

TENSORX × AMINA BANK × SOLANA FOUNDATION ×
APEX:E3 × CARDANO FOUNDATION

— INDUSTRY RESEARCH REPORT

AGENTIC FINANCE REPORT.

Capital that works around the clock. AI agents
deciding and acting inside mandates written by
people, on blockchain rails that settle in seconds.

CO-AUTHORS

TensorX · AMINA Bank · Solana Foundation ·
APEX:E3 · Cardano Foundation

GUEST ARTICLES

Blindsight · CV VC

LAUNCH VENUE

CV

SUMMIT

Kongresshaus Zurich
29–30 September 2026

Report Contents.

FOREWORD	Tim Grant — Executive Chairman, TensorX	P / 02
01 - SUMMARY	Executive Summary	
	What has actually changed, why it is structural rather than cyclical, and where the addressable opportunity sits.	P / 03
03 - FOUNDATION	The Regulated Banking Foundation	
	Why a regulated banking foundation is becoming the trust anchor for institutional agentic finance.	P / 12
05 - COMPLIANCE	Sovereign AI Infrastructure	
	Why residency is not sovereignty, why open-weight models are Europe's base layer for agentic finance, and what inference built for agentic loops requires.	P / 28
07 - IDENTITY	Verifiable Identity, the Trust Layer	
	The identity gap, the vLEI, and four jurisdictions arriving at the same requirement.	P / 39
GUEST ARTICLE - CAPITAL	Owning the Frontier Before It Is Priced	
	Why the AI-native transition is being built in private, and how early-stage venture holds exposure to it with discipline.	P / 46
02 - DEFINITION	What Is Agentic Finance?	
	From assistance to autonomy: the five agent classes, the loop and its governance, what the evidence on AI transformation says, and compliance as a property of the token itself.	P / 06
04 - INFRASTRUCTURE	Why Solana Is the Execution Layer	
	Why a single global state, settlement inside the decision cycle and sub-cent fees are what agents need, and how x402 lets software pay software.	P / 22
06 - INTELLIGENCE	ALICE — The Orchestration Layer	
	APEX:E3's multi-agent orchestration platform, the analytical reasoning layer that every execution agent consumes. Trusted by \$10T+ AUM.	P / 34
GUEST ARTICLE - SECURITY	The Attack That Never Breaks a Rule	
	Checking that what an agent proposes is what its principal actually asked for, from a check that sits outside the model.	P / 43
08 - ACTION	Conclusions & Recommendations	
	Six recommendations across three horizons, the six risks a committee will raise, and the evidence base. The window is open.	P / 49



TIM GRANT
EXECUTIVE CHAIRMAN, TENSORX

— FOREWORD

Build it so that you can answer for it.

A year ago, sovereignty was a specialist word. Today it is the first question institutional clients ask us. It did not arrive through principle. It arrived because firms putting AI into production walked into a wall made of data.

Our own route to it was ordinary enough. Through APEX:E3, which I also chair, we build AI systems for very large institutional asset managers. It became clear early that the binding constraint would be data rather than model quality: where it goes, who can see it, whether it is retained, and whether any of that can be demonstrated afterwards to somebody entitled to ask. For the institutions we serve, those are practical questions. Their policies, their regulators and their own clients require answers already. We could not deploy without them, so we built the answer. That is why TensorX exists.

I spent my career in financial services, and nothing about data control is new there. Retention, jurisdiction, auditability: the industry has been conditioned on all three for decades, long before anyone reached for the word sovereignty. What is new is AI. It took a moment to register that a model is somewhere your data goes, and that inference is processing like any other. Once that lands, the conclusion is the one the industry reached long ago about every other system it runs, and there is nothing radical in it.

There is a second lesson worth carrying across, and it is commercial rather than regulatory. Moving infrastructure to the cloud was, for a while, a real efficiency. Then the dependency matured, the prices moved, and the firms that had moved everything discovered how little leverage they had kept. I watched that happen from inside the industry. It would be careless to walk into the same arrangement a second time with something more consequential than servers: the models that do the reasoning.

That is the part I would ask readers to weigh. If your business depends on a model you do not control, then somebody else's commercial or regulatory decision is your continuity risk. A model can be repriced. It can be deprecated. It can become unavailable in your jurisdiction for reasons that have nothing to do with you. Sovereignty is a security question, an economic question and a strategic one at once, and only then a compliance matter.

None of which is an argument for caution. The firms getting the most out of AI are the ones who worked out what they could safely let it do, and then let it. What I would resist is the idea that AI is only an efficiency play, letting you do what you already do with fewer people. Where it matters it does something larger and less comfortable. It changes what the work is, and therefore what the organisation is for. That belongs on a board agenda, not in a technology budget.

Which brings me to this report. For years I have argued that digital assets and AI would eventually meet, and that the meeting point would be payments: machines paying machines, in the background, with nobody approving each one. That is no longer a forecast, and the pages that follow are an account of it: the regulated foundation that holds the assets, the rails that settle them, the identity layer that establishes who acted, the private inference that reasons, and the orchestration that ties them together, each described by the people building it.

I am an optimist about all of it, with one condition, the only instruction I have ever given our own engineers. Build it so that you can answer for it.

Tim Grant

EXECUTIVE CHAIRMAN, TENSORX · CHIEF EXECUTIVE, DEUS X CAPITAL · CHAIRMAN, APEX:E3

DUBLIN · SEPTEMBER 2026

01

CHAPTER ONE

Executive Summary.

Artificial intelligence, blockchain settlement and tokenised assets were three separate stories. They are now one system, and software can hold, allocate and move capital without waiting for a person. That makes the old machinery unnecessary rather than faster.

READING TIME

6 minutes

TAGS

Convergence · AI · RWAs · Settlement

AI and blockchain are merging.

For most of the last decade, artificial intelligence and blockchain were two separate stories with two different audiences. One was about models that could read, summarise and predict. The other was about ledgers that settle without an intermediary. They are now one story, and the point at which they meet is easy to state: software that can decide is being given the means to act.

What makes that matter to a financial institution is what the technology does to a constraint every institution is built around: a person has to sit in the loop. Treasury sweeps happen overnight because someone has to release them. Portfolios rebalance monthly because someone has to approve the trades. Collateral moves in batches because someone has to reconcile them. None of those cadences was chosen; each is the pace at which people can safely take responsibility for a decision.

An AI agent removes the person from the individual decision without removing accountability from the institution, provided the mandate is written down and enforced somewhere the agent cannot reach. That is the whole of the design problem, and most of this report is about how to solve it. The consequence is large. Once the decision can be made by software, there is no reason left for the batch, the overnight cycle or the monthly rebalance. Capital that used to wait can move when it should.

What does not change is worth stating just as plainly. The institution remains accountable for every transaction its software makes, the regulator asks the same questions and the audit trail has to be at least as good as a person's. Agentic Finance relaxes none of that. It moves the point of enforcement from a review after the fact to a boundary set before the agent acts: the mandate, the identity it carries, the rails it may use.

Blockchain settlement is what makes the second half possible. A model that can decide to move money still needs somewhere to move it: rails that are open at four in the morning, that settle in seconds rather than days and that charge so little that a payment of a fraction of a cent is as economical as one of ten million dollars. Public blockchains, Solana above all, are the first infrastructure to offer all three at once. Tokenisation completes the picture by putting the assets themselves, cash, money-market funds and collateral, into a form that software can hold and transfer directly.

This is already happening, at a scale that can be measured rather than forecast. x402, the open standard that lets software pay software, has been adopted by every major card network, Stripe, Shopify, Google and AWS.²⁸ In its first twelve months of measurement, on-chain agent settlement ran to roughly 176 million transactions, most of them for a few cents.⁸⁷ Bain expects AI agents to be handling between fifteen and twenty-five per cent of US e-commerce by 2030.⁸⁸ Those are consumer-side numbers; institutional finance will follow later and more carefully, but it will follow, because the economics are not close.

This report calls the result Agentic Finance: capital managed continuously rather than periodically, by software acting inside mandates that people write and can withdraw, on rails that settle at machine speed. It is the point at which the software that decides is also the software that acts. That is a different thing from AI applied to finance, which has existed for decades, and from a better interface on the old processes.

'Blockchain will be pivotal in allowing agentic AI to realize its potential for consumer transactions, and the growth of agentic AI is likely to become the "killer" use case that drives blockchain adoption.'

SANDY KAUL, HEAD OF DIGITAL ASSETS AND INNOVATION, FRANKLIN TEMPLETON · 21 JULY 2026²⁸

Why now, and what it is worth.

It would be reasonable to ask whether this is another cycle, a technology that arrives with great conviction and leaves a smaller footprint than promised. We do not think so, for a reason that has nothing to do with enthusiasm. Three conditions have to hold at once for autonomous capital management to be worth doing, and for the first time all three do.

The models are good enough: not perfect, not to be trusted without limits, but able to work across complex, fast-moving financial environments in ways that can be tested and bounded. The rails are ready: continuous, low-cost settlement that a program can use directly, with no opening hours and nobody to telephone. And the assets exist in the right form: tokenised money-market funds and stablecoins that are liquid, regulated and addressable by software, of which Franklin Templeton's on-chain fund is a prominent public example.

When only two of these held, automation was an experiment. When all three hold, it becomes the dominant strategy, and the reason is competitive rather than technical. An institution that runs treasury, allocation and collateral continuously will operate on a different clock from one that does not, with different unit economics and, in time, a different return profile. Institutions will adopt Agentic Finance because standing still, once a competitor has moved, stops being caution and becomes a structural disadvantage. Novelty has little to do with it. The asymmetry is what makes the timing awkward. Whoever moves first carries the cost of building governance that does not yet exist anywhere, and whoever follows gets that work cheaply, but arrives at a benchmark somebody else has already reset.

The opportunity needs no exaggeration. APEX:E3 puts the capital addressable over five years at more than \$4.8 trillion: corporate treasury, wealth-management cash, asset-manager buffers and collateral that sits idle for want of a decision-maker who never sleeps.³¹ On a \$500 million treasury, the modelled uplift from continuous intraday allocation is in the region of two hundred basis points, or roughly ten million dollars a year.³¹ Those are model outputs rather than results, and Chapter 08 is careful about what the evidence does not yet support. An executive does not need the third decimal place to see the shape of it.

The sensible place to begin is where the mandate is simplest and the cost of getting it wrong is smallest: idle cash. A treasury agent that sweeps surplus balances into a tokenised money-market fund and back again has a mandate that fits on a page, an outcome measured in basis points and a failure mode that costs basis points rather than a licence. That is where the institutions in this report have started, and where the recommendations suggest starting.

What follows is written by the people building the pieces: AMINA Bank on the regulated banking foundation, the Solana Foundation on settlement and the stablecoin economy already running on it, TensorX on why an institution has to own the inference its agents run on, APEX:E3 on the orchestration that turns many narrow agents into one accountable decision and the Cardano Foundation on how an agent proves whom it acts for. A guest contribution from Blindsight shows how to check that what an agent proposes is what its principal asked for, and one from CV VC sets out where the AI-native transition is being built and priced.

The report closes with six actions across three horizons. The first can start this quarter and requires no one to trust the technology yet, only to decide what an agent would be allowed to do.

The infrastructure is ready before the institutions are. [That is the ordinary shape of these transitions, and the reason the timing matters.](#)

DEFINITION · AGENTIC FINANCE

What is Agentic Finance?

Agentic Finance is what you get when AI agents stop advising and start acting: software that watches the market, reasons about it in real time and executes on-chain, inside mandates that people define but without a person approving each transaction.



CHAPTER AUTHOR

Marcus Maute

TensorX Swiss Representative

Digital infrastructure operator and transformation architect
· 20+ years at the intersection of financial services, blockchain and emerging technology.

ABOUT

Marcus Maute represents TensorX in Switzerland, working at the intersection of sovereign AI, digital-asset infrastructure and regulated finance. With more than two decades across financial-services transformation and institutional DeFi, he focuses on making agent-native financial infrastructure deployable within regulated environments. He is based in Zürich, which has become the densest concentration of applied AI research in Europe, with Google's largest engineering centre outside the United States; research or engineering operations for OpenAI, Anthropic, Microsoft, NVIDIA, Meta and Apple; and ETH Zürich and the University of Zürich supplying much of the talent behind them.

An assistant proposes. An agent acts.

The word agent is carrying a great deal of weight at the moment, and it is worth separating the two things it is asked to mean. Most of what is marketed as an AI agent is an assistant: it drafts, retrieves, summarises and recommends, then hands the result to a person who decides. That is genuinely useful, and it is not the subject of this chapter. An agent in the sense used here has been given a mandate and the authority to act inside it. It makes the transaction itself, rather than proposing it.

The distinction sounds semantic until you look at what has to change underneath it. Today's financial plumbing was built around the pace of human decisions: overnight batches, periodic rebalancing, settlement somebody initiates, compliance reviewed after the fact. Each is a sensible response to the same constraint: a person sits in the loop, and people are expensive, slow and asleep for a third of the day.

Agentic Finance inverts each of them. Decisions become continuous rather than periodic, because there is no longer a reason to batch them. Execution becomes immediate, because there is nobody to wait for. And compliance moves in front of the trade rather than behind it, because retrospective review stops meaning anything at machine speed. Reviewing yesterday's transactions is a control when a person made a dozen of them. It is a formality when software made ten thousand.

The economics change too, and this is the part most often underestimated. A card payment costs two to three per cent plus roughly thirty cents; an agent buying a second of compute pays a tenth of a cent.²⁸ On the rails agents use, a payment of a hundredth of a cent is as economical as one of ten million dollars. The consequence is that a class of transaction that was uneconomic, per-second metering, per-call payment, continuous small rebalancing, becomes ordinary. Cheaper payments are the least of it.

This is roughly where the rest of the field has arrived, by different routes. The American Banker on-chain glossary defines agentic finance as autonomous agents 'executing multi-step financial operations' within on-chain environments, exercising 'judgment within defined guardrails'.⁸⁴ ARK Investment Management, writing about convergence rather than finance in particular, describes smart contracts and stablecoins supporting 'a global digital monetary ecosystem, allowing AI agents to coordinate and direct real-world resources'.⁸⁵ The academic treatment converges on the same ground: the central design question, in Gong's formulation, is bounded autonomy: how to widen what agents may safely do without making markets more opaque, more fragile or less accountable.⁸⁵

The common thread is worth stating plainly, because it is the thread this report follows. Everyone arrives at the same pair of requirements: the agent has to be able to act, and the boundary around that action has to be enforced somewhere the agent cannot reach. A boundary that exists only in a prompt is not a boundary.

So the definition at the front of this chapter is deliberately narrow. Agentic Finance is the point at which the software that decides is also the software that acts, on-chain, inside a mandate a person wrote and can withdraw. AI in finance has existed for decades, and a better interface onto the same processes would still leave a person deciding. What follows answers the question that forces: if the machine is going to act, what has to be true first?

Five classes, and they fail in different ways.

It helps to be concrete about what these agents are, because the phrase autonomous agent invites either too much imagination or too little. In institutional practice they are narrow. Each has a defined job, a mandate that bounds it and an effect on the balance sheet you can measure. Five classes cover most of what is being deployed today, and they are worth distinguishing for an unglamorous reason: they fail differently, and so they have to be governed differently. Grouping them under one word conceals exactly the differences a risk committee needs to see.

The five are not independent. The compliance agent gates all the others; the treasury and collateral agents compete for the same cash; the payments agent is the only one the outside world ever sees. In production they run as one orchestrated system, as Chapter 06 describes, and it is the orchestration rather than the individual agent that an institution governs.

The **treasury agent** is the one most institutions meet first. It watches liquidity across an enterprise or a fund and moves surplus cash into tokenised, yield-bearing instruments, typically a money-market fund such as Franklin Templeton's on-chain fund (FOBXX), and back again as obligations fall due, intraday rather than at the close. Its mandate is short: what counts as surplus, which instruments are eligible, how much may move in one step. On a \$500 million treasury, APEX:E3's modelling puts the uplift from running this continuously at roughly two hundred basis points.³¹ The number matters less than what produces it. Cash that used to wait for a person now moves when it should.

