures, tied the agreement to on chain escrow, and released payment automatically on milestone acceptance. ClawBank's agent had earlier completed the legal formation of its own US company, receiving a federal EIN and opening an FDIC-insured bank account. These remain demonstrations built on human-designed scaffolding, and should be read as early signals rather than regulated deployments. Their significance is directional: agent-to-agent coordination is acquiring legal and financial form, and corporate law has met a company whose active decision-maker is not a person.^[22]

16 June: An AI-enabled adviser entered a fiduciary structure, and surfaced a contradiction

Coinbase Advisor launched as the first in-app conversational AI adviser offered through an entity registered simultaneously with the SEC as an investment adviser and with the CFTC and NFA as a commodity trading adviser.^[28] Days earlier, Coinbase for Agents connected AI agents directly to customer accounts to trade, pay, and execute within user-defined limits.^[13]

Coinbase deserves credit for the harder path. Most generative AI in finance shelters behind the same sentence: for informational purposes only, not a recommendation. Registration as an investment adviser is not a label. Under the Investment Advisers Act, the registered advisory entity owes fiduciary duties, including duties of care and loyalty. Coinbase placed an AI-enabled advisory experience inside that regulated structure.

In doing so, it surfaced the contradiction the whole field now has to resolve.

The registration places fiduciary responsibility on the regulated advisory entity. The product's disclosures state that outputs may be inaccurate or incomplete and emphasize the customer's continuing responsibility for investment decisions. Fiduciary obligations at the entity level therefore sit alongside disclosures about model limitations and user responsibility, and the legal treatment of that structure has not yet been meaningfully tested in public.

This is a clear public example of a structural problem, not a Coinbase failing. Every element is individually compliant. A registered adviser generated the recommendation. The disclosures were shown. The customer clicked approve. Yet when an outcome proves unsuitable, responsibility scatters: the fiduciary registration points at the firm, the disclaimer points at the customer, and the

approval click looks like human authority while functioning as a confirmation rather than an independent judgment about whether the advice met a fiduciary standard in that customer's situation. The action is perfectly locatable. The accountability is not.

This is the pattern I have termed a delegation failure in my research: every actor in the chain performs its delegated role in a compliant manner, and responsibility for the outcome still cannot be located. The pattern, and the term, will recur.

Robinhood's model allocates responsibility differently and no more comfortably: it states plainly that users bear full responsibility for outcomes and that it does not supervise, control, or guarantee the performance of any AI agent.^[17] Between a fiduciary who disclaims the result and a platform that disclaims supervision, the industry has not yet produced a settled answer to a question every institution will eventually be asked.

May to July: Retail execution went agentic

Robinhood launched Agentic Trading and an Agentic Credit Card in late May, with more than 70,000 agentic accounts created since.^[17] On 1 July, Robinhood Chain, an Ethereum Layer 2 built for financial services, went live on mainnet with tokenized stock tokens trading around the clock in more than 120 countries, processing 17 million transactions from nearly 350,000 wallets in its first week.^[29] On 10 July, Robinhood announced Agentic Trading for crypto, with rollout to eligible US traders beginning. Kraken unveiled its own agentic trading plans the same month.^[30] Agent execution is emerging as a competitive direction across major platforms.

10 June: The payment networks committed

Mastercard launched Agent Pay for Machines with more than 30 industry participants, built for verified agents making high-frequency payments across cards, accounts, and stablecoin rails.^[7] Visa's Intelligent Commerce Connect provides a single network, protocol, and token-vault-agnostic on-ramp for agentic commerce.^[8] Earlier, Ant International open-sourced the Agentic Mobile Protocol, the first agentic payment framework for mobile wallets and super apps, with settlement designed for nano-transactions and a full Know Your Agent framework built into the rails.^[25]

Through Q2: Machine payments found their measure

The x402 protocol, which embeds payments directly into web interactions, passed 100 million processed payments, with USDC carrying the large majority of measured onchain agentic commerce volume.^[6] AWS integrated x402 into Amazon Bedrock AgentCore, letting enterprise agents settle in USDC. The honest counterweight: monthly volume remains modest. The rails are being laid at industrial speed while machine-native payment activity is still early. For institutions, that gap is not a reason for comfort. It is the window in which authority, controls, and evidence can be designed into the system rather than retrofitted after scale.

16 June: Commerce got its universal translator

Adyen launched Adyen Agentic, a three-layer suite across product feeds, carts, and payments, positioned to let merchants integrate once and participate across evolving agent platforms, with Visa, Mastercard, and American Express as launch partners.^[5] OpenAI and Stripe's Agentic Commerce Protocol has been live inside ChatGPT since late 2025; Shopify and Google's Universal Commerce

Protocol brought a major retailer coalition in early 2026.^{[3][4]} A useful calibration from inside the industry: Adyen's agentic commerce lead rates the field's maturity at roughly 0.5 on a five-point scale, noting the constraint was never the AI but the plumbing of catalogs, trust, and fraud.^[23]

3 July: The UAE cleared money built for machines

DDSC, the regulated dirham-backed stablecoin supported by IHC, First Abu Dhabi Bank, and Sirius International Holding, received a no-objection certificate to partner with selected VARA-regulated exchange platforms, having been approved on ADI Chain in February. Its backers describe it as programmable money built for secure, automated value transfer, explicitly naming machine-to-machine transactions and trade between AI agents as design goals.^{[20][21]} A central-bank-approved dirham payment token backed by major UAE institutions and explicitly positioned for machine-to-machine activity is a significant global market signal. In the same stretch, Standard Chartered became the first Global Systemically Important Bank licensed to provide institutional clients with integrated access to Circle's USDC minting and redemption, launching the capability out of its DIFC operations in Dubai.^[36]

June and July: Regulated money moved closer to the core

Circle received final approval from the U.S. Office of the Comptroller of the Currency to establish a national trust bank.^[14] The Bank of England published its policy statement and draft rules for systemic sterling stablecoins.^[15] Franklin Templeton, with more than 1.5 trillion dollars under management, launched Franklin Crypto to bring institutional strategies and tokenized funds onchain.^[31] Banco Santander signed a memorandum of understanding with Abu Dhabi's G42 on agentic financial services at global-banking scale, and separately open-sourced its own AI governance and safety tooling for any rival or regulator to adopt.^{[32][27]}