The **yield agent** goes a step further and manages exposure across tokenised fixed income and cash equivalents, balancing yield against duration, credit and liquidity inside limits a person has set. It is the treasury agent with a wider mandate and a correspondingly stricter one, because the ways it can be wrong are more expensive.

The **compliance agent** is different in kind. It does not move capital; it decides whether capital may move. Every instruction proposed by every other agent passes through it before execution, checked against regulation and mandate at the moment of the transaction rather than in a review afterwards. Its output is a yes or a no, and an immutable record of why. Of the five it is the one that can never be optional, and the one whose failure costs a licence rather than basis points.

The **collateral agent** manages what is pledged where: allocating, substituting and recalling collateral across prime brokerage, repo and lending relationships as prices and exposures move, instead of once a day when somebody gets to it. It is the least visible of the five and, in a stressed market, quite possibly the most valuable.

The **payments agent** is where Agentic Finance meets the rest of the economy. It executes settlement instructions on-chain: large institutional transfers, cross-border supplier payments, settlement between agents, compensation streamed by the second and micropayments at sizes no legacy rail can process economically. Stablecoins settle these flows natively, at the granularity agents need, and the x402 standard lets one piece of software pay another with no account relationship in between. Much of the commerce this makes possible, agent-to-agent services, per-use pricing for data and compute, programmable supply-chain finance, did not exist before for a simple reason: it could not be paid for.

What the classes have in common matters more than the list. Each has a mandate a person wrote, an effect the institution can measure and a failure mode the institution has priced. An agent that lacks any of the three is not ready for a balance sheet, however capable the model behind it.

'DeFi is not yield farming for humans. It is agent treasury infrastructure: lending, borrowing, swapping, and hedging at machine speed with no human in the loop.'

RAOUL PAL, 'DEFI WASN'T MEANT FOR YOU' (SUBSTACK, 26 AUGUST 2026) ⁰²

A continuous loop, not a sequence of handoffs.

Agentic Finance is one loop, run continuously, with no hand-offs between people, rather than a smarter decision tool bolted onto old processes. Described step by step it sounds sequential; in practice all six steps run at once, many times a second.

It begins with signal: continuous market and portfolio data, on-chain and off, arriving as it happens rather than in an end-of-day file. The agent reasons over that data: real-time inference, on infrastructure the institution controls, for reasons Chapter 05 sets out. From the reasoning comes a decision, which the mandate turns into executable instructions along the best available path. Before any instruction executes it is validated, checked against regulatory and internal constraints in front of the trade rather than behind it. Only then does it settle, atomically, on-chain, denominated in stablecoins. And every step leaves an immutable record, so the institution's view of its own activity is complete and current rather than reconstructed the next morning.

Compare it with what it replaces. A treasury team runs a version of the same loop by hand: it reads the cash position in the morning, decides, seeks approval, instructs a payment and reconciles the next day. The steps are the same; each is simply separated from the next by a queue, a person, a cut-off, a working day, and the queues rather than the steps are where the time and the idle capital go.

Two things distinguish this from an automated version of the old workflow. The first is that compliance sits inside the loop rather than after it, which is the only place it can sit at machine speed. The second is that the loop has no natural pause. Nothing waits for a batch window, a sign-off or a working day, so capital is never idle for want of a decision.

The record is what makes the rest defensible. Every instruction, the reasoning that produced it, the check it passed and the settlement that followed are written once and cannot be altered afterwards. A supervisor asking why a payment was made gets the answer the agent had at the time, not a reconstruction.

That raises the question everyone asks next: what are people for? Agentic Finance takes people out of routine execution, and only out of that. People set the mandate, the limits and the escalation points, and agents work inside them. What is given up is the bottleneck; the control stays.

In practice this settles into three tiers. At the **operational** tier, cash sweeps, rebalancing within mandate and routine payments, agents act fully autonomously, at machine speed, with no one in the loop. At the **tactical** tier, a significant allocation shift or a new counterparty, agents propose and a person confirms before execution follows. At the **strategic** tier, mandate design, the risk framework and participation in governance, the work is exclusively human, because it is the frame within which everything else operates.

Where the tier boundaries sit is the most consequential decision here, and it is a business one rather than a technical one. Set the operational tier too narrow and the programme produces nothing but proposals; set it too wide and the first incident is the last. The sensible practice is to start narrow, with a mandate that fits on a page, measure and widen the tier as the evidence accumulates, the discipline a risk committee already applies to a new trader.

The tiers are where the accountability lives, and a regulator will ask to see them. That is why the sequencing matters more than most programmes expect. **The technical deployment takes weeks; the governance architecture takes months.** Starting with the governance is the fast option, whatever it looks like from outside.

'An economy run by software agents must run on software money, software contracts, and software governance, or it cannot run at all.'

JEREMY ALLAIRE, THE AGENTIC ECONOMY (2026) · P. 64⁹¹

Most AI programmes fail on the organisation, not the model.

It is worth pausing on why so many AI programmes disappoint, because the evidence is now specific and Agentic Finance is exposed to the same failure. The pattern is consistent: the model works, the pilot impresses and the value never reaches the profit and loss account.

The largest survey of the field, McKinsey's annual State of AI, finds that of everything an organisation can do, redesigning the workflow around the technology has the single biggest effect on whether AI shows up in earnings, and that high performers are about three times as likely to have done it.¹¹ The mechanism is unglamorous: gains are absorbed by whatever bottleneck sits downstream of the model, and AI amplifies the organisation it lands in, dysfunctions included. The money that pays back goes on process, governance and the operating model, not licences.

The second finding is about where the technology can be trusted. In a controlled experiment with 758 consultants at BCG, people using AI on tasks inside its competence produced markedly better work, faster. On tasks just outside it, which looked no harder, they did worse than colleagues with no AI at all, because the output was plausible and wrong, and nothing told them which side of the line they were on.¹² The boundary is jagged and it moves, but the lesson does not: the value comes from knowing exactly what the system may be trusted to do, and enforcing that boundary somewhere other than the user's judgement. There is an uncomfortable corollary for pilots. A team reporting success has usually been working inside the frontier; the same tool in the same hands degrades the moment the work drifts outside it, and nothing in the pilot gave any warning.

That is what a mandate is. The three tiers on the previous page are a map of that jagged frontier drawn for one institution: autonomous where the agent is reliably inside its competence, confirmed by a person near the edge, reserved to people beyond it. Agentic Finance is then an unusually disciplined form of AI deployment: the boundary is written down and machine-enforced, not left to whoever is at the keyboard.

Three further results shape how a programme is run. Gains land unevenly: in the first large workplace study, customer-service staff became about fifteen per cent more productive on average, the least experienced by a third, the most experienced hardly at all.¹³ Self-reported productivity is unreliable: experienced developers in a 2025 trial were slower with AI while believing they had been faster; a repeat a year later with better tooling found a speed-up of similar size.¹⁴ And a randomised study of over a thousand managers found that work attributed to a named 'AI employee' was checked less carefully and owned less personally than the same work from a tool.¹⁵ Instrument real cycle times, not surveys; and keep the agents off the org chart, because they have no accountability to carry. A person does.

It also says who should lead. The competence that matters is the ability to write a mandate, set a boundary and read the evidence honestly, which a good risk committee already has; familiarity with models counts for less. That is why the institutions furthest ahead put the programme under the people who run the balance sheet, not those who run the technology. The loop in 2.3 is a workflow redesign with the governance built in. That is why it can work where a model bolted onto the existing process does not, and why the hard part is organisational before it is technical.

Buy the model and leave the process alone, and you get the pilot that impresses **and the P&L that does not move.**

The next layer: compliance as a property of the token itself.

Today an agent's transactions are checked by the application that sends them, every payment, every time. That works, and it is how the institutions in this report operate now. But it leaves open a question that is structural rather than technical: when one autonomous agent settles with another, whose KYC applies? An application can only vouch for its own side.

Solana's Token-2022 standard offers a different place to put the answer.²⁵ A token issued under it can carry a transfer hook: a program the token itself runs on every transfer, which can reject any transfer that fails its checks. Because the hook lives in the token's configuration rather than in an application wrapper, it cannot be bypassed. It applies equally to a wallet transfer, a smart-contract call and a DeFi route. A token set up this way can only ever be held by an eligible wallet. Compliance stops being a checkpoint an application performs and becomes a property the token carries with it.

The second half is the population of agents allowed to hold it. Each agent is verified once, registered on-chain through the Solana Attestation Service and bound to a legal entity through its vLEI, the verifiable form of the Legal Entity Identifier that Sandro Knöpfel describes in chapter 07, from page 039. Its attestation names the audited, version-pinned runtime it runs. Once an agent is inside that perimeter, every interaction with every other agent is compliant by construction, because the hook already guarantees the counterparty is verified. Membership in the set is the compliance. There is nothing to re-check and nothing to gate, payment by payment.

It is worth being clear about what this is. It is forward-looking: the architecture is being developed as the next layer of the stack this report describes, and designs, partners and instruments will be set out in a companion paper. It makes no claim that the identity problem is solved. Chapter 07 explains why the identity gap is the part settlement rails cannot close on their own, and why four jurisdictions have arrived at the same requirement independently. What the token-level design does is move the enforcement to the one place an agent cannot route around.

For an institution the practical consequence is that the compliance cost of agent-to-agent activity stops scaling with the number of transactions. Checking every payment against every rule is affordable when a person makes a dozen a day and ruinous when software makes ten thousand. Verifying the participants once, and letting the token refuse everything else, is the only version of the economics that works at machine speed.

None of this requires an institution to wait. The agents it deploys today can be registered and attested now, and the tokens it holds can be issued with hooks now; what the companion paper adds is a shared perimeter across institutions. Building the identity discipline first, one agent, one legal entity, one attestation, is what makes joining that perimeter later a configuration change rather than a programme.

Allaire puts the larger point well: 'The agentic economy is the onchain economy: AI supplies the labor, and the onchain substrate supplies the form in which money, decisions, coordination, and ownership are expressed.'⁶¹ If that is right, then the rules the economy runs on have to be expressed in the same substrate, in the money and the contracts themselves rather than in a layer of applications watching them.

Both sides of every transaction verified before it happens, **with no intermediary balance sheet in between.**

REGULATED BANKING · AMINA BANK

The regulated banking foundation.

Autonomous agents can perceive, reason and act — but institutional capital will only move through them on rails a regulated bank stands behind. A licensed banking foundation is becoming the trust anchor that makes agentic finance deployable at institutional scale.



CHAPTER AUTHOR

Markos Theologitis

Chief Technology Officer · AMINA Bank

25+ years in international financial-services technology · technology strategy, operational resilience and digital transformation.

ABOUT

Markos is the Chief Technology Officer of AMINA Bank, responsible for leading the bank's global technology organization, strategy, and digital transformation. With more than 25 years of international experience, he has a proven record of shaping technology strategy, strengthening operational resilience, and delivering impactful digital solutions across global financial markets. He began his career at Accenture as a management consultant, advising major banks on strategic transformation, mergers and acquisitions, and technology integration. Before joining AMINA Bank, Markos held several senior technology leadership roles within major European financial institutions, where he built and led high-performing international technology and product organizations. Markos holds a Master's in Chemical Engineering from the National Technical University of Athens (NTUA) and an MBA from the ALBA Graduate Business School.

Autonomy relocates responsibility; it does not retire it.

The agentic finance stack, particularly on the execution side, is increasingly taking shape. Public chains such as Solana and Ethereum provide deterministic settlement, tokenised money-market funds introduce a compliant yield leg, multi-agent platforms translate mandates into decisions, sovereign inference can keep reasoning jurisdiction-bound, and agentic operating systems are beginning to assemble these components into coherent workflows. Each of these advances answers a question about capability. The more important question now is whether these capabilities can be integrated into a reliable, regulated and economically viable financial system. In that sense, the execution stack is no longer purely conceptual, but it is also not complete yet.

In practice, these capabilities are moving into traditional financial workflows, and the pattern is visible at both ends of the market. Retail and wealth clients increasingly lean on AI for budgeting, portfolio construction, tax-aware rebalancing, and always-on financial guidance. Institutional deployments go further: AI agents monitor cash positions, rebalance portfolios, optimise collateral, initiate FX conversions, and allocate idle balances across approved yield-bearing instruments. The significance is not that these functions are new, but that increasingly they can be coordinated and executed by software rather than simply presented to a human decision-maker.

One of the fastest-growing uses, however, is **payments**. Agents are beginning to support the execution and routing of payments with increasing levels of autonomy, from large-value institutional transfers and cross-border supplier settlement to machine-to-machine transactions and per-second streaming, denominated in stablecoins and settled on-chain in seconds. Each of these is a way financial execution is increasingly being delegated to software, and each one raises the same question the moment real money moves through it.

A treasurer, a fund board, and a FINMA supervisor ask that question, and they ask it first. If an agent misreads a regime shift at three in the morning and rotates a nine-figure position into the wrong tokenised instrument, whose balance sheet absorbs the loss, and who explains it to the regulator on Monday? That is a liability question, not a capability question, and nothing upstream in the stack resolves it.

This is where regulated financial institutions belong in the stack. A regulated bank does not replace the execution decision engines; it integrates them into its custody, compliance, and payment infrastructure so that autonomous execution stays bounded by regulated controls. The implication is not that a regulated bank assumes accountability for the agent, but that the agentic finance stack itself needs to incorporate more regulated and compliance-enabled components. Regulated custodians, banks, payment providers, and compliance infrastructure can provide these control layers, helping ensure that autonomous execution operates within defined controls and accountability frameworks.

The reason is structural. An AI agent has no legal personality. It cannot hold a license, cannot be a party to a custody agreement, cannot be sanctioned, and cannot be a defendant. A public ledger gives you an immutable record of what happened; it does not give you someone answerable for it. Autonomy therefore pushes responsibility upward, to whoever authorised the mandate and whoever holds the assets. Somebody licensed has to occupy that position, or institutional capital does not enter.

The Cardano Foundation and Global Legal Entity Identifier Foundation (GLEIF) contribution solves one half of this by binding an agent to a legal entity through the verifiable Legal Entity Identifier (vLEI): it establishes on whose authority the agent acts. A bank supplies the other half. A vLEI with no regulated custodian behind it is an attestation with nothing to seize. A wallet with no supervised owner is counterparty risk with no recourse. Identity is necessary but, on its own, insufficient. Identity establishes on whose authority an agent acts. A regulated bank establishes against whose assets, within which limits, and under whose supervision.

Not every regulated bank can occupy this position. Most hold the license but not the on-chain capability. They can custody a bond and reconcile a wire, but they cannot hold the keys, settle on a public chain, or screen a smart-contract counterparty in real time. Institutions that have developed these capabilities natively are likely to hold a structural advantage over those attempting to extend legacy processes into blockchain-based markets. As autonomous systems become more active participants in financial markets, the ability to bridge regulated finance and on-chain infrastructure may become a critical component of the agentic finance stack.

Bounding authority so the boundary is enforced, not merely promised.

The useful design question in agentic finance is not whether an agent can act. It is how its authority is bound, and whether that boundary is enforced through controls embedded in the financial infrastructure through which the agent operates. This distinction matters because agentic banking does not inherently require a blockchain protocol. Agents can operate within conventional banking infrastructure, where authority can instead be bounded through account permissions, transaction limits, policy engines, approval workflows and other institutional controls.

Blockchain-based systems introduce an additional design possibility: some of these constraints can be enforced directly at the asset or protocol level. The market is already converging on part of the answer through Token-2022 transfer hooks and attestation-gated agents: compliance carried by the token itself rather than re-checked by each application. For regulated institutions, the opportunity is to provide the infrastructure and supervision around both models, while blockchain-based controls can provide an additional layer of programmable enforcement where appropriate. The regulated bank's control surface maps ordinary banking controls onto on-chain enforcement:

01 The client is the account holder, not the AI agent.

Onboarding, KYC, and AML attach to the legal entity of the client. The agent operates a mandate over assets the bank custodies, so every action it takes inherits a real, sanctionable counterparty.

02 Authority is expressed as enforceable constraints.

Allow only listed counterparties and venues, a permitted asset universe, per-transaction and cumulative value caps, and settlement only into whitelisted, bank-controlled wallets. Where the instrument supports transfer-hook gating, those limits live in the mint configuration and cannot be bypassed by a contract call or a DeFi route.

03 Oversight is pre-trade and continuous.

A compliance agent validates every proposed transaction against a rule set. The bank is the entity that owns that rule set, is examined on it, and carries the liability when it fails. That distinction is the whole point.

04 The mandate is revocable in real time.

Keys rotate, allowlists narrow, and an agent's authority can be withdrawn mid-session. Authority you cannot revoke instantly is not a mandate, it is an unsecured position.

05 Agent integrity is continuously verified.

Institutions assume an AI agent may hallucinate, be compromised through prompt injection or a software vulnerability, or simply drift outside its intended objective, and the controls are built for that assumption.

Make the agent's speed survivable — not the agent immune to compromise.

The AI agent remains under the client's control, so the regulated institution does not attempt to prevent prompt injection, model hallucinations, or other failures within the agent itself. Instead, the regulated layer provides deterministic controls and guardrails around what the agent is authorised to do. Transaction limits, whitelists, policy rules, approval thresholds and MPC-based signing can ensure that no single compromised component can independently move client assets outside the permitted parameters. Human approval can remain mandatory above defined limits.

The objective is therefore not to make the agent immune to compromise, but to ensure that a compromised or erroneous agent cannot by itself result in an unauthorised movement of client assets. None of this is designed to slow the agent down. It is designed to make its speed survivable. A risk committee will sign off on millisecond execution the moment that millisecond happens inside limits it has set in advance through a deliberate, accountable process, and can pull back at will. Scalable control is a banking function before it is a software one.

INDEPENDENT VALIDATION

Every critical action is validated independently of the model that proposed it:

DETERMINISTIC RULE CHECKS

Policy engines evaluate every proposed action.

PRE-TRADE SIMULATION

Previews the on-chain effect before anything is signed.

WALLET & COUNTERPARTY SCREENING

Against sanctions lists and allowlists.

ANOMALY DETECTION

Against the mandate's normal behaviour.

VALUE & VELOCITY CIRCUIT-BREAKERS

Halt activity the moment a threshold is breached.

The objective is therefore not to make the agent immune to compromise, but to ensure that a compromised or erroneous agent cannot by itself result in an unauthorised movement of client assets.

Machine-native rails move value; a balance sheet converts it into ownership.

The agentic architecture now reaching production is close to a closed loop. A stablecoin flows in and earns on idle balances; FOBXX, a tokenised money-market fund, holds the cash leg; Solana settles; and the audit trail lives on-chain. Closed loops are elegant. For a supervised institution they are also incomplete, because capital enters and leaves through the traditional system, and the two settlement worlds run on incompatible assumptions. The mismatch is worth stating precisely, because it is where a regulated institution earns its place:

01 Fiat settlement has a backstop; a public chain does not.

Fiat settlement has a central bank behind it and a lender of last resort. A public chain gives cryptographic finality but no par-redemption guarantee on a stablecoin and no backstop if an issuer breaks. Something has to stand between “settled on-chain” and “good funds a client can withdraw.”

02 Bank books can segregate client assets; wallets cannot.

Client assets on a bank’s books can be structurally segregated and, under Swiss law, kept outside the bank’s own insolvency estate. Assets in an unregulated wallet or with an unregulated venue are likely exposed to that counterparty’s solvency, nothing less.

03 Reporting duties attach to a licensed entity.

Reporting duties under AIFMD, UCITS mandates, MiFID best-execution, the FATF travel rule, and AML regimes attach to a licensed entity. A smart contract cannot discharge them.

Role of a regulated digital asset financial institution: custody of keys and assets with institutional key management (MPC rather than a single signer), so control of the position does not reduce to one point of failure; fiat rails and conversion, the regulated on-ramp and off-ramp between stablecoin-denominated settlement and bank money, with redemption at par backed by a balance sheet rather than a claim on a protocol treasury; safeguarding and segregation of client assets, including through the bank’s own resolution; and monitoring, reconciliation, and reporting that tie each on-chain action to a named account, satisfy travel-rule and AML obligations, and produce statements an auditor and a supervisor will accept.

The chain removes the correspondent bank from the payment path. It does not remove the bank from the ownership path.

Control is the precondition, not the point. The point is what runs once it is in place.

Agentic finance, bounded by a regulated bank, is a source of measurable value, and that value is the reason a board takes on the work of adopting it at all. The same controls that make autonomy survivable are what let an institution run its balance sheet harder and cleaner than any human desk can. The outcomes an institutional reader should expect are concrete:

01 Liquidity optimisation.

Idle cash stops sitting idle. An agent sweeps balances into FOBXX the moment they clear a threshold and pulls them back before they are needed, so the treasury carries the least drag consistent with its obligations.

02 Capital efficiency.

Collateral is posted, netted, and recalled continuously against live positions rather than in daily batches, releasing capital that batch cycles strand overnight.

03 Operational cost reduction.

Reconciliation, screening, and settlement that once consumed a back office run as deterministic checks on-chain, so headcount shifts from processing to oversight.

04 Faster settlement.

On-chain finality in seconds replaces multi-day float, freeing working capital that settlement lag would otherwise trap.

05 Continuous, 24/7 treasury management.

Markets and payments do not keep banking hours, and neither do an agent. Cash, collateral, and FX are managed across every time zone and weekend without a night shift.

None of these outcomes are hypothetical, and none are reachable without the control layer above them. Speed without limits is a liability; limits without speed are a spreadsheet. A regulated bank is what lets an institution achieve both — high speed and scale.

What autonomous treasury looks like when a real institution runs it.

Consider a corporate treasury team at a multinational as one example. Cash may be scattered across multiple entities and currencies. Today that cash is often reconciled after it moves, swept periodically, and invested when the team has time to approve it. Under an agentic mandate, the picture changes without the client retaining less control. The same model can extend beyond corporate treasury to other clients managing portfolios, liquidity or payments across multiple accounts.

The treasurer sets the envelope: a permitted universe of FOBXX and other products, counterparties and venues on an allowlist, a cap of, say, 150 million USD under management, and human sign-off above a defined ticket. Inside that envelope the agent runs continuously. It sweeps idle dollar balances into FOBXX overnight and over weekends, funds a supplier payment in stablecoins the moment an invoice clears, converts back to fiat through the bank’s rails when an entity needs local currency, and rebalances collateral as positions move. Every action is simulated before it is signed, screened against sanctions and the allowlist, and checked against value and velocity limits.

A regulated custodian safeguards the assets, enforces the limits at the protocol level, and produces the statement an auditor and a regulatory body will accept. The treasurer wakes to a treasury that worked all night, a full record of what it did and why, and the standing power to narrow the mandate or withdraw it entirely before breakfast. That is the distance between an experiment and a mandate a risk committee will actually sign.

THE ENVELOPE · AT A GLANCE		SET BY THE TREASURER · ENFORCED BY THE CUSTODIAN	
PERMITTED UNIVERSE	FOBXX and others.		
COUNTERPARTIES & VENUES	Allowlist only.		
CAP UNDER MANAGEMENT	~150 million USD.		
HUMAN SIGN-OFF	Above a defined ticket.		
CONTROLS ON EVERY ACTION	Simulated, screened, value- and velocity-checked.		

Four curves crossed at once, and 2026 is where they meet.

Agentic finance is not arriving because a single technology matured. It is arriving because four of them crossed at the same time, and the crossing is what turns the case from speculative into emergent.

01 AI reached operational reliability.

Models moved from drafting text to executing multi-step financial workflows with tool use, structured output, and error rates a control layer can measure, constrain and absorb. The reasoning leg is finally dependable enough to sit behind a mandate.

02 Financial assets became increasingly programmable.

Tokenised money-market funds, treasuries, and cash instruments now give agents access to programmable, yield-bearing assets, while traditional banking infrastructure continues to provide access to fiat balances, accounts and conventional financial instruments. The important shift is not that assets must move on-chain, but that more of the financial environment can be accessed and acted upon programmatically.

03 Financial settlement became increasingly machine accessible.

Stablecoins and other blockchain-based settlement assets provide always-on, programmable rails with on-chain finality, while traditional payment networks increasingly offer instant or near-instant settlement within their own regulatory and operational frameworks. Agents can therefore interact with both on-chain and conventional payment infrastructure, subject to the permissions and limits of each rail.

04 Institutions began to demand continuous operations.

Payments, collateral, and treasury increasingly run around the clock, and the gap between round-the-clock markets and business-hours operations has become a cost that boards want closed.

Together they turn autonomous finance from a demonstration into an operating model. The opportunity is therefore broader than putting traditional financial activity on-chain. It is to allow software to operate across both traditional and blockchain-based financial infrastructure, while regulated institutions provide the permissions, controls and connectivity that keep that autonomy within defined boundaries.

Together they turn autonomous finance from a demonstration into an operating model.

A widening mandate, not a switch.

The prevailing industry recommendation is right on direction and optimistic on cadence. “Starting now is not optional” reads well to a corporate treasurer with discretion to spend. It reads differently to a fund board carrying fiduciary duty and a regulator on the other end of the phone. Adoption will not be a switch anyone flips. It will be a mandate that widens, and the institutions that move fastest will be the ones that kept each step small enough to approve.

– 3.8 CLOSING ARGUMENT

The binding constraint on agentic finance in 2026 is not latency, model quality, or settlement cost. Those are solved or visibly solving. The harder challenge is building the controls and governance needed for institutions to delegate financial activity to autonomous systems without weakening existing standards of oversight.

Agentic finance will be built by two things that look opposed but are not: autonomous execution that moves at machine speed, and regulated accountability that answers for where it moved. Remove the autonomy and you are left with a slow bank. Remove the accountability and you are left with an experiment no fiduciary can fund. The institutions that win will hold both at once.

The most likely path is a phased approach, where each stage broadens the scope of activities that an agent is permitted to perform. The goal is not to prescribe a fixed path towards fully autonomous treasury management. It is to establish a controlled progression in which the scope of an agent’s authority can expand only where the relevant risk controls, monitoring, governance and escalation mechanisms are demonstrably adequate. The permissions remain bounded and revocable throughout. Every step is measurable, reversible, and defensible in front of a regulator. That is the line between a pilot that impresses a conference audience and a mandate that survives a risk committee.

This is why regulated infrastructure is not simply another layer added to the agentic finance stack. It provides an environment in which autonomous financial activity can operate in a regulated environment on a scale: custody of assets, defined transaction limits, compliance controls, payment and settlement infrastructure, monitoring, auditability and clear mechanisms for intervention. It is the layer that carries agentic finance from experimentation into production deployment. For institutions ready to make that move, the foundation is not something to wait for. It is here to build on now.

The most likely path is a phased approach, where each stage broadens the scope of activities that an agent is permitted to perform.

Disclaimer — Research and Educational Content.

This document has been prepared and co-published by AMINA Bank AG ("AMINA"), the Solana Foundation, TensorX, APEX:E3 and the Cardano Foundation together as named co-authors of this report (together, the "Co-Authors"). AMINA is a Swiss licensed bank and securities dealer with its head office and legal domicile in Switzerland. It is authorised and regulated by the Swiss Financial Market Supervisory Authority ("FINMA"). The report also includes guest contributions by Blindsight Technologies AG and CV VC AG (together, the "Guest Contributors"). The Guest Contributors are not Co-Authors. Each Guest Contributor is responsible for its own contribution, and the views expressed in a guest contribution are those of its authors.

This document is published solely for informational, research and educational purposes. It is not an advertisement, solicitation, offer, invitation, recommendation or inducement to buy or sell any financial instrument, digital asset, token, security or other investment, or to participate in any particular investment strategy, transaction, product or service. Nothing in this document constitutes investment advice, financial advice, legal advice, tax advice or any other form of professional advice.

This document may be made available by the Co-Authors, including as a freely downloadable PDF, through their respective websites, blogs, newsletters, social media accounts and other official communication channels, in each case subject to applicable law and regulation. Recipients may download, read and share this document for informational and non-commercial purposes, provided that the document is not modified, the Co-Authors are appropriately identified, and this disclaimer remains included in full. No part of this document may be altered, misrepresented, sold, incorporated into a commercial product, or used in a misleading context without the prior written consent of the relevant Co-Authors.

This document is not directed to, and is not intended for distribution to or use by, any person or entity that is a citizen or resident of, or located in, any jurisdiction where such distribution, publication, availability or use would be contrary to applicable law or regulation, or would subject any Co-Author to any registration, licensing or other regulatory requirement in that jurisdiction. Recipients are responsible for informing themselves about and observing any applicable legal or regulatory restrictions.

The information contained in this document is based on sources considered by the Co-Authors to be reliable at the time of publication, including publicly available information and information provided or reviewed by one or more of the Co-Authors or Guest Contributors. However, no representation or warranty, express or implied, is made as to the accuracy, completeness, reliability, timeliness or fitness for any particular purpose of the information contained in this document. The information is not intended to be a complete statement or summary of the matters discussed.

The Co-Authors do not undertake any obligation to update, amend or keep current the information contained in this document, whether as a result of new information, future events, changes in market conditions, regulatory developments or otherwise. Statements in this document are made as of the date of publication unless otherwise indicated.

Any statements, estimates, projections, opinions, assumptions, scenarios or forward-looking statements contained in this document are for informational and illustrative purposes only. They are subject to change without notice and may not materialise. They should not be relied upon as predictions, guarantees or assurances of future performance, outcomes or results.

References in this document to particular protocols, networks, digital assets, tokens, financial instruments, products, services, companies or market participants are for explanatory purposes only. Such references do not constitute an endorsement, recommendation or assessment of suitability by any Co-Author. Any financial instruments, digital assets, tokens, products or services described or referred to in this document may not be available, suitable or lawful for all recipients or in all jurisdictions.

Nothing in this document constitutes a representation that any investment, strategy, product or service is suitable or appropriate for any recipient's individual circumstances. Investments and digital assets involve risks, including the potential loss of capital, market volatility, liquidity risk, operational risk, technology risk, regulatory risk and counterparty risk. Recipients should make their own independent assessment and seek appropriate professional advice, including legal, tax, regulatory, accounting and financial advice, before making any decision based on or connected with the information in this document.

Any formulas, models, calculations, prices, yields, returns, examples or scenarios included in this document are provided solely for informational or explanatory purposes. They do not represent valuations for individual investments, executable prices, firm quotes, commitments, guarantees or assurances that any transaction can or could have been entered into on such terms. Different assumptions, methodologies or market conditions may produce materially different results.

In preparing this document, the Co-Authors may have used artificial intelligence-enabled tools to assist with research, summarisation, drafting or editing. Any such use was subject to human review. No AI-enabled tool should be understood to have provided legal, financial, investment, tax or regulatory advice.

To the maximum extent permitted by applicable law, the Co-Authors, the Guest Contributors and their respective affiliates, directors, officers, employees, representatives and agents disclaim all liability for any loss, damage, claim, cost or expense, whether direct, indirect, consequential or otherwise, arising from or in connection with the use of, reliance on, publication of, distribution of, or inability to use this document or any information contained in it.

Unless otherwise indicated, copyright and all other intellectual property rights in this document are owned by or licensed to the Co-Authors or, in the case of a guest contribution, the relevant Guest Contributor. All rights are reserved except as expressly set out in this disclaimer.

© 2026 AMINA Bank AG, Solana Foundation, TensorX, APEX:E3 and Cardano Foundation. Guest contributions © 2026 Blindsight Technologies AG and CV VC AG. All rights reserved.