Through 2026: The quieter deployment inside regulated banks

The headlines belong to crypto-native platforms. The deployments belong somewhere less visible. Inside regulated institutions, agentic AI has landed first where the work is procedural and auditable: financial crime detection, anti-money laundering, regulatory-change triage, controls monitoring, continuous transaction monitoring, and reconciliation. Core providers have embedded agents into the systems banks already run. FIS brought agentic AI to banking in partnership with Anthropic, beginning with financial crimes and naming BMO and Amalgamated Bank among the first institutions to deploy its financial crimes agent; Fiserv launched an operating system for agentic AI in banking with third-party partners.^[41] A Wolters Kluwer survey found 44 percent of finance teams expected to use agentic AI during 2026.^[40]

The pattern across the cited deployments is an emphasis on governed environments, traceability, and human review. This is where a real bank's first agent lives, and it is also where the accountability question bites hardest. An agentic AML workflow that wrongly escalates a customer, or fails to escalate one it should have, is a supervised failure with a named accountable officer, a regulator entitled to an explanation, and a customer entitled to redress. The consequences are not merely reputational. They can be regulatory, operational, financial, and customer-facing.

EXHIBIT 2 · NINETY DAYS THAT MOVED THE MARKET

27 MAY TO 10 JULY 2026

27 May	• Robinhood opens to agents Third-party AI agents execute trades and card purchases
10 June	• Mastercard Agent Pay for Machines Thirty-plus partners; settlement across cards, accounts, stablecoins
10 June	• The supervisors converge FSB: some agent actions may be impossible to override or remedy
16 June	• An AI adviser accepts fiduciary duty Coinbase Advisor: registered with the SEC, CFTC and NFA at once
16 June	• Adyen Agentic launches Merchants integrate once across competing agent protocols
18 June	• Software becomes a counterparty Two AI entities sign and settle a contract, no human in the loop
1 July	• Robinhood Chain goes live Seventeen million transactions in week one; tokenized stocks
3 July	• The UAE clears money built for machines DDSC reaches VARA exchanges, naming agent-to-agent trade
10 July	• AI moves closer to allocating capital JPMorgan agents outperform a 60/40 portfolio in backtests
10 July	• Circle wins national trust bank approval Programmable money moves toward the institutional core

SOURCES: COMPANY, REGULATORY & INSTITUTIONAL
ANNOUNCEMENTS CITED IN THIS REPORT

THE STATE OF AUTONOMOUS FINANCE 2026 ·
AMNA USMAN CHAUDHRY

2026: The agentic economy met the onchain economy

Circle's recent work argues that the agentic economy will require financial infrastructure designed for software acting autonomously. In this model, AI agents increasingly transact, enter agreements, and coordinate economic activity, while stablecoins, programmable settlement, and onchain infrastructure provide the rails through which value can move.^[51]

The significance is that agentic AI is beginning to be framed not as a new interface for finance, but as a new class of economic participant. That brings the governance question into sharper focus: institutions will still need to determine whose authority an agent exercises, what limits apply, what evidence is retained, and who remains accountable when several systems contribute to an outcome.

June and July: Governance moved closer to runtime

The Financial Stability Board proposed twelve sound practices for responsible AI adoption, directed to boards and senior management.^[9] IOSCO issued a supervisory toolkit for AI use in capital markets.^[10] Singapore's IMDA updated its Model AI Governance Framework for Agentic AI, emphasizing responsible deployment and continuing human accountability.^[11]

In July, the Monetary Authority of Singapore and industry partners introduced Safeguards for Agentic Finance at Runtime (SAFR). The framework moves governance closer to the point of action through checkpoints that verify and record proposed agent actions before execution.^[48]

The security community is approaching the same problem through the OWASP Agentic AI Top 10, CSA's MAESTRO framework, and related work on risks created when agents gain access to tools, systems, and other agents.^[26]

These initiatives do not create one unified regime, but they point in the same direction: governance is moving from high-level principles toward operational controls. Human accountability remains the institutional obligation; runtime safeguards are beginning to provide part of the machinery.

Read together, these developments show the distance between financial decision and financial action collapsing while governance moves closer to execution. The remainder of this report maps the system now assembling and the questions institutions should ask.

SECTION 04

Defining Autonomous Finance

Most discussion of AI in financial services is a discussion about tools. This report proceeds from a different premise.

Autonomous Finance is the rise of non-human financial agency.

Finance has always assumed a human being at the point of initiation. Someone with a name, a signature, a legal identity, and the capacity to be held answerable. Every layer of the system was built on that assumption, from the account application and the identity check to the contract, the liability chain, and the court that enforces it. The assumption runs so deep that it is rarely stated. What is changing is not the sophistication of the tools institutions use. It is the identity of the actor initiating the transaction.

Building on the definition introduced in my earlier KYS: Know Your Swarm research, this report broadens Autonomous Finance beyond the multi-agent systems that formed the governance focus of that work and defines the category at the level of financial agency itself.

The definition

DEFINITION · AUTONOMOUS FINANCE

Autonomous Finance is the evolution of financial systems in which financial agency extends beyond direct human operation to autonomous systems, whether software, embodied, or collective, that can interpret intent or conditions, make or coordinate decisions, and initiate financial actions across digital, monetary, market, and public infrastructure, with varying degrees of human authorization, supervision, and accountability.

CHAUDHRY, 2026

Four elements of the definition carry the weight.

Agency, not automation. The distinction is not between manual and automated processes. Finance has been automating for fifty years. The distinction is between a system that executes a decision a human has made and a system that interprets conditions and initiates the action itself. Algorithmic trading automates a human strategy. An agent that reads a customer's circumstances, forms a recommendation, and acts on it is exercising something closer to agency, and agency is what the trust apparatus of finance was built to govern.

Software, embodied, or collective. Autonomous financial agency is not confined to a single agent with a payment credential. It includes software agents, systems embedded in physical infrastructure and devices that transact on their own account, and collectives of interacting agents whose behavior is not reducible to any single component. The collective case is the one existing governance frameworks were least designed for, and it is where this report concentrates.

Across digital, monetary, market, and public infrastructure. The category is not a product line. It spans payment rails, programmable money, market infrastructure, and increasingly the public infrastructure through which citizens access services. A definition confined to commerce or trading would miss where the consequences land.

With varying degrees of human authorization, supervision, and accountability. Autonomy is a gradient, not a state, and the gradient is where the governance question lives. The useful question is never whether a system is autonomous. It is how much authority was delegated, by whom, within what limits, with what supervision, and who answers for the result.

That final clause is why the definition matters commercially rather than only conceptually. The accountability question is not an ethical addendum to Autonomous Finance. It is constitutive of what the category is.

Agentic Finance and Autonomous Finance

The two terms are closely related, but they describe different scopes.

Agentic Finance refers to the use of AI agents to perform, support, coordinate, or execute financial functions. An agent may advise a customer, initiate a payment, execute a trade, monitor risk, negotiate with another agent, or participate in a wider financial workflow. The focus is the use of agentic systems within finance: what the agent can do, whose authority it exercises, what mandate it holds, and within what limits it operates. The governance question is delegated agency.

Autonomous Finance describes the broader evolution of financial systems as agency, decision-making, coordination, and execution become increasingly distributed across humans, AI agents, autonomous machines, programmable money, digital assets, institutions, and connected infrastructure. It includes interactions between humans and agents, among agents, between machines and financial institutions, and across programmable economic infrastructure. Its implications extend beyond individual products or workflows to the structure of markets, institutions, labor, liquidity, economic participation, governance, and accountability. The governance question is how authority, oversight, accountability, and redress are preserved across a wider system of human and non-human actors.

Agentic Finance is therefore an enabling capability within Autonomous Finance; it is not the whole category. Within the six-layer stack presented in the next section, Agentic Finance is horizontal rather than separate. Agentic capabilities can operate across every layer: accessing tools and data in Context & Action; coordinating across systems in Coordination; participating in Commerce; initiating value transfer through Payments, Programmable Money & Settlement; operating under defined identities, mandates, and permissions in Identity & Authorization; and remaining subject to institutional oversight, assurance, accountability, and redress in Governance & Trust.

Autonomous Finance is the wider financial system that evolves as these capabilities interact with people, institutions, programmable assets, digital infrastructure, and the physical economy. An institution can govern its agents well and still be exposed through the system in which they operate.

Agentic Finance is about the use of AI agents in finance. Autonomous Finance is about what finance becomes as financial agency is increasingly shared between human and non-human actors.

What Autonomous Finance is not

Autonomous Finance is not a chatbot that answers banking questions, a recommendation engine that never acts, traditional workflow automation with a new label, algorithmic trading alone, agentic checkout alone, stablecoins or blockchains alone, or one fully autonomous system replacing a financial institution.

Each of these may be an application, a rail, or a component. None is the category on its own.

The category begins where financial agency, the capacity to interpret conditions, make or coordinate decisions, and initiate financial action, ceases to be exclusively human. It becomes more consequential as that agency interacts across people, institutions, autonomous systems, programmable assets, and financial infrastructure, creating new questions of authority, oversight, accountability, and economic impact.

SECTION 05

The Map: Six Layers of the Autonomous Finance Stack

The moments discussed above look like separate headlines. They are not. Each one slots into a layer of the same emerging system, and the market becomes far easier to navigate once the layers are visible.

The six layers are:

- 01 Context & Action
- 02 Coordination
- 03 Commerce
- 04 Payments, Programmable Money & Settlement
- 05 Identity & Authorization
- 06 Governance & Trust

The lower layers expand what autonomous systems can do. The upper layers determine whether institutions can authorize, supervise, understand, and stand behind what they do.

Applications sit across the stack rather than inside one layer. Banking, payments, wallets, wealth, trading, treasury, DeFAI, commerce, insurance, and public-sector finance each draw on several layers at once. The same is true of crypto and blockchain infrastructure: stablecoins, wallets, onchain settlement, smart contracts, and tokenized assets appear across payments, identity, coordination, and applications rather than forming one isolated layer.

SECTION 06

Capability domains

LAYER	CORE QUESTION	SELECTED CAPABILITY DOMAINS
Governance & Trust	Can the institution supervise, evidence, account for, and stand behind outcomes?	Enterprise alignment, oversight, assurance, accountability, redress
Identity & Authorization	Who is acting, for whom, under what authority, and within what limits?	Agent identity, principal identity, mandates, credentials, permissions
Payments, Programmable Money & Settlement	How does value move securely inside an agentic workflow?	Agent payments, stablecoins, tokenized money, multi-rail settlement
Commerce	How do agents discover, compare, and complete purchases?	Machine-readable catalogs, carts, checkout, merchant controls
Coordination	How do agents discover, communicate, delegate, and collaborate?	A2A protocols, orchestration, agent-to-agent agreements, multi-agent workflows
Context & Action	How do agents access the tools, data, and systems required to act?	MCP, enterprise connectors, runtimes, sandboxes, tool use

Layer 1: Context & Action

Autonomous Finance begins when an AI system can do more than generate text. The Model Context Protocol has become a widely adopted interoperability layer connecting AI applications to external tools and data, while managed agent platforms increasingly ship persistent sessions, tool permissions, sandboxed execution, policy enforcement, tracing, and human approval patterns as standard equipment.^[1]

For financial institutions, the key point is simple:

Tool access is not merely a technical integration. It is a delegation decision.

An agent with read-only access to a knowledge base creates a different risk profile from one that can update customer records, submit a payment, place a trade, alter a limit, or initiate a treasury action. Controls should reflect what the system can do, not only what model it uses. Institutions need a clear inventory of what each agent can access, what it can change, which actions require approval, and which systems demand stronger controls.

EXECUTIVE TAKEAWAY

the risk boundary moves with the action boundary.

Layer 2: Coordination

Financial work is already distributed across specialist functions: onboarding, identity, fraud, compliance, payments, liquidity, treasury, advice, execution, servicing. Agentic systems are beginning to mirror that division of labor. The Agent2Agent Protocol provides an open standard for agents to discover one another and collaborate across systems; in April 2026 the Linux Foundation reported support from more than 150 organizations and active enterprise production use.^[2] And as the ClawBank and Shodai example shows, coordination is beginning to produce formal agreements designed to carry legal and financial effect between agent-led entities.^[22]

The significance is not the number of agents. It is the changing shape of the workflow. One agent may gather information, another check eligibility, a third calculate options, a fourth execute, a fifth monitor. The institution gains speed and modularity, but the outcome can become harder to understand if responsibility, context, or authority moves across several systems. Not every multi-agent workflow is a "swarm," and not every interaction requires the same oversight. But material workflows should be understood as systems, not as collections of isolated tools.