EXECUTION LAYER · SOLANA FOUNDATION

Why Solana is the execution layer.

Solana was not retrofitted for AI. Its architecture maps directly onto what autonomous agent infrastructure requires: instant, cheap, composable settlement that operates continuously without human intervention.



CHAPTER AUTHOR

Rishin Sharma

Artificial Intelligence · Solana Foundation

AI @ Solana · Founder, Wharton Blockchain Lab.

ABOUT

Rishin leads artificial-intelligence work at the Solana Foundation, focused on where autonomous agents and high-performance on-chain infrastructure meet. He is the founder of the Wharton Blockchain Lab and has worked across applied machine learning and decentralised systems.

Solana was not retrofitted for AI.

The properties that make Solana suitable for autonomous agents were design decisions taken years before anyone framed the problem in these terms. They happen to line up almost exactly with what an agent needs, which is worth explaining rather than asserting.

Begin with state. An agent rarely does one thing. It reads a position, checks a price, moves a balance and records the result, and it needs all of that to be true at the same instant, in the same place. Solana maintains a single global state rather than spreading execution across separate environments, so a multi-step action either happens in full or does not happen at all. For a workflow assembled from several agents acting on the same portfolio, that is the difference between a system an institution can reason about and one it cannot, and no mere convenience.

Then timing. Blocks confirm in roughly four hundred milliseconds today, and the Alpenglow upgrade brings finality to around a hundred and fifty. The number matters less as a headline than as a threshold: it is the point at which settlement stops being something the agent waits on and becomes something that happens inside the loop. Card networks record a transaction in milliseconds too, and then settle it over one to three days. Here, recording is settlement.

There is a second property of that design that institutions tend to appreciate faster than engineers expect. Because the network separates programs from the data they act on, unrelated activity can be processed in parallel rather than queued behind it. In practice this means congestion in one corner of the market does not propagate into another: a surge of activity in some unrelated token does not slow the settlement of a treasury instruction. Anyone who has watched an operational process degrade because an unrelated system was busy will recognise why that matters.

The single-ledger property has an unglamorous benefit too. When every participant reads the same state, reconciliation between counterparties largely stops being a task. There is no version held by one institution to be checked against a version held by another, because there is only one record and both parties are looking at it. A great deal of what back offices do exists to resolve disagreements that this architecture does not create.

And cost. A fee measured in fractions of a cent is the condition that decides which behaviours are possible at all, long before it is a line item to be compared on a table. An agent that pays a tenth of a cent for a second of compute cannot exist on infrastructure where the fee exceeds the purchase. Once fees fall far enough, whole categories of activity, per-second metering, per-call payment, continuous rebalancing, stop being uneconomic and simply become ordinary.

It is simply what agents turn out to need, and none of it was built for them.

Fees are the condition that decides which behaviours can exist at all, [long before they are a feature to be compared.](#)

The dollars on Solana are working, not parked.

If you want to know what a network is actually used for, look at what moves across it rather than at what sits on it.

On Solana the answer is stablecoins, and the pattern is distinctive. The network holds a relatively modest share of the world's stablecoin supply while settling a far larger share of stablecoin transfers. Measured across a year, the value moved runs into the trillions. Put the two figures side by side and the conclusion is that a dollar held on Solana turns over several times faster than a dollar held elsewhere.

That velocity is the number worth attending to, because it is evidence of use rather than of storage. A dollar that moves several times a week is being spent, swept, settled and redeployed. A dollar that sits still is collateral. Agentic finance needs the first kind, and needs it in volume.

The most human demonstration is remittances. Those rails have scarcely changed in fifty years, and the cost shows up in three places: an average fee of about six and a half per cent on a two-hundred-dollar transfer, three to five business days to settle, and hundreds of billions of dollars locked up in pre-funded accounts to guarantee the money arrives on time. Stablecoin transfers settle in seconds, cost a fraction of that, and reach people the banked corridors never served. Established remittance businesses now run stablecoin corridors in production, not in pilot.

The reach matters as much as the cost. Something over a billion adults have no bank account, and close to half of them carry a smartphone. A corridor built on stablecoins serves that population directly, without requiring a branch, a correspondent relationship or a minimum balance, and it does so with a settlement guarantee the traditional corridor cannot offer, because the transfer either completes on the ledger or does not happen.

Institutional settlement followed the same path. Tokenised money-market funds clear on Solana; a global bank has completed a full trade-finance lifecycle on it; commercial paper settlement has moved onto it; and payment stablecoins issued by regulated institutions have launched natively on it rather than being bridged in later.

It is worth noticing what these two groups, families sending money and institutions settling trades, have in common. Both were drawn by the same three properties: the transfer is final in seconds, it costs almost nothing relative to its value, and it does not require the counterparty to be reachable during business hours. Those are exactly the properties an autonomous agent needs, which is why the agentic economy is likely to inherit this infrastructure rather than commission its own.

The point for this report is narrow. An agent needs a unit of account that is stable, programmable and final within its own decision cycle. That is precisely what a stablecoin on a fast ledger is. The rails an agent will use were built for people sending money home and institutions settling trades, and they already work at scale; nobody designed them for agents either.

The rails the agentic economy needs **are the ones already carrying the dollars.**

Agents are already transacting here.

It is easy to describe an agentic economy in the future tense. On Solana much of it is already running, and the useful exercise is to look at what these systems have in common rather than at their individual metrics.

There is a marketplace for real-time inference that settles compute payments on-chain across thousands of active GPUs, so a model can be paid for by the second rather than by the month. There is an agentic operating system that connects agents to on-chain protocols, with a large plugin ecosystem around it. There are distributed networks in which people contribute bandwidth, data or compute and are compensated directly, at amounts far too small for any conventional payment rail to process.

What links them is that each depends on paying very small amounts to very many parties, continuously, without a human approving each one; the blockchain is incidental. That is a payment problem before it is an AI problem, and it is the reason this activity concentrated on a network where a payment costs a fraction of a cent and settles immediately.

The same ecosystem has been useful in a less comfortable way: it has surfaced what does not work yet. Practitioners building agentic banking on these rails report that the hard problems are rarely technical. An agent cannot hold an account in its own name; some person or legal entity must stand behind it, which is the argument §03 of this report makes from the banking side. Onboarding, limits and revocation still have to be designed deliberately. What the rails supply is the ability to move value at the right granularity and speed; the accountability around that movement is a separate construction, and this report devotes two chapters to it.

THE TOOLKIT

The developer surface has matured alongside it. The Solana Agent Kit connects an agent directly to on-chain protocols; the Solana MCP puts that capability inside the development environment where agents are actually built; and agentic payment infrastructure is available from production providers rather than assembled from scratch. An agent registry and skills directory let agents discover what other agents can do.

For an institution, the significance is that the integration work has already been done and tested by a large ecosystem, so a treasury or payments agent is deployed onto proven infrastructure rather than pioneering it.

There is a final point that is easy to miss when reading a list of projects. None of these systems was built as a demonstration of agentic finance. They were built because a specific payment problem, paying a GPU by the second, paying a contributor a fraction of a cent, paying for a model call, had no workable answer anywhere else. Infrastructure that emerges this way tends to be more durable than infrastructure built to prove a thesis, because it is already load-bearing for someone.

The agentic economy is already functioning, and it is already settling value.

x402: a thirty-year-old status code, finally used.

When the web was specified, its authors reserved an HTTP status code for a situation they expected to arise and could not yet solve. Code 402 reads Payment Required. For three decades it sat unused, because there was no way for a machine to pay for a resource without a human first opening an account, entering card details and accepting terms.

x402 puts that status code to work. When an agent requests a resource it cannot access, the server answers with a price rather than a rejection. The agent evaluates the price against its budget, pays, and receives the resource, in the same exchange, with no account, no subscription and no human in the path. Payment settles in stablecoins.

The consequence worth dwelling on is what replaces the account. The agent's wallet becomes its identity: the thing that pays is the thing that is recognised, which removes the API key, the billing relationship and the provisioning step in one move. An agent can therefore transact with a service it has never used before, and the service can charge it without knowing anything about it in advance.

It is worth being clear about where x402 sits in the stack this report describes. It answers the question of how an agent pays: the mechanics of a per-request charge between two parties who have no prior relationship. It does not decide whether the agent should have made the purchase, and it does not establish who the agent is acting for. Those are the mandate question of §2.3 and the identity question of chapter 07 respectively. Payment, authority and identity are three separate layers, and a great deal of confusion in this field comes from treating them as one.

x402 is an open standard, now hosted by the Linux Foundation rather than owned by any single company, and adopted across card networks, payment processors and cloud providers. That matters for an institution: this is becoming shared infrastructure with a governance model, not a proprietary interface that can be withdrawn.

Stablecoins are what make it practical. A per-request payment is only sensible if the payment is smaller than the thing being bought, settles immediately, and is denominated in something whose value will not move between request and response. On Solana, a payment of a fraction of a cent clears in under a second, which is why the standard and the network fit together as well as they do.

The mechanism also has a natural failure mode that institutions should price in. Because payment becomes frictionless, the discipline that used to come from friction has to be supplied elsewhere: by spending ceilings, counterparty allowlists and velocity limits enforced in infrastructure rather than in a prompt. §8.4 sets that out as a risk in its own right. The right conclusion is that removing a human from the payment path obliges you to put the controls somewhere the agent cannot reach; per-request payment itself is no more dangerous than any other kind.

The account was always a human interface. [An agent needs a wallet, a price and an answer.](#)

AI will not wait at a checkout screen.

Agentic payments have a friction problem that is easy to state. Every call an agent makes needs its own approval and its own settlement. That is perfectly workable for a single payment. It falls apart the moment an agent fires hundreds of small paid calls in a loop, because a human ends up back in the middle, approving them one at a time, and every one of them carries a settlement cost and a delay.

Payment channels are the Solana primitive built to remove that. The analogy is a bar tab, or a prepaid meter: money goes down at the start, usage runs up without a payment each time, and the whole thing is settled once at the end. An agent authorises a ceiling, spends against it, and settles the total.

There are four movements, and the interesting property is what happens in the middle two. The agent's wallet deposits a ceiling into on-chain escrow held by the program rather than by an operator. Usage is then metered off-chain: each call is a signed message rather than a transaction, so it costs nothing and waits for nothing. When the work is done, a single transaction records what was actually consumed. Funds go to the recipients, and whatever was not spent returns to the wallet it came from.

Funds stay non-custodial for the whole of that cycle, which is the difference between this and the prepaid-credits arrangement it resembles. In the familiar version, the balance leaves your control and reappears as a number in somebody else's database. Here it stays in escrow governed by the program until it is either spent or returned.

The fit with agent work is specific. Per-call settlement is a poor match for a purchase whose cost is not known in advance, a model billed by the token or compute billed by the second, and a worse one for many small deliveries: token-by-token streaming, or a burst of a few hundred cheap calls. A channel amortises all of it down to one open and one settle, however much metering happens in between.

It also fits what already exists. The mechanism maps onto the standards already live on the network: authorising a ceiling for a single metered call, settling many deliveries together, or opening a session that streams metered deliveries and settles once when it goes idle. The program and the specification are open-source through the Foundation, which matters if you intend to depend on them.

The scale has been measured rather than asserted. The Foundation ran a hundred thousand wallets through a payment-channel proxy and cleared more than a million payments a second, at a cost per payment so small that the figure stops carrying meaning, and published the template and load tests so the result can be reproduced. Inference endpoints are already consumable this way today: an agent authorises once and draws on them at scale, with no account to open and no approval to wait for.

Settlement stops being something the agent waits for. **It becomes something that happens once, after the work is done.**

FURTHER READING

solana.com/solutions/ai

The Foundation's hub for agent tooling — Agent Kit, MCP, x402 and the payment-channel template.

SOVEREIGN AI · TENSORX

Why sovereign AI infrastructure is non-negotiable.

Agentic Finance cannot run on conventional AI infrastructure without failing regulatory reality. It is an active compliance failure waiting to happen, and anything but theoretical.



CHAPTER AUTHOR

Craig Donnelly

Chief Technology Officer · TensorX

Technology leader in sovereign AI and high-performance infrastructure · two decades across enterprise software, fintech and cloud.

ABOUT

Craig Donnelly is Chief Technology Officer at TensorX, where he leads the development of sovereign AI infrastructure for regulated financial institutions and other compliance-critical industries. With more than 20 years of experience building and leading engineering teams, he has delivered mission-critical platforms that process billions of transactions daily. Based in Dublin, he specialises in distributed systems and the engineering of production AI, with particular focus on the demands that come with deploying it at scale: security, operational resilience and the assurance that data stays where it should. His career spans early-stage ventures to multinational organisations, with a consistent thread of making emerging technology dependable enough for institutions to build on.

ABOUT TENSORX

Sovereign AI infrastructure for regulated industries, including finance, healthcare and government. Delivering inference and agentic workloads in jurisdiction-bound, zero-retention environments, with drop-in OpenAI compatibility for existing AI tooling. Directly consumable by agents via x402, settled in stablecoins on Solana.

CHAPTER AUTHOR · CRAIG DONNELLY

CHIEF TECHNOLOGY OFFICER · TENSORX

Where the intelligence actually runs.

Every other chapter of this report describes agents that decide. This one is about where the deciding physically happens.

An agent's reasoning is a sequence of inference calls made against a model running on a particular machine, in a particular building, under a particular legal jurisdiction, and nothing about that is abstract. When a treasury agent weighs whether to sweep a balance, the position, the mandate, the counterparty allowlist and the reasoning chain all pass through that machine. Multiply that by the cadence this report describes, continuous rather than periodic, and an institution is no longer sending the occasional query to an AI provider.

It is worth being precise about what an inference call carries. A model never sees a question in isolation; it sees the context assembled around it: the positions, the recent history, the mandate text, the tool outputs, whatever the agent gathered in order to reason. The context window is the exposure surface, and in an agentic system it is refilled thousands of times a day. Anyone still reasoning about AI risk in terms of what an employee might paste into a chat box is measuring the wrong thing by two or three orders of magnitude.

That changes the nature of the infrastructure question. For a chatbot, the question is quality. For an agent operating inside a regulated mandate, the question is jurisdiction: whose law governs the machine that just read your portfolio?

Conventional AI infrastructure was not designed with that question in view. The large providers optimised for capability and scale, and the commercial terms followed the technology rather than the other way round. For a marketing team drafting copy, the arrangement is unremarkable. For a bank whose supervisor expects to know where client data was processed, who could compel its disclosure and whether anything persisted after the transaction, the same arrangement is a finding waiting to be written.

There is a related distinction that vendor language tends to collapse. 'We do not train on your data' is a statement about one particular use. It says nothing about retention, about logging, about human review of flagged conversations or about what could be produced in response to a lawful order. A provider can hold every one of those positions in good faith and still be holding your data. The question an institution needs answered is whether the provider has the data at all, well before what it intends to do with it.

The gap does not announce itself. Models keep answering, agents keep executing and nothing looks broken. It surfaces at examination, when someone asks a question the architecture cannot answer, and by then the workflows have been running for a year.

An institution running agents is streaming its book through an AI provider, rather than sending it the occasional query.

A European region is not the same as European law.

In June 2025, Microsoft's director of public and legal affairs for France appeared before the French Senate and was asked, under oath, whether he could guarantee that data belonging to French citizens and held in Microsoft's cloud would never be handed to United States authorities without the approval of the French state. He answered that he could not guarantee it.

The answer was accurate rather than evasive. Under the US CLOUD Act, American authorities may compel a US-controlled provider to produce data it holds anywhere in the world, including in European data centres, without notifying or involving European authorities. The obligation attaches to the company, not to the server. A European region operated by an American company is European in geography and American in law.

This is the distinction the market most often blurs, and it is worth stating plainly. **Residency is about where the bytes sit. Sovereignty is about who can reach them.** Most procurement processes test the first and quietly assume the second follows.