EXECUTIVE TAKEAWAY

one agent may be governed as a tool. A network of agents must also be governed as a system.

Layer 3: Commerce

Agentic commerce changes the interface between demand and supply. ACP lets agents complete purchases while merchants retain control of orders, fulfilment, and payment relationships.^[3] UCP gives agents a shared language for transacting with merchants while preserving merchant-specific rules.^[4] Adyen Agentic adds a modular enterprise layer across feeds, carts, and payments so merchants integrate once and participate across evolving agent platforms.^[5] Intent-to-payment loops are closing in specific verticals: Mindtrip, with Sabre and PayPal, now completes travel search, comparison, and payment without leaving the conversation.^[24]

The broader shift is from a marketplace designed for people to browse toward one that machines can query, compare, and act within. That does not remove the merchant, the customer, or the financial institution. It changes the interface among them. When the visible checkout page becomes less central, institutions still need reliable signals of customer intent, merchant legitimacy, authorization, fraud risk, and post-transaction support.

EXECUTIVE TAKEAWAY

the interface may become invisible. Institutional responsibility does not.

EXHIBIT 3 · THE AUTONOMOUS FINANCE STACK

SIX LAYERS

AGENTS ACT

04 Payments, programmable money and settlement

How value moves inside an agentic workflow

03 Commerce

How agents discover, compare and complete purchases

02 Coordination

How agents discover, delegate and collaborate

01 Context and action

How agents access the tools and data required to act

----- THE ACCOUNTABILITY LINE -----

INSTITUTIONS ANSWER

05 Identity and authorization

Who is acting, for whom, under what authority

06 Governance and trust

Can the institution stand behind the outcome

THE SIX-LAYER AUTONOMOUS FINANCE
STACKTHE STATE OF AUTONOMOUS FINANCE 2026 · AMNA USMAN
CHAUDHRY

Layer 4: Payments, Programmable Money & Settlement

Payments are where agentic systems become financially consequential. x402 embeds payments directly into web interactions and has passed 100 million processed payments.^[6] Agent Pay for Machines enables verified agents to make high-frequency payments across cards, accounts, and stablecoin rails.^[7] Intelligent Commerce Connect provides a single integration across networks, protocols, and token vaults.^[8] The Agentic Mobile Protocol brings agent payments, nano-settlement, and built-in Know Your Agent to mobile wallets and super apps.^[25] Coinbase for Agents connects AI agents directly to accounts within user-defined limits.^[13]

Programmable settlement is moving toward the institutional core: Circle's national trust bank approval, the Bank of England's systemic stablecoin framework, and in the UAE, DDSC.^{[14][15][20][21]} The Central Bank of the UAE has also completed Project Aperta, the BIS-led initiative for secure, interoperable cross-border open finance.^[54] While not an agentic-finance initiative, its infrastructure could support future agent-enabled financial workflows.

These developments do not mean machine-initiated payments have reached mass adoption. They show that payment capability is becoming something an agent can call inside a workflow rather than

a destination a person must visit. The institution still needs to know who authorized the action, what limits applied, what evidence is retained, how exceptions are handled, and who is responsible when an outcome is disputed.

EXECUTIVE TAKEAWAY

faster execution increases the value of well-designed authority, controls, and evidence.

Layer 5: Identity & Authorization

As agents gain the ability to act, identity and authorization become distinct from payment capability. The questions are straightforward: which agent is acting; for whom; under what delegated authority; with what scope, duration, and limits; and can the receiving party verify those claims?

The institutional case is now being made at the highest level. The IMF has identified Know Your Agent as an emerging regulatory consideration, calling on regulators to examine a shift from Know Your Customer toward Know Your Agent requirements, including verifiable identities for financial agents linked to legal entities.^[47]

The market is beginning to answer. Ant International's Agentic Mobile Protocol ships with a full Know Your Agent framework and a dynamic agent trust rating, an early production implementation rather than a proposal.^[25] Visa's Trusted Agent Protocol, developed with Cloudflare, gives merchants a way to distinguish verified agents from malicious bots at the point of checkout.^[8] Mastercard's Agent Pay for Machines credentials agents to operate within defined permissions.^[7] And SAFR, an agentic-finance safeguards framework published by MAS with industry participation, brings agent identity into a runtime control architecture that applies controls to each agent action independently.^[48]

Identity infrastructure at national scale exists, though built for a different subject. The Central Bank of the UAE's consent-based eKYC platform establishes verified identity for customers across the financial system.^[49] Read together, the three layers are visible: eKYC establishes who the customer is, Know Your Agent extends identity to the individual agent, and the layer above them, where several agents interact to produce a single outcome, remains open.

Identity is necessary, but it is not the whole governance problem. A payment credential proves an instrument can be used. An agent identity establishes who is acting. Authorization defines what the agent may do. Institutions must also understand how authority is used across a workflow, how material decisions are produced, and how accountability is preserved when several agents contribute to an outcome.

EXECUTIVE TAKEAWAY

capability answers "can it act?" Identity and authorization answer "should this action be accepted?"

Layer 6: Governance & Trust

Governance is not a final approval step. It is the operating environment that allows useful AI to scale.

The FSB's twelve sound practices address organization-wide governance and the AI lifecycle for boards and senior management.^[9] IOSCO's toolkit equips supervisors for AI in capital markets.^[10]

Singapore's framework emphasizes meaningful human accountability and the responsibilities of deployers.^[11] The CBUAE's guidance addresses responsible adoption, consumer protection, accountability, explainability, human oversight, and privacy.^[12] The security industry converges from below through OWASP, MAESTRO, and NIST.^[26] Santander has open-sourced its own governance tooling, an early signal that some banks see baseline governance as shared infrastructure rather than competitive edge.^[27]

There is a precedent for the structural problem, and it is recent. In June 2026 the IMF published new guidance on BigTech in financial services, concluding that activity-based regulation "assumes that activities generate the same risks and entities do not pose systemic risks," and that BigTech's combination of services is challenging those assumptions.^[50] The Fund's diagnosis is that an approach focused solely on individual activities can underplay risks that accumulate across a group, and its remedy is to expand the regulatory perimeter and move toward group-wide supervision. The parallel to multi-agent systems is my own, not the Fund's, but it is difficult to miss: governing each agent action independently rests on the same assumption the IMF has found wanting in the BigTech context, that risk assessed one activity at a time captures risk produced by activities in combination.

Do not build an agentic AI governance island. Agentic AI introduces new capabilities and new questions. It does not create a separate institution. New use cases should connect to existing structures: board accountability, business ownership, enterprise risk, model risk, cybersecurity, data governance, conduct, compliance, third-party risk, audit, resilience, and redress. Existing governance may need to evolve; the principle is that evolution should remain connected to the institution's accountabilities, risk appetite, and customer obligations.

EXECUTIVE TAKEAWAY

governance should enable responsible scale, not become a check at the end.

SECTION 07

The Case for Waiting

What the skeptics get right

A report that only presents the case for movement is advocacy, not analysis. The case for caution is strong and any institution weighing this should hear it in full.

The volume is not there yet. x402 has processed more than 100 million payments, but measured monthly transaction volume remains modest against any meaningful benchmark in payments. Infrastructure announcements are not adoption.

The market's own builders rate it as immature. Adyen's Global Head of Agentic Commerce places the field's current maturity at roughly 0.5 on a five-point scale, expecting about 1.5 within a year, and notes that the constraint was never the AI but the underlying plumbing of catalogs, trust, and fraud.

[23]

This has failed once already. Autonomous checkout cooled earlier in 2026 after a major retailer found that checkout-in-chat cratered conversion, and the experiment was withdrawn. The agents could complete the purchase. What stalled adoption was trust: the reliability and answerability that customers and institutions require before handing real money to a process they cannot fully supervise.

Most institutions are not treating this as strategic. Fifty-two percent of financial services firms are piloting or deploying agentic AI, yet only 14 percent see AI as transformational to organizational strategy and competitive advantage, and four of the top five use cases remain back-office functions.

[42] Current deployment is concentrated in internal operations rather than business model change.

The regulatory position is unsettled. The revised U.S. interagency model-risk guidance of April 2026 states expressly that generative and agentic AI are novel, rapidly evolving, and not within its scope, directing institutions to apply existing risk-management practices instead. Institutions are being asked to govern agentic AI before a prescriptive framework exists.

Taken together, the skeptical case is that the infrastructure is running well ahead of demonstrated economic value, that early experiments have already failed on trust rather than capability, and that a patient institution may let others absorb the cost of learning.

That case is largely correct on the facts, and it leads to a different conclusion than it appears to. If adoption were already at scale, governance would be a retrofit problem: expensive, contested, and late. Because adoption is early, authority, controls, and evidence can still be designed into the operating model while the cost of doing so is low and the number of live agents is small. The window the skeptics correctly identify is not only a reason to wait on deployment. It is the reason to move now on governance.

SECTION 08

The KPI Is Not the Number of AI Agents

AI agent counts may indicate experimentation. They do not prove transformation, customer value, resilience, or return on investment.

The better question is not how many AI agents an institution has deployed, but which outcomes are improving and what level of autonomy is actually required to achieve them.

The measures that matter are the ones an institution would want to hold up under scrutiny. Tasks automated is a count; outcomes improved is a result. Pilots launched is an intention; use cases in production is a commitment. And an incident log that is clean so far says far less than the ability to demonstrate that a workflow operated within its approved boundaries.

Autonomy is a design choice, not a KPI. The goal is not more autonomy. It is better outcomes through appropriate autonomy.

SECTION 09

Safe Adoption as an Operating Model

Safe adoption means creating value while keeping risk understood, proportionate, bounded, evidenced, and institutionally owned. Six elements:

- 01 **Start with the outcome.** Define the customer, operational, risk, inclusion, or financial value before deciding how autonomous the system should be.
- 02 **Match autonomy to materiality.** The appropriate degree of human involvement should reflect the impact, reversibility, complexity, customer consequence, and regulatory sensitivity of the use case.
- 03 **Extend existing governance.** Map the use case into current accountabilities, policies, risk classifications, committees, assurance processes, and incident pathways. Adapt where necessary, but avoid creating a disconnected governance structure.
- 04 **Make authority explicit.** The institution should be able to determine who or what authorized an action, what scope applied, and when authority should expire, be reduced, or be withdrawn.
- 05 **Preserve evidence.** Retain enough information to supervise performance, investigate material events, support audit and assurance, and provide meaningful redress.
- 06 **Evolve with the use case.** Review controls as AI agents gain new tools, interact with new systems, operate across new jurisdictions, or move into more consequential decisions.

The cycle is:

Responsible scale begins with value, matches autonomy to the use case, aligns deployment with existing institutional governance, applies proportionate controls, preserves evidence, and learns before expanding.

This is how governance becomes an enabler of adoption and ROI rather than a check at the end.

SECTION 10

Beyond Individual Agent Governance

From KYA: Know Your Agent to KYS: Know Your Swarm


Know Your Agent asks foundational questions: which agent is acting; who or what does it represent; what authority does it hold; what permissions and limits apply; can those claims be verified?

As workflows become more connected, a second set of questions emerges: which agents contributed to a material outcome; how did context, responsibility, or authority move across the workflow; can the institution understand material interactions and handoffs; can it determine where oversight, escalation, or accountability should apply?

KYS: Know Your Swarm is a proposed governance lens for this next layer: extending accountability from individual agent identity toward interactions and collective outcomes in multi-agent systems.

KYS is not intended to replace KYA, nor to create a governance regime separate from the institution. It is better understood as an extension test for existing governance:

Can the institution's current accountability, risk, audit, and redress frameworks still answer the necessary questions when more than one autonomous agent contributes to an outcome?



This remains one of the least resolved governance questions in Autonomous Finance. Existing frameworks largely address individual systems, agents, and actions; they do not yet provide a widely established approach to accountability for outcomes produced across interacting agents. Institutions deploying multiple agents into consequential workflows may therefore find themselves operating ahead of specific guidance and relying on existing governance principles to address distributed action.

This is the swarm accountability gap: the space between governance frameworks built for individual agents and outcomes produced by agents in interaction. I introduced the term in *The Swarm Accountability Gap in Autonomous Finance* (SSRN, 2026), and it is the gap the seven questions in the next section are designed to surface.