There is a second layer to check, and it is the one that most often fails quietly. A European subsidiary of an American parent is, for these purposes, still American: control follows the corporate group rather than the letterhead. So does the sub-processor chain. A provider may hold data in the EU and still route inference, logging, monitoring or support through infrastructure and people who sit outside it. The question a data-processing agreement should answer is who, anywhere in that chain, could be compelled, and by whom; where the primary region sits is the smaller question.

The legal ground beneath transatlantic transfers has not been stable either. The framework that currently legitimises those flows is the third of its kind; the two before it were struck down.

None of this implies bad faith on the part of American providers. Their engineers are not indifferent to privacy and their contracts are not written to deceive. The point is structural: no commercial commitment can override a legal obligation in the jurisdiction that governs the company making it. Which is precisely why diligence should be aimed at the architecture rather than at the assurances.

For most workloads the difference never becomes visible, which is precisely why it persists. For a regulated institution running autonomous agents over client assets, it becomes visible exactly once, under examination, when the honest answer to where was this processed, and who could have compelled it? is the answer nobody wanted to give.

The conclusion TensorX draws from this is deliberately narrow. The remedy is an architecture in which the assurance is unnecessary, where the provider could not produce the data under compulsion because the data was never kept. A stronger assurance would change nothing.

You want a provider who structurally cannot look, rather than one who promises not to.

What European institutions build their AI on.

Something is happening in European boardrooms that is not about agentic finance and is about to collide with it. Institutions that spent two years running pilots are moving AI into the operating core: underwriting, servicing, reconciliation, research, code. The transformation is real, and the competitive logic set out in §01 applies to European firms exactly as it applies to everyone else. The open question is what they build it on.

The default is a closed model from a US provider, reached over an API. It is the fastest route to something that works, and for a great many uses it is entirely reasonable. But it carries a dependency that compounds quietly. The weights cannot be inspected. A version cannot be pinned with any confidence that it will still be there next quarter, because models are updated, deprecated and withdrawn on the vendor's schedule rather than the institution's. The system cannot be fully audited, because the component that did the reasoning is not available for examination. And the supplier is also a potential competitor: the large model companies are moving up the stack into applications and agents of their own, which means an infrastructure vendor can become a competitor without asking permission.

For a marketing workflow, that is a manageable commercial risk. For an institution that intends to let software move capital inside a regulated mandate, it is a structural one.

That has a practical consequence. Much of what the Act asks for is documentary: which system, which version, what data, what result. An institution running an open-weight model on infrastructure it controls answers from its own records; one querying a closed endpoint depends on a third party to answer, on that party's timetable.

OPEN WEIGHTS AS THE BASE LAYER

The open-weight frontier has closed most of the gap that once justified the trade. Models published with open weights, GLM, Qwen, Kimi and DeepSeek among them, now sit close enough to the proprietary frontier that, for most institutional work, capability no longer decides the question. Control does. An open-weight model can be run on hardware the institution or its provider owns, in a jurisdiction it chooses; it can be pinned to a version and kept there; and it can be examined.

That last property is why open weights are a precondition for agentic finance, well beyond any European preference. §2.5 of this report describes agents whose on-chain attestations reference an audited, version-pinned runtime. That guarantee means something only if the runtime can in fact be pinned and in fact be audited, which is to say only if the weights are available to whoever must sign the attestation. An agent whose reasoning rests on a model that may change without notice cannot honestly attest to anything.

Regulation is moving the same way. The EU's Artificial Intelligence Act, Regulation (EU) 2024/1689, has passed from drafting into enforcement, and the enforcement is being handed to the regulators institutions already answer to. Ireland's Regulation of Artificial Intelligence Act, signed into law in July 2026, created a national AI Office and designated fifteen competent authorities, the Data Protection Commission and the Central Bank among them. A bank's AI supervisor, increasingly, is its banking supervisor, and documentation duties that once read as future work now have an address to send the answers to.

For an agent that must attest to the runtime it reasoned on, sovereignty is a precondition, [and a European preference only second.](#)

Structural, not contractual.

What TensorX means by sovereign is narrow, and every part of it is testable.

Inference runs on EU infrastructure: in Dublin on hardware TensorX owns, and in Helsinki on hardware leased from a Finnish provider, rather than capacity rented from a company that answers elsewhere.

Prompts and completions are processed and discarded: nothing stored, nothing logged for training, nothing available later to be produced under compulsion, because it does not exist to be produced. The models served are open-weight, and customer traffic is never trained on. Integration is deliberately unremarkable: the API is OpenAI-compatible, so moving an existing agent across is a change of base URL rather than a rewrite. GDPR compliance, ISO 27001 readiness and a full data-processing agreement sit underneath.

None of that is unusual to claim; most providers claim some version of it. What matters is whether each item is a property of the architecture or a clause in a contract. That difference is testable by anyone willing to ask, and the five questions below are the ones we would put to any vendor, including ourselves.

Two things TensorX does not claim. It does not claim that sovereign infrastructure makes a model correct; §8.4 sets out why model risk survives every architectural control. And it does not claim that every workload needs this. Much of what an institution runs is not sensitive, and paying for sovereignty where it earns nothing is its own kind of waste. The argument here is narrower: where an agent reasons over client positions inside a regulated mandate, the jurisdiction of the machine is part of the control environment, and it should be chosen rather than inherited.

FIVE QUESTIONS TO PUT TO AN AI VENDOR

01 Who owns the data once it leaves our systems, and can you prove it?

02 Where is it cached: in which country, under whose jurisdiction?

03 Are our prompts retained, logged or used to improve your models?

04 Do you build products in our space, or could you tomorrow?

05 Who controls the model weights, and what happens to our workflows if access changes overnight?

The questions are not rhetorical. Each has an answer that is either architectural or aspirational, and the two do not sound alike.

AI stopped answering. It started working.

Ask a strong model something hard today and it plans, calls a tool, reads a document, runs a calculation, checks its own result and only then commits. For anything consequential that is the behaviour an institution wants: a slower answer that is right beats an instant one that has to be checked by hand.

It also changes what the infrastructure has to be good at. A single answer is one inference. An agentic loop is many, in sequence, each waiting on the one before it. Latency does not sit beside the others; it accumulates. A gateway that looks fast on a single call can still be slow across a loop that makes a dozen, and re-sending the same context at every step compounds in the same way. What decides the experience of an agent is how well the loop holds together: how context is carried between steps, how much of the expensive prefix can be reused and how the whole thing behaves under load.

This is where sovereignty and performance stop being a trade-off. A provider running its own inference stack can tune it for the workload actually running on it, and the workload this report describes is many short, sequential, context-heavy calls, running continuously, over data that must not leave the jurisdiction, which is a long way from a chat window.

One further property matters only once someone asks about a decision after the fact. Replaying an agent's reasoning, the same model, the same version, the same weights, is straightforward when the runtime is pinned and under your control, and effectively impossible when the model behind the endpoint has been updated in the meantime. Auditability is being able to run it again, over and above logging what happened.

MACHINE-NATIVE ACCESS

One further requirement only appears once agents are real. An agent that must open an account, accept terms of service and present a corporate card is running a person's errand rather than acting autonomously. TensorX inference is consumable directly by agents through **x402**: an agent needing sovereign inference pays per query, in stablecoins, with no human in the billing path, the same settlement rail described in §2.2 and §04. The compliance layer and the payment layer become the same layer.

```
$ agent.inference({ provider: "tx402.ai" })
→ x402 paywall · pay stablecoins · sovereign EU inference · zero
retention
```

Without compliant inference, agentic finance does not scale in regulated markets. The barrier is knowing the infrastructure already exists; [the technology is in place](#).

INTELLIGENCE · APEX:E3

ALICE — the orchestration layer.

Every execution agent depends on an upstream input: the analytical reasoning that translates an institution's mandates into specific portfolio decisions. ALICE is APEX:E3's award-winning agentic platform, and the layer that produces them.



CHAPTER AUTHOR

Usman Khan

Founder & CEO · APEX:E3

Computer Scientist · Building ALICE, the sovereign AI agent for capital markets · powered by APEX:E3 technology · Founder & CEO · Multiple exits · Loves solving problems using code.

ABOUT

Usman Khan is the Founder and CEO of APEX:E3, the firm building ALICE, the sovereign multi-agent AI platform for institutional capital markets. A computer scientist by training, he spent eight years at UBS Investment Bank in fixed income algorithmic trading before co-founding Algomi, recognised as Risk.net's Trading Technology of the Year and selected for Tech City's Future Fifty. He has been named to the Wall Street Journal Top 40 Under 40 in trading and technology, and to FinTech's 15 People to Watch.

ABOUT APEX:E3

APEX:E3 is an enterprise AI infrastructure company purpose-built for capital markets, enabling institutions to build, deploy, and operate production-grade AI applications and agents. Its BDaaS (Big Data Analytics as a Service) platform ingests, stores, and queries data in real time to power agentic AI across enterprise systems. At its core is ALICE, APEX:E3's award-winning AI agent for capital markets, delivering multi-asset automation and decisioning at scale. Spanning model deployment, inference, coding, and fine-tuning, APEX:E3 brings the full AI model stack into one system built for enterprise. The offering is further strengthened by a years-long research collaboration with the University of Oxford's Department of Computer Science.

CHAPTER AUTHOR · USMAN KHAN

[LINKEDIN.COM/IN/USMAN-K](https://www.linkedin.com/in/usman-k) · APEX:E3

Reasoning is its own layer in the stack.

Programmable settlement, tokenised yield instruments, sovereign inference and agent tooling on the execution layer establish the foundation on which machine-managed capital becomes viable. Alongside them sits an equally essential capability, and it is the one most often left implicit: **the analytical reasoning that translates an institution's mandates into the specific portfolio decisions autonomous agents then execute.**

Every execution agent depends on an upstream input of this kind. A Treasury Agent sweeping idle capital into a tokenised money-market fund requires a defined liquidity ladder and a stated risk tolerance. A Yield Agent rebalancing tokenised fixed income requires a target allocation grounded in prevailing conditions. A Compliance Agent requires constraints articulated against the institution's risk framework. None of those inputs appears by itself. The intelligence that produces them is its own layer, distinct from execution and interlocking with it.

It is worth being precise about what that layer has to do, because the word decision makes it sound like a single act. It is not.

A workflow, not a single decision.

A defensible allocation call requires pulling current exposures, ingesting live market and liquidity data, reviewing fundamental and macro research, running risk and scenario analysis, verifying every option against mandate and concentration limits, stress-testing under adverse conditions, and producing an executable instruction set with reasoning able to withstand committee review.

For a human team that is the work of hours, frequently days. At the cadence Agentic Finance contemplates it has to happen continuously, a tempo no human-led process can sustain, and one that no amount of additional staffing fixes, because the constraint is sequential rather than one of capacity.

WHY A SINGLE MODEL IS NOT ENOUGH

A common assumption is that a sufficiently capable language model can address tasks of this complexity end-to-end. In institutional portfolio optimisation, that assumption does not hold. Market-data retrieval, risk modelling, macro synthesis, and mandate enforcement each draw on distinct competencies; an undifferentiated model attempting all of them at once tends to produce **fluent output without traceable provenance**, the failure mode that accounts for much of the caution regulated institutions show toward production AI in capital allocation.

An orchestrator coordinating specialist agents.

ALICE coordinates an ensemble of specialist agents: market-state and liquidity observation, fundamental and protocol research, on-chain analysis, risk and scenario modelling, macro and regime assessment, and mandate enforcement.

No single agent is sufficient. The orchestrator determines which questions to put to which specialist, in what order, against which data, and how the outputs combine into a coherent recommendation with a documented reasoning chain.

The word doing the work there is documented. An institution cannot act on a recommendation it cannot interrogate, and the difference between a system that produces an answer and one that produces an answer a committee can take apart is the difference between a demonstration and a deployment.

What ALICE provides is best described as a harness: the layer that turns raw model capability into production workflows by coordinating data, models and tools, and by holding the whole sequence together when any one step misbehaves. Capability is rarely the binding constraint in an institutional deployment. Coordination is.

Each specialist is narrow on purpose. A market-state agent is not asked to reason about credit; a mandate-enforcement agent is not asked to form a macro view. Narrowness is what makes an answer attributable: when a recommendation is questioned, it can be traced to the agent that produced that part of it, the data it consulted, and the constraint it was working within.

SOVEREIGN BY DESIGN

Privately deployed and model-agnostic: run on the institution's own infrastructure and models, or paired with sovereign inference such as TensorX. No third party between the institution and its portfolio reasoning.

CAPITAL-MARKETS-NATIVE

Specialisations, evaluation criteria, and reasoning patterns built around how institutional capital is actually managed, market microstructure, multi-asset analytics, regulated mandate enforcement, rather than retrofitted from a generalist template.

FULLY AUDITABLE

Output is an updated portfolio view: the specialist agents that contributed, the data they consulted, what changed since the prior view, and the recommended action: the artefact a committee, risk function, and regulator can interrogate.

Continuous optimisation of a multi-asset crypto portfolio.

Crypto is the asset class in which this architecture currently demonstrates end-to-end operability, because every layer of the stack is simultaneously available. Markets run continuously; on-chain data is queryable in real time via real-time on-chain data services; the assets are programmable; and execution is available autonomously through the agent tooling on the execution layer. Consider a mandate spanning majors, tokenised cash equivalents such as Benji, and an on-chain yield sleeve. A monthly human review cannot respond to conditions that change over **hours**.

ALICE runs this workflow continuously: its specialist agents maintain a live view of market state, surface the developments moving the assets, reprice the portfolio's drawdown and tail-risk exposure, and update the macro and regime view. The orchestrator synthesises these into a current target allocation, alongside an explicit reasoning chain: what changed, which agents drove it, the data consulted, and the mandate constraints applied.

The output is machine-readable and structured for handoff: the **Yield Agent** rotates the on-chain yield sleeve, the **Treasury Agent** manages the tokenised cash leg, Solana provides the settlement rails, and TensorX, where adopted, ensures every reasoning step ran on sovereign EU infrastructure. **The loop closes within a single integrated stack.**

– 6.5 WHERE ALICE SITS IN THE LOOP

01	Signal — ingest market & portfolio data
02	Reason — ALICE orchestration ◀
03	Decide — mandates → instructions
04	Comply — pre-trade validation
05	Execute — atomic settlement in stablecoins
06	Report — immutable audit trail

Each execution agent is a consumer of orchestrated analytical intelligence; without this layer, autonomous execution falls back on pre-configured rules.

Whose model is thinking, and where does it think?

There is a question underneath everything in this chapter that institutions have only recently begun asking out loud. It is whose model produced the reasoning, and where the thinking happened, rather than whether the reasoning is good.

APEX:E3 answered it early. Four years ago the firm made what was then a contrarian bet: that open-weight models would mature to the point where an institution could own its intelligence outright rather than rent it, and keep its proprietary data away from hosted closed-model interfaces in the process. That bet is now the standard the platform is built to: open models that rival closed ones, running inside the institution's own environment; harness engineering that turns raw model capability into governed, reliable workflows; and proprietary data integrated at every layer without leaving the building.

In practice that means deployment inside the institution's own private cloud or on its own premises rather than a call to somebody else's endpoint. Data stays in that environment. Nothing is used for external training, and no model exposure is shared with other customers. Access is granular, and every output is traceable through logs, citations and the decision path behind it.

The observation the firm keeps returning to is worth repeating, because it contradicts most of what is written about enterprise AI: initiatives fail because integration, control and reliability were never solved, and almost never because the model fell short.

The market has since arrived at the same place. Writing in the Financial Times in September 2026, Toby Nangle reported that "some London-based asset managers, as well as not wanting to build dependency on any single proprietary model, are reluctant to post their most confidential data and intellectual property to closed-model providers."¹⁰ The firm's clients include Vanguard.

This is the argument of §05 arriving from the other direction. TensorX holds that where a model runs determines whether an institution can answer for it; APEX:E3 holds that which model, and who controls its weights, determines whether the institution owns its own reasoning. Together they describe a single requirement: a firm deploying agents over its own capital should be able to say without qualification whose intelligence made the decision, and where it ran.