The objective is not to govern every interaction equally, nor to assume every multi-agent system carries the same risk. It is to preserve proportionate visibility and accountability for the interactions that matter, and to be able to demonstrate it. Identity establishes who is acting. Accountability must survive the moment action becomes distributed.

SECTION 11

Seven Questions Every Financial Institution Should Ask

These questions follow directly from accountability obligations already in force, and they will arrive from somewhere: a board, a regulator, an auditor, a corporate client's procurement team, or a customer contesting an outcome. Most institutions cannot answer them today. Eighteen percent are confident they could pass an independent audit of their AI controls.^[38]

01 What measurable outcome does this improve, and what level of autonomy does that actually require?

Start with value rather than capability. Maximum autonomy is not automatically better; the right level is the least independence needed to produce the outcome without exceeding risk appetite. Autonomy is a design choice, not an achievement.

02 Which AI agent is acting, for whom, and within what authority?

Identity, principal, mandate, permissions, limits, duration, and revocation, clear enough for the use case and verifiable by the party on the other side of the action.

03 Where does this sit inside our existing governance, and what did we have to change?

Determine what extends, what adapts, and where genuinely new capability is required. An answer that describes a separate agentic AI governance structure is a warning sign, not a plan.

04 What evidence would show the system operated within approved boundaries?

Consider what a business owner, a risk function, an auditor, a supervisor, and an affected customer would each need to understand after the fact. Evidence that satisfies engineers may not satisfy any of them.

05 When multiple AI agents contribute to an outcome, can we reconstruct how it was produced?

Which agents participated, how context and authority moved between them, which interaction produced the result. This is a question current frameworks do not yet address consistently, and it becomes unavoidable when a material outcome is contested.

06 Where can a human meaningfully intervene, and would that intervention actually work?

Oversight is not the presence of a person. It is whether that person has enough information, enough time, and enough authority to change the outcome. A human who approves without the capacity to refuse is a control on paper only.

07 When it goes wrong, who owns the outcome, the incident, and the customer's remedy?

By name, by function, by institution. Responsibility should remain locatable across business units, control functions, vendors, platforms, and third-party agents. "The system did it" is not an answer any supervisor will accept, and it is not an answer any customer can appeal.

These questions are deliberately technology-agnostic. They apply across banking, payments, wealth, trading, wallets, treasury, commerce, insurance, and public-sector finance, and they remain valid as models, protocols, and vendors change.

An institution that can answer all seven is not slower than its competitors. It is the one that can deploy into consequential workflows at all.

SECTION 12

The UAE Advantage

The jurisdiction where deployment and accountability are advancing together

The opening of this report set out the conditions. This section sets out why the same conditions that create the pressure also create the advantage.

The UAE combines three conditions that rarely appear at the same speed anywhere in the world: deployment ambition backed by national mandate, institutional and regulatory capacity with accountability principles already in force, and a stated commitment to responsible, human-centered adoption. In June 2026, the Artificial Intelligence and Data Authority was established as a single national body, consolidating the Artificial Intelligence Office, the digital government function of the TDRA, and the Emirates Data Office, and reporting directly to Cabinet.^[19] Dubai Chambers has formed an Executive Committee for Agentic AI to drive private-sector adoption, giving the mandate institutional machinery to match the target.^[53]

The regulatory capacity is layered rather than centralized. The CBUAE supervises at the federal level, VARA oversees virtual assets in Dubai, and the DFSA and FSRA supervise the DIFC and ADGM financial free zones respectively. For institutions, that means several regulated environments, each with established innovation testing pathways, in which agentic use cases can be developed under supervision while the wider rules are still forming.

The financial infrastructure is moving in step. DDSC brings central-bank-cleared programmable money explicitly designed for machine-to-machine transactions.^{[20][21]} Standard Chartered's integrated USDC minting and redemption capability launched first out of DIFC.^[36] Inception, part of G42, is expanding sovereign enterprise agentic AI infrastructure in the UAE, including through the launch of InceptionClaw.^[44] Microsoft has separately committed US\$15.2 billion in investment to the UAE between 2023 and 2029.^[46]

The country's banks are moving with the state. Emirates NBD ranked first in the inaugural Evident AI Index for banks in the Middle East and Africa and launched a partnership with Techstars targeting agentic finance startups.^[34] First Abu Dhabi Bank was recognized in the same index for agentic AI already operating inside core banking, including payment processing and relationship management.^[35]

The strongest opportunity will be the ability to develop use cases and governance together. The swarm accountability gap does not yet have a widely established regulatory answer; the rules are still forming globally. A jurisdiction where deployment is fastest, and where the accountability principle arrived before the agents did, is positioned to deploy agentic systems and to help shape how Autonomous Finance should be governed in practice. Where infrastructure, institutions, regulators, and market participants can learn in parallel, safe adoption becomes a competitive capability instead of a compliance afterthought.

SECTION 13

What to Watch Next

Five questions for the next 12 to 18 months

- 01 Will AI agent enabled activity convert into meaningful economic volume? Infrastructure is moving faster than mass adoption. The market will begin separating durable, high-value use cases from demonstrations and novelty.
- 02 Which autonomy models will customers and institutions trust? The winning design may not be maximum autonomy. It may be autonomy that is contextual, transparent, bounded, easy to understand, and easy to revoke.
- 03 Where will responsibility, liability and redress settle? Agent builders, model providers, platforms, financial institutions, merchants, and users may all contribute to a transaction or outcome. Real adoption requires responsibility to remain clear enough for supervision, dispute resolution, and customer trust.
- 04 Will open protocols create interoperability, or new complexity? MCP, A2A, ACP, UCP, x402, and emerging identity standards can reduce fragmentation. They can also create more connections, dependencies, and interactions to govern.
- 05 Can governance keep pace without becoming a barrier? The most effective frameworks will support experimentation, learning, and ROI while preserving accountability, consumer protection, and institutional resilience.

The institutions that move first may not be those that deploy the most agents. They may be those that learn fastest while maintaining the trust required to scale.

SECTION 14

Conclusion

The first half of 2026 made one point increasingly clear: artificial intelligence is no longer limited to generating financial insight. It is beginning to participate in financial action, under a national clock and a standing accountability obligation.

The strategic question is not how many AI agents an institution can deploy. It is whether the institution can identify use cases with real value, select an appropriate level of autonomy, align adoption with its existing governance, verify identity and authority, understand material interactions as workflows connect, retain meaningful human accountability, demonstrate performance and ROI, and stand behind the outcomes that follow.

The technology will continue to change. The institutional questions will remain recognizable: Who is acting? Under whose authority? Toward what outcome? Within which limits? With what evidence? And who is accountable when the result affects a customer, a market, or the financial system?

Those questions do not slow innovation. Answered well, they are what make responsible scale possible, at the speed this market now demands.

Autonomous Finance will not scale on capability alone. It will scale where value, trust, and accountability advance together.

The future is Autonomous Finance. The perspective must remain human.

THE AUTHOR



Amna Usman Chaudhry

Financial Economist & Frontier
Finance Strategist

LinkedIn [linkedin.com/in/](https://www.linkedin.com/in/amnausmanchaudhry)

[amnausmanchaudhry](https://www.linkedin.com/in/amnausmanchaudhry)

X [@Amna_uc](https://twitter.com/Amna_uc)

Autonomous Finance

autonomousfinance.substack.com

About the Author

Amna Usman Chaudhry is a financial economist and frontier technology strategist defining Autonomous Finance, the structural shift toward AI-coordinated financial systems.

She is the author of *KYS: Know Your Swarm and The Swarm Accountability Gap in Autonomous Finance* (SSRN, 2026), which introduce a governance framework for multi-agent AI in regulated finance and the concept of the swarm accountability gap. Her work spans fintech, AI, blockchain, digital assets, financial inclusion and responsible innovation.

Amna is a graduate-level lecturer in frontier finance and emerging technologies, and serves on boards and advisory councils across fintech, technology and responsible innovation. She is a UNESCO Women4EthicalAI Expert and has twice been listed on the Innovate Finance Women in FinTech Powerlist.

She has spoken at and contributed to United Nations convenings and multi-stakeholder platforms on algorithmic accountability, digital assets, AI governance and the future of Autonomous Finance.

Research. Strategy. Dialogue.

The seven questions in this report are the ones I am putting to institutions across the market. If yours is working through them, I would like to hear how you are answering them.

ABOUT FINTECH TUESDAYS

Fintech Tuesdays connects the MENA fintech industry by bringing founders, financial institutions, payment companies, investors, regulators, and ecosystem leaders into the same room.

Founded in the UAE in 2019, FTT has grown from an informal gathering of fintech professionals into a regional community of more than 1,600 members, supported by a wider network of over 10,000 LinkedIn followers. Through Business Breakfasts, Soirées, bespoke events, and an active vetted WhatsApp community, FTT creates space for practical conversation, collaboration, opportunity, and access.

Its mission is to inspire, unite, and build one of the region's most impactful fintech communities through inclusion, connection, and shared momentum.

WEBSITE

fintechtuesdays.com

LINKEDIN

[linkedin.com/company/fintech-tuesdays](https://www.linkedin.com/company/fintech-tuesdays)

CONTACT

info@fintechtuesdays.com

APPENDIX A

Methodology and Scope

This is an executive synthesis rather than an exhaustive market census.

The report draws primarily on official publications from technology companies, payment networks, financial institutions, standard-setting bodies, regulators, and UAE government sources. It focuses on developments that materially advance one or more layers of the proposed Autonomous Finance stack.

The six-layer framework is intended as a practical analytical model. The boundaries between layers are not always rigid. A single product may operate across several layers, and some capability domains will evolve as standards and business models mature.

The report distinguishes among enabling infrastructure, live products and pilots, early market adoption, regulatory and supervisory guidance, and author analysis. Where the report makes an inference, it is presented as analysis rather than as a settled market fact.

APPENDIX B

Notes and Sources

[1] **Anthropic**. “Introducing the Model Context Protocol.” 25 November 2024; and “Code Execution with MCP.” 4 November 2025.

[2] **Linux Foundation**. “A2A Protocol Surpasses 150 Organizations, Lands in Major Cloud Platforms, and Sees Enterprise Production Use in First Year.” 9 April 2026.

[3] **OpenAI**. “Buy It in ChatGPT: Instant Checkout and the Agentic Commerce Protocol.” 29 September 2025.

[4] **Shopify**. “Agentic Commerce for Every Developer: The Spring '26 Edition.” 17 June 2026; and “Universal Commerce Protocol.”

[5] **Adyen**. “Adyen Announces Adyen Agentic as the Universal Translator for the Next Era of Commerce.” 16 June 2026.

[6] **Coinbase**. Q1 2026 Earnings Deck. 7 May 2026; and x402 official documentation.

[7] **Mastercard**. “Mastercard Launches Agent Pay for Machines to Unlock Super-Fast, Always-On Payments.” 10 June 2026.

[8] **Visa**. “Visa Opens the Door to AI-Driven Shopping for Businesses Worldwide.” 8 April 2026; and “Visa Introduces Trusted Agent Protocol: An Ecosystem-Led Framework for AI Commerce.” 2025.

[9] **Financial Stability Board**. Sound Practices for Responsible Adoption of Artificial Intelligence. Consultation Report. 10 June 2026.

[10] **International Organization of Securities Commissions**. Supervisory Toolkit for AI Use in Capital Markets. Final Report. May 2026.

[11] **Infocomm Media Development Authority of Singapore**. Model AI Governance Framework for Agentic AI. Updated 20 May 2026.

[12] **Central Bank of the UAE**. Guidance Note to Protect Consumers and Ensure Responsible Use of Artificial Intelligence in the Financial Sector. 23 February 2026.

[13] **Coinbase**. “Coinbase for Agents: Your AI Agent Can Now Trade and Pay with Coinbase.” 11 June 2026.

[14] **Circle**. “Circle Receives Final OCC Approval to Establish National Trust Bank.” 10 July 2026.

[15] **Bank of England**. Sterling-Denominated Systemic Stablecoins: Policy Statement and Draft Code of Practice. 22 June 2026.

[16] **Central Bank of the UAE**. Payment Token Services Regulation. Introduced June 2024; updated official CBUAE Rulebook materials accessed July 2026.

[17] **Robinhood**. “Robinhood Is Now Open to Agents.” 27 May 2026.