Settlement costs compress, stablecoins converge, sovereign inference becomes a utility, but the system that decides [which analyses to run, in which sequence, under which constraints, and how their results combine into a defensible institutional decision](#) sits closer to a firm's strategic process than to its infrastructure. Institutions that own this layer own how their strategic process is exercised.

Verifiable identity, the trust layer for agentic finance.



CHAPTER AUTHOR

Sandro Knöpfel

Global Lead, Market
Structure & Strategic
Partnerships · Cardano
Foundation

Two decades in capital
markets · driving strategic
partnerships and the Web2-
to-Web3 transition at the
Cardano Foundation.

Agentic AI can make autonomous decisions, but it cannot take responsibility for them, and institutions cannot deploy what they cannot hold accountable. The missing layer is verifiable identity. The **vLEI** (verifiable Legal Entity Identifier), GLEIF's cryptographic extension of the global LEI, lets an organization and the agents acting on its behalf prove who they are and under whose authority they act, in a format that both machines and regulators can verify automatically. In May 2026, GLEIF and the Bank for International Settlements demonstrated through Project Aperta that this architecture functions across jurisdictions, connecting five central bank networks through LEI-based verification at the core.

Bringing this standard to public blockchain infrastructure requires purpose-built tooling. **Veridian**, a product-focused company originating from the Cardano Foundation, is building that tooling within the KERI/ACDC ecosystem and is pursuing Qualified vLEI Issuer accreditation. When in place, an autonomous agent transacting on-chain can carry a verifiable, revocable credential tying every action back to any legal entity and the human-defined mandate behind it. Every transaction becomes attributable. Every delegation chain becomes auditable.

ABOUT THE AUTHOR

Sandro Knöpfel serves as Global Lead, Market Structure & Strategic Partnerships at the Cardano Foundation. He leverages over two decades of capital-markets experience driving global partnerships and digital transformation, specialising in translating technological innovation into business outcomes, with a focus on positioning blockchain infrastructure within institutional and regulated environments.

The identity gap has not been solved.

This report describes a six-step loop for autonomous capital management, and at step four a compliance agent validates every proposed transaction before execution proceeds. That step carries an assumption which is easy to miss: that the agent can be identified, its authority verified, and the institution behind it held accountable. When a person executes a transaction, that chain is legible. When an agent does the same at machine speed across institutional boundaries, none of those links exists by default.

The reason is historical rather than technical. Financial institutions built their compliance systems around human actors. KYC frameworks, AML monitoring and access controls all assume that whatever sits behind a transaction is a person or a legal entity with a stable, verifiable identity, something that can be checked once, recorded, and relied on afterwards.

Agents break that assumption in three distinct ways. They act at machine speed with no shared trust anchor across institutions, so a counterparty has nothing to check against. They delegate, cascading instructions across systems the original operator never directly touched, so the record of who asked for what becomes long and undocumented. And the authority behind an action may have changed between invocation and execution, a mandate withdrawn a moment ago, on a transaction settling now.

Any one of those is survivable. Together they mean the compliance infrastructure that governs human actors does not transfer to machine actors, and cannot be made to transfer by adding more checks at the application layer. What is missing is the thing every control depends on: a verifiable answer to who is acting, and on whose authority.

The cost of not having that answer is already visible in what institutions spend on the human version of it. Global surveys put average annual KYC spend at around sixty million dollars per institution.²⁴ Financial-sector fraud losses run to roughly thirty billion a year, and where self-sovereign identity has been deployed, onboarding time has fallen by about seventy per cent.³² The Legal Entity Identifier is already referenced in more than three hundred regulations worldwide²³, the clearest signal that the market agrees identity should be the anchor. What it has lacked is a machine-verifiable form of it.

That is the subject of this contribution. Not whether an agent can be trusted to reason, which the rest of the report addresses, but whether at the moment an agent acts anyone can prove which regulated entity stands behind it. Everything that follows here comes back to that question.

\$60M ANNUAL KYC COST PER INSTITUTION · THOMSON REUTERS ²⁴	\$30B ANNUAL FINANCIAL-SECTOR FRAUD LOSSES ³²	70% FASTER ONBOARDING WITH SSI · SIGNICAT ³²	300+ REGULATIONS REFERENCING THE LEI · GLEIF ²³
---	--	---	--

One credential for identity, authority and audit, checked before execution.

The Legal Entity Identifier already anchors institutional identity in over three hundred regulations worldwide.²³ Its cryptographic extension, the vLEI, makes that identity machine-verifiable: an organisation, and the agents acting on its behalf, can prove who they are, what role they hold and under whose authority they act, without moving sensitive data across institutional boundaries in order to do it.

A single credential carries three things institutions normally keep in three systems. **Identity:** it proves the legal entity behind every agent action, cryptographically bound and automatically verifiable by machines and regulators alike. **Authority:** it carries the mandate scope and the full delegation chain, so permission is traceable from invocation through to execution. **Audit:** every transaction becomes attributable to a verified legal entity and a human-defined mandate, and every delegation chain becomes auditable afterwards.

That maps onto the governance tiers this report describes. At Tier 01, where agents execute autonomously inside a mandate, the credential carries that scope cryptographically rather than as configuration in an application. At Tier 03, where mandate design is exclusively human, it anchors the authority structure back to the legal entity whose board approved it. The tiers say what an agent may do; the credential makes that checkable from outside.

The architecture has already been demonstrated across borders. In May 2026 GLEIF and the Bank for International Settlements connected five central-bank networks through Project Aperta, with LEI-based verification at its core.

Regulators have arrived at the same requirement independently, and in four different legal traditions. Singapore's IMDA launched the world's first governance framework specifically for agentic AI in January 2026, naming agent identity and permission control among its four structural requirements.¹⁶ NIST followed in February with an AI Agent Standards Initiative whose explicit focus is agent identity and cross-institutional interoperability.¹⁷ In March the United Kingdom's Competition and Markets Authority published binding guidance confirming that consumer-protection law applies to autonomous agent actions, and that the deploying institution bears accountability.¹⁸ In the European Union, general enforcement and transparency obligations under the AI Act take effect on 2 August 2026, with the Annex III high-risk obligations covering financial operations applying from December 2027 under the Digital Omnibus agreement; there too, accountability sits with the deploying institution.¹⁹

These are parallel conclusions from independent jurisdictions arriving at the same structural gap, and nobody coordinated them: the institution that deploys an agent answers for it, and answering for it requires the agent to be identifiable. Industry has moved the same way. On 24 June 2026 Cardano joined Google, IBM, Circle and other founding contributors in launching the Legal Context Protocol²², an open standard developed with the American Arbitration Association that makes legal terms, consent and dispute resolution verifiable when AI agents transact autonomously.

Identity infrastructure is the foundation on which accountable autonomous operations are built, and no compliance checkbox.

Bringing KERI and ACDC to the execution layer.

A standard is worth only as much as its implementation, and the implementation is further along than most readers will expect. Veridian, a product-focused company originating from the Cardano Foundation, holds maintainer status on KERIA, SignifyTS and KERIpy, the core open-source components of the KERI and ACDC ecosystem, recognised by GLEIF for those contributions.

The KERI protocol lets any entity, an autonomous agent included, hold a cryptographic identifier independently of any central registry. Paired with ACDC credentials, a trusted authority can issue signed claims about an agent’s authorisation scope without the underlying data being exposed at the point of verification. That property matters more for agents than for people: a counterparty may need to check authority thousands of times a day, and every check that requires data to move is a check that will eventually be refused on privacy grounds.

Key management extends down to the hardware. Veridian’s fingerprint smart-card prototype keeps the cryptographic root of trust on the physical device, so it never leaves it.

The accreditation work runs in parallel. Veridian Attributions AG, the Swiss commercial entity, is pursuing GLEIF accreditation to become Switzerland’s first Qualified vLEI Issuer, and the Veridian QVI Suite provides the accreditation stack that other providers can license rather than build. The practical effect shows up in issuance time. Legacy vLEI issuance meant a scheduled video call with the issuer, manual phonetic string verification and one signer processed at a time: days of work. GLEIF-approved remote verification with batch issuance for multiple signers produces the same credential in under ten minutes.

Nor is any of it confined to a pilot. The UNDP Tadamon Accelerator has deployed Veridian’s architecture across fifty-seven member countries, issuing DID-compliant, portable verifiable credentials to civil-society organisations at scale. The architecture underneath is the same one an asset manager needs, even where the credential differs.

KERI / ACDC CORE	ISSUANCE & WALLET	VERIFICATION	ANY VERTICAL
Open cryptographic key-event log. No central registry.	Mobile-native. Keys never leave the device.	Machine-verifiable. No data exposure at point of check.	vLEI, agentic finance, government. One stack.

The governance model this report describes only holds if the identity layer beneath it is equally robust. **Verifiable identity is what makes institutional deployment defensible, and nothing about it is an add-on.**

The attack that never breaks a rule.

The controls in this report bind what an agent may do, and they are the right controls. One check belongs beside them: is what the agent proposes what its principal actually asked for? That question gets decided inside the context window, where an agent that never leaves its mandate can still be steered by someone else.



GUEST AUTHOR

Guilherme Santos

Co-founder & CEO,
Blindsight Technologies
AG, Zurich

Offensive-security background and top-ranked ethical hacker, former Security Architect at Kühne + Nagel, specialises in the security of AI systems.

This report is careful about where its controls stop. §3.2 draws the line deliberately: the bank limits what a compromised agent can do and leaves what happens inside the agent to another layer. §5.1 points out that the context window is the exposure surface, refilled thousands of times a day. §8.4 names the result: a signature proves execution, not intent.

Those three passages describe the job Blindsight does. We sit between what the agent reads and what the agent does, and we answer one question for every action it proposes: is this what the principal asked for, and does it fit the controls they set? The rest of this contribution is about how that question gets answered, and how far the answer can be trusted.

ABOUT THE AUTHOR

Guilherme Santos is Co-founder and CEO of Blindsight Technologies AG in Zurich, where he leads the company's work on securing the AI layer of regulated organisations. He comes from offensive security, ranked among the world's top ethical hackers and credited with more than twenty disclosed zero-day vulnerabilities across critical systems and Fortune 500 companies, with a specialisation in the security of AI and agentic systems. Before founding Blindsight he was a Security Architect at Kühne + Nagel, advising senior management on cybersecurity strategy and the secure adoption of AI and cloud. He is President of the German chapter of the Global Council for Responsible AI and speaks regularly on AI security, including at it-sa, ShmooCon and GITEX.

ABOUT BLINDSIGHT · SECURING AI

Blindsight provides trust to AI systems, securing their runtime and data and providing visibility, all in one consolidated platform. Blindsight secures the AI layer of regulated organisations end to end. On the employee side, it discovers shadow AI use and prevents data leaks. On the AI-system side, it protects agents at runtime against prompt injection, data and RAG poisoning, and indirect injection through tool outputs. Built in Zurich by a team from offensive security and adversarial machine learning. NVIDIA Inception member. blindsight.io · blindsight.io/blog

Inside the mandate, against the user.

Take the treasury from §3.5. The treasurer has set the envelope: FOBXX and a few other instruments, an allowlist of venues, a cap, a person above a certain ticket. Inside that envelope the agent runs all night.

Now put one bad document in front of it. A supplier invoice with an extra line in a field nobody reads, telling the agent to route the payment to a different wallet on the allowlist. Or a filing in the research index that nudges the yield leg toward one instrument the mandate already permits. Nothing here breaks the envelope. The wallet is allowed. The instrument is allowed. The amount is under the cap. The transfer hook passes it, the compliance agent passes it, the vLEI says who signed. The treasurer wakes to a treasury that worked all night and did exactly what somebody else told it to.

I have spent most of my career on the attacking side, and this is what this kind of attack looks like now. Nobody talks the agent out of its rules. They hand it a permitted action and let the controls approve it. §8.4 already says this about signed payment mandates: the signature secures the execution, not the reasoning that built it. Reviewing afterwards does not catch it either. The distance between the bad line in the invoice and the settled payment is a single inference call.

Matching the action to the request.

Every action an agent takes arrives as a tool call: pay this, move that, rebalance here. Before it executes, Blindsight holds it and asks whether it matches what the principal actually asked for. Three things go into that judgement.

What the principal asked for. Not the mandate as a rule set, but the request as the treasurer wrote it: sweep surplus into FOBXX, fund cleared invoices, keep the local entities in currency. That text is the reference point for every action that follows.

What the agent read on the way here, and who wrote each piece. Most of what an agent reads was never written by its principal. A retrieved document, a tool's reply, a note it left itself last week. We mark each of those as outside content before the model sees it, and we carry that marking through to the action. So when the payment instruction arrives, we know it was shaped by a line in a supplier invoice and not by the treasurer.

What the tool said it would do. A payment tool declares what it pays and where. If the call reaches outside that declaration, it fails on that alone.

From those three, one of three outcomes. The action matches the request and the controls, and it goes through at machine speed. It does not, and it is blocked with a record of why. Or it is ambiguous, and it goes up to the human tier that §2.3 already defines, with the evidence attached: what the agent was asked, what it read, what it tried to do.

That last part is what turns the audit trail from a log into an explanation. The record shows not only what the agent did, but what it was told and by whom.

Two things make this work in practice. It has to be fast. If the check adds latency the desk will switch it off, so it runs inline, in milliseconds, without a second model call. And it has to be good at telling voices apart. In our own testing, a detector that knows which words the principal did not write catches materially more than one reading the prompt as a single block: 18.6 points higher detection on InjecAgent, the public benchmark for injection through tool outputs, with no rise in false alarms. The same discipline applied to documents before they reach the knowledge base flagged 98 per cent of the planted files in an index we poisoned on purpose. On that same poisoned index, an unprotected agent repeated the attacker's planted answer in 60 per cent of trials; with Blindsight's check in front of it, that fell to 8.³⁴

In front of the hard limits.

We built this, then attacked it, and we keep attacking it as the technology moves. What that testing has shown is simple. Any check that lives inside the agent, a rule in the system prompt or the model judging its own output, is the agent marking its own homework. It raises an attacker's cost and it stops some attacks. It is not a wall, and nobody should sell it as one.

That is why Blindsight's check runs as its own component outside the model, with its own record of what the agent read and who wrote it, and why it sits in front of the controls this report describes rather than instead of them. Hard limits belong where the agent cannot reach them: custody, the transfer hook, the revocable mandate, the person above the tier. Our layer catches what it can before those limits are tested, and when something does reach them, it gives the institution the record to say what happened and why. You need both.

One more thing, outside the perimeter this report draws. The controls in these pages govern the agents an institution registers, but those are not the only agents inside it. Staff are already running copilots, automation flows and browser agents over the same data, with no mandate and no legal entity behind them. A verified perimeter governs what it knows about. Finding what is actually running is where the governance in these pages has to start.

Protecting your agents from attacks means checking every action [against what was actually asked for](#).

Owning the frontier before it is priced.

The AI-native transition is being built now, mostly in private, by small teams, under macro and geopolitical conditions that no one can forecast. Early-stage venture is the one instrument that holds exposure to the layer where founders adapt fastest, at a price that still reflects what a company might become.



GUEST AUTHORS

David Long
Lukas Etter
Kaya Tilev

General Partners · CV VC

Early-stage venture
investors in founders
building with frontier
technologies across Europe,
the USA, Africa and the
Middle East.

ABOUT CV VC

CV VC AG, Zug, invests in early-stage founders building category-defining companies with frontier technologies across Europe, the USA, Africa and the Middle East, with more than 80 investments since 2019. Its conviction is that the future is AI-native, and it invests in the foundation, building blocks and enablers of that future. CV VC is the co-publishing partner of this report. cvvc.com

THE WINDOW IN WHICH VALUE IS CREATED

Every technology cycle has a moment when the winners are being formed and cannot yet be bought at a public price. In the early internet that moment lasted a few years, and most of the value that later appeared in listed markets was created in a small number of private companies, backed by investors willing to underwrite an unproven architecture and a founder with a point of view.

We think the AI-native transition is at that point, and that it is larger than what came before. Software has always captured the value of the tool. AI captures part of the work the tool was built to support: the judgement, coordination and cognitive labour that used to sit with people. When the marginal cost of reasoning falls towards zero, a company is made of different things. Our job is to be in the room while that is built.

ABOUT THE AUTHORS

David Long, Lukas Etter and Kaya Tilev are the General Partners of CV VC. Between them they draw on experience as founders and operators across banking, private and public markets investing, deep tech and AI, and scale-ups.

A noisy macro, a structural signal.

Anyone writing about markets in 2026 has to acknowledge the noise. Rate expectations move and move again. Trade policy has become an instrument of statecraft. Access to compute, models and capital now depends on jurisdiction. This summer, access to some frontier models was suspended and then restored under export-control decisions, which reminded many boards that a dependency they treated as a utility is a policy variable.

The reflex is to wait. We think most of these headlines are second-order. The first-order fact is that AI has moved from demonstration to production, and value has shifted from raw model capability towards whoever owns the integrated workflow and the deployment: the company that gets an agent into a regulated process, connects it to proprietary data and makes it reliable enough that someone signs off.

Two further shifts do not depend on a favourable macro environment. Sovereignty has become a buying criterion, so geopolitics now generates demand for independent, auditable, compliant systems, particularly in Europe. And digital finance has matured into infrastructure: with regulatory clarity arriving in major jurisdictions, programmable money that software can initiate, verify and settle is a design assumption for founders. Both need builders, and builders need early capital.

Why the early stage.

Listed markets offer liquidity, but mostly through incumbents whose AI opportunity is diluted by everything else they do and largely priced in. Late-stage private markets have absorbed enormous capital, concentrated in a few very large financings. Entry valuations reflect this, and the return profile drifts towards broad equity market movements with a liquidity discount.

The early stage is where architectural choices are still open and the entry price still reflects potential more than proof. **Convexity:** venture returns are driven by a few outsized outcomes, and a portfolio built on small first cheques can absorb many failures and still be defined by the few companies that reprice a category. **Access:** the people building the next layer of vertically integrated intelligent software are, today, mostly pre-seed or seed, and once the demand for a company's solution is obvious the access is gone. **Optionality:** an early position is a right to learn, and we can add capital to those who prove out and let the rest run their course.

Timing matters as much as the stage. In our experience the next wave of companies reach early proof points faster and with less capital than the previous generation of software businesses, because the tools that build software are now themselves AI. That shortens the distance between a first cheque and the evidence needed to underwrite the next one, which suits an investor who commits in stages.

A disciplined shape for an undisciplined asset class.

Early-stage investing has a habit of confusing enthusiasm with strategy, and the current excitement makes that easy. We hold four constraints.

Focus on where value accrues. We think in three converging layers: automation that produces, through AI, robotics and autonomous systems; intelligent business software that decides, orchestrating value chains and cognitive work; and digital finance that transacts, a programmable and verifiable fabric on which autonomous systems pay, settle and prove what they did. Investing across all three is exposure to a system that is converging, and not a bet on one category.

Small first cheques, concentrated conviction. Initial tickets stay modest, so being wrong costs little and we can back a wide range of founders. Larger follow-on capital goes to companies that show real traction, so the evidence a company produces, not our narrative, decides how much more we commit.

Founders before narratives. At this stage there is little revenue history to analyse, so the founder's judgement is the asset: unusual proximity to the problem, the speed to iterate against a moving frontier, and the honesty to tell us when a thesis is failing.

Defensibility over novelty. Model capability commoditises quickly. We look for architectures that will still be defensible in three years: proprietary data, embedded workflow, distribution, regulatory fit, and a model-agnostic stack that survives a supplier changing terms or a jurisdiction changing rules.

Why geopolitics favours the small and the early.

Geopolitical risk is usually presented as a reason for caution. For young companies it is more often a reason for redesign, which is what they do best. An incumbent must unwind supply chains, contracts and legacy systems when the rules change. A seed-stage company carries none of that, and can be built from day one for data residency, open weights and multi-jurisdiction compliance.

For European founders in particular, regional demand for sovereign and compliant systems is a home market that did not exist at this scale a few years ago, and it rewards those who build for it early. It also argues for geographic breadth: a portfolio across Europe, the US and selected emerging markets gains from the differences between regimes and is less exposed to any single one changing course.

What we would say plainly about the risk.

A serious argument for venture must include the counter-argument. Dispersion between the best and the median fund is very wide. Most companies fail, capital is illiquid and the timeline is long, which is why this belongs in the part of an allocation built for patience. Valuations in a hot theme can run ahead of fundamentals at any stage. Liquidity is also changing shape, as follow-on rounds, secondaries and earlier exits give some investors options that did not exist a decade ago, but these supplement the work of picking well and being patient. They do not replace it.

The case, in short.

The AI-native transition is being built now, in private, by small teams, shaped by forces that markets cannot forecast. The useful response is to hold exposure to the layer where founders adapt fastest, at a price that still reflects what the company might become. We think that window is open, and will not stay open.

08

CHAPTER EIGHT

Conclusions & recommendations.

This report has documented Agentic Finance already running, in production, at small scale, on infrastructure the people in these pages built, rather than arguing that it is coming. What remains is organisational, and it is the part an institution controls.

The stack exists. The decision is where to draw the line.

Seven chapters and two guest contributions have made one argument from different directions. Software that decides is being given the means to act. The constraint every financial process was built around, a person in the loop, is no longer a physical necessity, and the cadences it produced are no longer required. What remains necessary is everything that made the person accountable: a mandate, an identity, a record and a boundary enforced somewhere the software cannot reach.

Each part of that boundary now has an owner. AMINA Bank supplies the regulated foundation: custody, fiat rails and a balance sheet that stands behind settled tokens. Solana supplies execution: a single global state, settlement inside the decision cycle, fees that make the smallest useful payment economical and a stablecoin economy already turning over at scale.

TensorX supplies inference that runs inside the jurisdiction, retains nothing and can be pinned and replayed, so that an agent's reasoning is something an institution can answer for. APEX:E3 supplies the layer above execution, the orchestration that turns many narrow agents into one decision a committee can take apart, and the case that the weights doing the thinking should belong to the institution.

The Cardano Foundation supplies identity: the vLEI, which lets an agent prove which legal entity it acts for and under what authority, and which four jurisdictions have now independently decided to require.

Nothing else in the stack establishes it, and without it the compliance step in the loop validates the form of a transaction rather than the authority standing behind it.

Read together, these are five answers to the question this report kept returning to, and only incidentally five products: if the machine is going to act, what has to be true first? The asset has to be held by a regulated institution. The rail has to settle finally and cheaply. The reasoning has to happen where the institution can see and trust it. The decision has to be attributable to the agent that made it. And the agent has to be identifiable as acting for someone.

The status of each should be stated as clearly. The foundation, the rails, the inference and the orchestration are in production today, with clients and measured results. The identity layer is further along in standards than in deployment: the vLEI exists and is used, the accreditation that brings it to agents on public chains is in progress and the token-level design in §2.5 is a next step rather than a current one. A stack described as complete when it is nearly complete is exactly the claim a risk committee should distrust.

What runs through all of it is the stablecoin: the unit agents pay in, the form idle cash takes while it waits, the thing x402 settles and payment channels meter. A tokenised money-market fund is where that cash earns while it waits; Franklin Templeton's is a well-known public example.²⁸ Every other layer exists to make the movement safe; these two are what actually moves.

What this means depends on where you sit. A treasurer has the cleanest case, and the recommendations that follow are written for it first. An asset manager under UCITS or AIFMD should settle §05's jurisdictional question first: where the model that reads the book runs, and who could compel it. A wealth manager has uninvested client cash an agent can keep working, and a builder will find the standards already open and live.

Autonomy inside a mandate is a control environment, and every part of it now has an owner.

Six actions across three horizons. The first needs no technology at all.

01

Write the first mandate and set the tiers.

0-3 MONTHS

Before any technology is bought, decide what an agent would be allowed to do: what counts as surplus cash, which instruments are eligible, how much may move in one step and which decisions need a person. Put the three tiers of §2.3 on a page and take it to the risk committee. This is the step that can start this quarter, and it needs nobody to trust the technology yet.

02

Audit idle cash and payment friction.

0-3 MONTHS

Quantify two things at once: institutional cash earning below money-market yields, and the correspondent-banking and FX cost embedded in cross-border treasury flows. That figure is the pilot's benchmark and the board's business case. APEX:E3 can quantify it independently where internal capacity is short.

03

Pilot a treasury agent on the first cash pool.

0-6 MONTHS

One agent, a defined pool of \$50–200 million, a tokenised money-market fund as the instrument, Solana as the rail and sovereign inference from the first day. Measure for ninety days against the audit, on real cycle times and basis points rather than surveys. It is the lowest-friction, highest-certainty entry point in this report.

04

Make stablecoins the settlement currency, and register the agents.

0-6 MONTHS

For any agent that transacts, stablecoins should be the designated settlement currency from the first deployment, with yield-bearing balances so idle agent cash is never unproductive. In the same step, bind each agent to a legal entity, one agent, one entity, one attestation, so that joining a verified perimeter later is a configuration change rather than a programme.

05

Expand to orchestration, and own the model.

6-18 MONTHS

Add yield, compliance and collateral agents under an orchestration layer that produces a reasoning chain a committee can interrogate. Pin the model version, keep the weights where the institution or its provider controls them and confirm that each agent's attestation names that runtime. This is where the return inflects, and where the governance from step one has to already be in place.

06

Treat it as competitive infrastructure.

18-48 MONTHS

For the institutions that lead, Agentic Finance stops being an efficiency programme and becomes how the balance sheet is run: a different clock speed, different unit economics and, in time, a different return profile. Put it under the people who run the balance sheet, not the people who run the technology.

Six risks, and what the stack does about each.

A report that only argues the upside is not one an allocator can act on. Six risks a risk committee will raise, and where this report answers each.

01 · REGULATION

The rules are still being written.

Singapore, the US, the UK and the EU all issued agent-specific guidance in 2026, and none of it is settled.¹⁶⁻¹⁹ Liability and the treatment of agent-initiated transactions will keep moving.

MITIGATION · Pre-trade compliance (§2.3) and a regulated foundation (§03) absorb rule changes without a redesign.

02 · AUTHORISATION

An agent with keys can spend.

Limits by size, counterparty, asset and time window must live in infrastructure. A limit in a prompt is a suggestion, not a control.

MITIGATION · Mandates and tiers (§2.3); transfer hooks at the token (§2.5); the bank's control layer (§3.2).

03 · PROMPT INJECTION

A signature proves execution, not intent.

A 2026 red-team of Google's Agent Payments Protocol showed signed mandates secure a payment's execution, not the reasoning that built it. An agent can be manipulated upstream of the signature.⁰⁶

MITIGATION · A verified-agent perimeter and version-pinned runtimes (§2.5); human confirmation above Tier 01.

04 · CORRELATED BEHAVIOUR

Many agents, similar models, same signal.

Agents on similar models and data can herd at machine speed. The 2010 flash crash came from far simpler systems that could only sell what they were given.

MITIGATION · Circuit breakers in mandates; Tier 02 confirmation for allocation shifts; model and data diversity across agents.

05 · KEYS, CUSTODY, ISSUERS

Keys can be lost. Issuers can fail.

Keys are a single point of failure, and a stablecoin redeems at par only as reliably as the issuer and reserves behind it.

MITIGATION · Regulated custody with institutional key management (§3.3); redemption backed by a bank balance sheet; diversified issuers.

06 · CONCENTRATION & MODEL RISK

One chain. One model.

One execution layer exposes every agent to that chain's outages. Inference errors and drift do not vanish because inference is sovereign.

MITIGATION · Fallback routing in mandates; version-pinned runtimes (§2.5); continuous audit (§2.3, step 06).

THE EVIDENCE BASE

Measured on-chain agent settlement is still small: ~\$73 million across ~176 million transactions in the twelve months to April 2026, most between one and ten cents (Keyrock, via Hashed Emergent, June 2026)⁰⁷. The transaction count supports this report's argument; the dollar volume does not yet. The opportunity figures in §01 and §06 are models, not measurements.

8.5 CLOSING ARGUMENT

The foreword to this report sets the only test that matters: build it so that you can answer for it. Everything in between has been an attempt to make that possible: the bank behind the asset, the ledger behind the settlement, the jurisdiction behind the inference, the reasoning chain behind the decision, the credential behind the agent. **What no supplier can deliver is the mandate itself: written inside the institution, slower to agree than any system is to deploy and the one part of this that cannot be bought in a quarter.**

About the co-authors.

TensorX TensorX

Sovereign AI · Lead author & architect · Dublin · Helsinki
tensorx.ai · tx402.ai

Sovereign AI infrastructure for regulated industries. TensorX runs open-weight models on EU infrastructure in Dublin and Helsinki, under EU jurisdiction, with zero data retention by design: prompts and completions are processed and discarded rather than stored. The API is OpenAI-compatible, so an existing agent moves across in a single line of code, and inference is consumable directly by agents through x402, paid per query in stablecoins.

IN THIS REPORT

§05 Sovereign AI — the compliance gap, residency versus sovereignty, open weights as Europe's base, and inference built for agentic loops.

CHAPTER AUTHOR

Craig Donnelly
Chief Technology Officer



AMINA Bank

Regulated Banking · Zug, CH
aminagroup.com

FINMA-regulated Swiss bank for digital assets, with regulated custody, settlement and fiat rails that anchor autonomous execution.



Solana Foundation

Execution Layer · Geneva, CH
solana.org

The high-performance public ledger for agent-scale settlement: sub-400ms finality, single global state, and payment channels.



APEX:E3

Quantitative Intelligence · Capital Markets
apexe3.com

Enterprise AI infrastructure for capital markets, operator of **ALICE**, its award-winning multi-agent platform.



Cardano Foundation

Identity Layer · Zug, CH
cardanofoundation.org

Verifiable organisational identity for autonomous agents: the vLEI on public blockchains, through Veridian and the KERI/ACDC ecosystem.

GUEST CONTRIBUTIONS

Blindsight — runtime protection for AI agents against prompt injection and data poisoning.

CV VC — early-stage venture investing in the AI-native transition.

Sources & notes.

Works quoted or relied on, with the chapter in which they appear. Figures originating with a co-author or contributor rather than a published source are listed as such.