[18] **Emirates News Agency (WAM)**. “UAE Government Framework to Deploy Agentic AI Across 50% of Government Sectors and Operations Within Two Years.” 23 April 2026.

[19] **Emirates News Agency (WAM)**. “Mohammed bin Rashid Approves Establishing Artificial Intelligence and Data Authority.” 14 June 2026.

[20] **Emirates News Agency (WAM)**. “CBUAE Approves Launch of UAE Dirham-Backed Stablecoin DDSC on ADI Chain.” 11 February 2026.

[21] **Emirates News Agency (WAM)**. “DDSC Receives Central Bank Approval to Partner with Selected VARA-Regulated Exchange Platforms.” 3 July 2026.

[22] **ClawBank and Shodai**. Announcement concerning the reported first Ricardian contract negotiated and signed between autonomous AI agents. 18 June 2026.

[23] **PYMNTS**. “Adyen on Agentic Commerce: The AI Was Never the Hard Part.” Interview with Karan Katyal, Global Head of Agentic Commerce, Adyen. June 2026.

[24] **Mindtrip, Sabre, and PayPal**. Official announcement concerning agentic travel search, comparison, and checkout. 2026.

[25] **Ant International**. “Ant International Introduces Agentic Payment Framework for Mobile Interfaces.” Agentic Mobile Protocol. 28 April 2026.

[26] **OWASP**. Agentic AI Top 10. 2026; Cloud Security Alliance. MAESTRO. 2026; National Institute of Standards and Technology. AI Risk Management Framework; and NHI Management Group. “AI Agents: The New Attack Surface.” 2026.

[27] **Banco Santander**. Open-source AI governance and safety tooling release. GitHub. 2026.

[28] **Coinbase**. “Coinbase Advisor.” 16 June 2026.

[29] **Robinhood Newsroom**. “Robinhood Chain Mainnet Launch.” 1 July 2026; and “Agentic Trading for Crypto.” 10 July 2026.

[30] **Kraken**. Official announcement concerning AI-powered agentic trading. July 2026.

[31] **Franklin Templeton**. Official announcement concerning the launch of Franklin Crypto. 2026.

[32] **Banco Santander**. “Banco Santander and G42 Sign Memorandum of Understanding to Explore Strategic Cooperation in Artificial Intelligence.” 3 June 2026.

[33] **Emirates News Agency (WAM)**. Official announcement concerning Dubai’s executive plan for agentic AI adoption in the private sector. June 2026.

[34] **Emirates NBD**. “Emirates NBD Ranked #1 in Inaugural Evident AI Index for Banks – Middle East and Africa.” 2 June 2026.

[35] **The National**. “Emirates NBD and FAB Lead Responsible AI Index as Region’s Banks Close Gap with Global Peers.” 2 June 2026.

[36] **Standard Chartered**. Official announcement concerning integrated USDC minting and redemption access for institutional clients through DIFC. 2026.

[37] OvationCXM. 2026 Agentic AI Readiness Survey: When AI Agents Become the Customer. Survey of 522 U.S. finance leaders, fielded April 2026. Published 23 June 2026.

[38] Grant Thornton. 2026 AI Impact Survey Report: Banking Insights. Survey of 950 business executives. April 2026.

[39] Bank Director. 2026 Risk Survey. Sponsored by Baker Tilly. March 2026.

[40] Wolters Kluwer. Q1 2026 Banking Compliance AI Trend Report. Survey of 148 financial institutions. February 2026.

[41] FIS and Anthropic. Official materials concerning the deployment of an agentic financial-crimes solution; and Fiserv. Official announcement concerning agentOS. 2026.

[42] Cambridge Centre for Alternative Finance. The 2026 Global AI in Financial Services Report. Cambridge Judge Business School. May 2026.

[43] Central Bank of the UAE. Digital Dirham: A Universal Payment Instrument. July 2025.

[44] Inception. "Inception Launches InceptionClaw, the First Sovereign, Enterprise-Grade AI Super Assistant Built on the Catalyst Platform." G42, 5 May 2026.

[45] Mordor Intelligence. Autonomous Finance Market: Size, Share, and 2026–2031 Growth Trends Report. Industry Analysis Update, January 2026.

[46] Microsoft. "Microsoft's US\$15.2 Billion Investment in the UAE." November 2025.

[47] Davidovic, Sonja, and Hervé Tourpe. "How Agentic AI Will Reshape Payments." IMF Notes, 2026, No. 004. International Monetary Fund, April 2026.

[48] Monetary Authority of Singapore and industry partners. Safeguards for Agentic Finance at Runtime (SAFR). July 2026.

[49] Central Bank of the UAE. Official announcement and related materials concerning the consent-based eKYC platform. 2026.

[50] Bains, Parma, Gina Conde, Nobuyasu Sugimoto, and Chao Wu. "BigTech in Financial Services: Emerging Regulatory Considerations." IMF Technical Notes and Manuals, 2026/09. International Monetary Fund, 30 June 2026.

[51] Circle. "What the Agentic Economy Means for Banks" and related publications concerning financial infrastructure for the agentic economy. 2026.

[52] G42. "AI Agents Can Now Apply for Jobs at G42." 27 February 2026.

[53] Dubai Chambers. Official announcement establishing the Executive Committee for Agentic AI. 2026.

[54] Central Bank of the UAE and BIS Innovation Hub. "CBUAE Marks Financial Innovation Milestone with 'Aperta' Project." 17 June 2026.

Foundational Research and Commentary by the Author

Chaudhry, Amna Usman. KYS: Know Your Swarm – A Governance Framework for Multi-Agent AI Systems in Autonomous Finance. SSRN, 2026.

Chaudhry, Amna Usman. The Swarm Accountability Gap in Autonomous Finance: Delegation Failure and Institutional Responsibility in Agentic Finance and Autonomous Financial Systems. SSRN, 2026.

Chaudhry, Amna Usman. Autonomous Finance. Independent research and commentary publication, 2026–present. autonomousfinance.substack.com

© 2026 Amna Usman Chaudhry. All rights reserved.

Selected original frameworks and diagrams may be reproduced for non-commercial presentations and discussion with clear attribution to Amna Usman Chaudhry, The State of Autonomous Finance 2026. Commercial reproduction, modification, or redistribution requires permission.