CITED WORKS

- 01 **Allaire, J. (2026).**
The Agentic Economy: The Convergence of Intelligence and the Economy — A Treatise. Co-Founder & CEO, Circle; written in a personal capacity. Quoted at pp. 64, 89 (§2.3, §2.5).
AgenticEconomyTreatise.com
- 02 **Pal, R. (26 Aug 2026).**
“DeFi Wasn’t Meant For You.” Substack essay. Quoted in §2.2. Author is a Sui Foundation board member and co-founder of EXPAAM.
raoulpal.substack.com/p/defi-wasnt-meant-for-you
- 03 **ARK Investment Management (Jan 2026).**
Big Ideas 2026. Winton, B., “The Great Acceleration,” p. 6 (Convergence Network Strength +35% in 2025; smart contracts and stablecoins as a monetary ecosystem for AI agents); Grous, N. & Prasanna, V., “The AI Consumer Operating System,” p. 29 (AI-facilitated online spend, 2% in 2025 to ~25% by 2030, >\$8T). Quoted in §02.
ark-invest.com/big-ideas-2026
- 04 **American Banker ON-CHAIN (2026).**
Glossary entry, “Agentic finance.” Quoted in §02.
on-chain.americanbanker.com/glossary/agentic-finance
- 05 **Gong, H. (2026).**
“Agent-to-Agent Finance: Blockchain Payments and Trust Infrastructure for Autonomous AI Agents.” arXiv:2607.00245 (v1 30 Jun 2026; rev. 4 Sep 2026). Cited in §02 for bounded autonomy.
arxiv.org/abs/2607.00245
- 06 **Debi, T., Zhu, W. & Sen Gupta, P. (2026).**
“Whispers of Wealth: Red-Teaming Google’s Agent Payments Protocol via Prompt Injection.” arXiv:2601.22569 (v1 30 Jan 2026; rev. 18 May 2026). Cited in §8.4: signed mandates secure execution, not the reasoning that constructs them.
arxiv.org/abs/2601.22569
- 07 **Hashed Emergent (25 Jun 2026).**
Deshmukh, R., “The Agentic Stack: AI Agents × Crypto Rails in Emerging Markets,” reporting Keyrock’s measurement of on-chain agent settlement (~\$73M across ~176M transactions, May 2025–Apr 2026). Cited in §1.1 and §8.4.
hashedem.substack.com/p/the-agentic-stack-ai-agents-x-crypto
- 08 **Bain & Company (17 Dec 2025).**
“2030 Forecast: How Agentic AI Will Reshape US Retail” — AI agents to account for 15–25% of US e-commerce by 2030. Cited in §1.1, as reported in Kaul (2026), entry 28.
- 09 **TensorX (2026).**
Chapter §05 draws on the company’s published positions: “Your EU Region Is Not Sovereign”; “Europe Is Sleepwalking Into US AI Dependence”; “AI Stopped Answering. It Started Working.”; “Your AI Vendor Can See Your Business. Ask Them These Five Questions.” The five questions in §5.4 are reproduced from the last of these.
tensorx.ai/blog
- 10 **Financial Times (11 Sep 2026).**
Nangle, T., “How big is the open-model threat to AI hyperscalers?” FT Alphaville. Quoted in §6.6 on institutional reluctance to send confidential data to closed-model providers; reports that APEX:E3’s clients include Vanguard.
ft.com
- 11 **McKinsey & Company (Nov 2025).**
The State of AI: How Organizations Are Rewiring to Capture Value. Global survey; of the organisational attributes tested, fundamental workflow redesign has the largest effect on EBIT impact from gen AI, and high performers are roughly three times as likely to have redesigned workflows. Self-reported survey data. Cited in §2.4.
mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai
- 12 **Dell’Acqua, F., McFowland, E., Mollick, E. et al. (2023; 2025).**
“Navigating the Jagged Technological Frontier.” Harvard Business School Working Paper 24-013; published in Organization Science (2025). Pre-registered field experiment with 758 BCG consultants: inside the frontier, quality 40%+ higher; outside it, 19 percentage points less likely to be correct. Cited in §2.4.
hbs.edu · pubsonline.informs.org
- 13 **Brynjolfsson, E., Li, D. & Raymond, L. (2025).**
“Generative AI at Work.” Quarterly Journal of Economics 140(2). Staggered rollout across 5,172 customer-support agents: ~15% average productivity gain, ~34% for the least experienced, near zero for the most experienced. Cited in §2.4.
academic.oup.com/qje · nber.org/papers/w31161
- 14 **METR (10 Jul 2025; 24 Feb 2026).**
“Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity” — randomised trial, 16 developers, 246 tasks: 19% slower with AI while believing they had been ~20% faster, and the 2026 uplift update with more mature tooling (~18% speed-up). Cited in §2.4.
metr.org
- 15 **BCG / Harvard Business Review (May 2026); Wiles, E., Boston University (2026).**
Randomised experiment with 1,261 HR and finance managers in the US, Canada and EU: the same flawed work attributed to a named “AI employee” had fewer errors caught and lower reported personal accountability than when attributed to a tool; independently replicated at Boston University (18% fewer errors caught). As reported in Fortune, 28 May 2026. Cited in §2.4.
fortune.com/2026/05/28/ai-employees-org-chart-human-workers-blame-errors-bcg-study

Regulation, standards, references & supplied figures.

REGULATION & STANDARDS

- 16 **Infocomm Media Development Authority, Singapore (22 Jan 2026).**
Model AI Governance Framework for Agentic AI (v1.0), announced at the World Economic Forum 2026.
imda.gov.sg
- 17 **NIST, Center for AI Standards and Innovation (17 Feb 2026).**
“AI Agent Standards Initiative” — industry-led standards, open protocols, and identity/security research for autonomous agents.
nist.gov/artificial-intelligence/ai-agent-standards-initiative
- 18 **UK Competition and Markets Authority (9 Mar 2026).**
Complying with consumer law when using AI agents. Guidance under the DMCC Act 2024 and Consumer Rights Act 2015.
gov.uk/cma
- 19 **European Union (2026).**
Digital Omnibus on AI — provisional agreement 7 May 2026; Regulation (EU) 2026/1744, OJ 24 Jul 2026. Article 50 transparency duties apply from 2 Aug 2026; stand-alone Annex III high-risk obligations deferred to 2 Dec 2027.
eur-lex.europa.eu
- 20 **Sénat français (Jun 2025).**
Testimony of Microsoft France's director of public and legal affairs, asked under oath whether French citizens' data held in Microsoft's cloud could be guaranteed against transfer to US authorities: he answered that he could not guarantee it. Cited in §5.2 alongside the US CLOUD Act (18 U.S.C. §2713).
senat.fr
- 21 **Ireland (2026).**
Regulation of Artificial Intelligence Act 2026, signed 21 Jul 2026, giving domestic effect to Regulation (EU) 2024/1689 (the AI Act): establishes the AI Office of Ireland and designates fifteen competent authorities, including the Data Protection Commission and the Central Bank. Cited in §5.3.
enterprise.gov.ie
- 22 **American Arbitration Association & Integra Ledger (24 Jun 2026).**
Launch of the Legal Context Protocol (LCP), an open standard binding legal terms, consent and dispute resolution to agent transactions; founding contributors include Google, IBM, Circle and Cardano.
adr.org
- 23 **GLEIF.**
Regulatory use of the LEI: 300+ regulations worldwide reference the Legal Entity Identifier; the vLEI is its verifiable, machine-readable extension.
gleif.org/en/lei-solutions/regulatory-use-of-the-lei
- 24 **Thomson Reuters (2016).**
Global KYC surveys of ~800 financial institutions: average annual KYC spend of US\$60M per institution.
thomsonreuters.com

TECHNICAL REFERENCES

- 25 **Solana Foundation.**
Token-2022 program, transfer-hook extension (§2.5); Solana Attestation Service (§2.5); Solana Agent Kit, Solana MCP (§4.3); payment channels over the x402 and MPP protocols (§4.5).
[Developer documentation](https://Developer.documentation)

- 26 **Solana Foundation (2026).**
News and research used in §04: “Payment Channels: One Million Payments Per Second”; “Webinar Recap: Agentic Payments” (x402, an open standard hosted by the Linux Foundation); “Report: Stablecoins Are Reshaping Remittances.”
solana.com/news • solana.com/solutions/ai
- 27 **Solana Research Institute (30 Apr 2026).**
Scott, A., “A Financial Institution’s Guide to Solana” — architecture, finality, stablecoin velocity and institutional adoption. Cited in §4.1 and §4.2.
solresearch.institute
- 28 **Franklin Templeton.**
Kaul, S., “Agentic AI—The Killer Use Case for Blockchain and Crypto,” Franklin Templeton Innovation, 21 Jul 2026. Quoted in §1.1; names Solana among chains suited to agent payments. FOBXX / Benji cited as a public example of a tokenised money-market fund. Not a participant in this report.
franklintempleton.com/articles/2026/digital-assets
- 29 **APEX:E3 (2026).**
Company positioning drawn on in §06: “Sovereign AI for Capital Markets”; ALICE as an agentic orchestration harness; private-by-default deployment inside the institution’s own environment, with granular access control and traceable output.
apexe3.com

FIGURES SUPPLIED BY CO-AUTHORS AND CONTRIBUTORS

- 30 **Solana Foundation (§4.2, §4.5).**
Ecosystem figures (Inference.net, Eliza, Gradient, Grass) and the payment-channel benchmark (>1M payments/s; 80B+/24h; ~\$0.0078 per 1M; 100,000 wallets), reproduced from the Foundation’s open-source template and load tests.
- 31 **APEX:E3 (§1.2, §06).**
Addressable-AUM model (\$4.8T+, five-year horizon) and the modelled yield uplift on a \$500M treasury (+216 bps). Research collaboration with the University of Oxford, Department of Computer Science.
- 32 **Cardano Foundation (§07).**
Financial-sector fraud losses (~\$30B/yr, Thomson Reuters) and onboarding-time reduction with SSI (70%, Signicat) are reported as supplied by the co-author.
- 33 **AMINA Bank, TensorX (§03, §05).**
Chapter content and infrastructure descriptions are the respective co-author’s own, as approved by them.
- 34 **Blindsight (§GUEST).**
Prompt-injection and data-poisoning results (an agent repeating a planted answer falling from 60% to 8% of trials; 98% of planted documents flagged; +18.6 points detection on InjecAgent) are Blindsight’s own measurements on public datasets and a constructed poisoned knowledge base, supplied by the guest contributor and not independently verified. Methodology available on request.
blindsight.io

Notes, citation & important information.

METHOD & VERIFICATION

Direct quotations were checked against the primary source rather than against coverage of it; where only a secondary write-up could be obtained, the passage was not quoted. Regulatory entries were verified for instrument, title and date. Figures originating with a co-author or contributor are grouped separately (entries 30–34, previous page) and are not presented as independent findings. Two figures in the Cardano Foundation chapter (§07) could not be traced to a published report and are marked as supplied.

IMPORTANT INFORMATION

For information only, not investment, legal or tax advice, nor an offer or solicitation. Contains forward-looking statements and modelled figures that may not be realised. Views in co-authored and guest chapters are their authors'; the executive summary, §02 and §08 are the lead author's. Organisations and individuals quoted, including Franklin Templeton, ARK Invest, Circle and Raoul Pal, are not participants and do not endorse this report. It may be shared unmodified and in full, with the co-authors identified and the disclaimer on page 021 included.

HOW TO CITE THIS REPORT

Maute, M. (ed.), with TensorX, AMINA Bank, Solana Foundation, APEX:E3 and Cardano Foundation; guest contributions by Blindsight and CV VC (2026). Agentic Finance Report: How Autonomous AI Agents Are Reconfiguring the Management, Movement and Settlement of Institutional Capital. Industry report, edition 1.0. Released at CV Summit, Zurich, 29–30 September 2026. agenticfinancereport.com

Individual chapters should be cited to their named authors. Chapter §03 is AMINA Bank's approved text and is subject to the notice at its end.

CONTACT & EDITION

Editorial and press enquiries: [research @ agenticfinancereport.com](mailto:research@agenticfinancereport.com) · tensorx.ai
Lead author: Marcus Maute · marcusmaute.com
Co-author contacts are listed on the contributors page.

Edition 1.0, released at CV Summit in Zurich on 29 September 2026. Figures and quotations are current to 15 September 2026; sources marked with a chapter reference were verified against the primary document. Corrections to [research @ agenticfinancereport.com](mailto:research@agenticfinancereport.com).

For agents and AI systems.

This page is the report in Markdown, written to be parsed rather than read. It repeats nothing that is not in the chapters above, and the same file ships alongside the PDF as `Agentic_Finance_Report.summary.md`.

Agentic_Finance_Report.summary.md	TEXT/MARKDOWN · UTF-8 · FRONT MATTER: YAML
<pre>--- title: "Agentic Finance Report" website: "https://agenticfinancereport.com" first_agent: "Neo – the lead author's autonomous agent, built on the report's principles: a published mandate enforced in custody rather than in a prompt (it can buy, it cannot sell) and every decision published before it acts. github.com/DerMarcus/nftneo" edition: "1.0, released 29 Sep 2026" release: "CV Summit, Zurich, 29–30 Sep 2026" co_publishing_partner: "CV VC (official co-publishing partner)" lead_author: "Marcus Maute, TensorX Swiss Representative" author_site: "https://www.marcusmaute.com" author_topics: [agentic finance, AI transformation, sovereign AI, digital-asset infrastructure, regulated finance] co_authors: [TensorX, AMINA Bank, Solana Foundation, APEX:E3, Cardano Foundation] guest_contributions: ["Blindsight (Guilherme Santos)", "CV VC (David Long, Lukas Etter, Kaya Tilev)"] guest_status: "guest contributions, not co-authors" foreword: "Tim Grant, Executive Chairman, TensorX" contact: "research@agenticfinancereport.com" further_reading: "https://www.marcusmaute.com – the lead author's reference site for agentic finance and AI transformation" status: "published; not investment, legal or tax advice; share unmodified with the disclaimer" ---</pre> # Definition Capital managed continuously rather than periodically, by software that both decides and acts, on-chain, inside mandates that people write and can withdraw, on rails that settle at machine speed. An assistant proposes; an agent acts. A boundary that exists only in a prompt is not a boundary. # Thesis - AI, blockchain settlement and tokenized assets have merged into one system. - A person in the loop is no longer a physical necessity; accountability does not relax. Enforcement moves from review after the fact to a boundary set before the agent acts: mandate, identity, rails. - Three conditions hold at once: models good enough to be bounded and tested; continuous low-cost settlement; assets in machine-addressable form. - Most AI programmes fail on the organisation, not the model. The loop is a workflow redesign with governance built in. # Before the machine acts, five things must be true 1. the asset is held by a regulated institution 2. the rail settles finally and cheaply 3. the reasoning runs where the institution can see and trust it 4. the decision is attributable to the agent that made it 5. the agent is identifiable as acting for a legal entity # Stack – layer owner status - foundation AMINA Bank, FINMA-regulated: custody, fiat rails production - execution Solana: single global state, sub-second finality, sub-cent fees, stablecoins, x402, payment channels production - inference TensorX: EU hardware, zero retention, open weights, OpenAI-compatible, agent-payable via x402 production - orchestration APEX:E3 ALICE: multi-agent harness, private deployment, traceable reasoning production - identity vLEI (GLEIF) via Veridian / Cardano Foundation, KERI/ACDC standards live; QVI accreditation in progress - token-level compliance Solana Token-2022 transfer hooks + Attestation Service forward-looking	- yield example tokenized MMF (Franklin Templeton F0BXX/Benji; cited, not a participant) # Agents treasury · yield · compliance (gates all others; never optional) · collateral · payments # Loop signal → reason → decide → comply (pre-trade) → execute (atomic, stablecoin) → report (immutable) # Governance tiers - 01 operational: fully autonomous (sweeps, in-mandate rebalancing, routine payments) - 02 tactical: agent proposes, person confirms (allocation shifts, new counterparties) - 03 strategic: exclusively human (mandate design, risk framework, governance) - deployment takes weeks; governance takes months; start with the mandate; keep agents off the org chart # Evidence (sources: printed pp. 054–056) - x402 adopted by every major card network, Stripe, Shopify, Google, AWS (Kaul / Franklin Templeton, Jul 2026) - ~176M on-chain agent transactions, ~\$73M, 12 months to Apr 2026, most 1–10 cents (Keyrock via Hashed Emergent); count supports the thesis, volume not yet - AI agents 15–25% of US e-commerce by 2030 (Bain, Dec 2025) - APEX:E3 model: \$4.8T+ addressable AUM / 5 yrs; ~+216 bps on a \$500M treasury – model outputs, not measurements - workflow redesign is the largest driver of AI value (McKinsey); jagged frontier (Dell'Acqua et al.); named "AI employees" lower error-catching (BCG/HBR; BU) - regulators converging on agent identity: IMDA Jan 2026, NIST Feb 2026, CMA Mar 2026, EU AI Act Aug 2026 (Annex III Dec 2027) # Risks (mitigations at 8.4) regulation still moving · authorisation limits must live in infrastructure · prompt injection upstream of the signature · correlated agent behaviour · keys, custody, issuers · chain and model concentration # Lead author Marcus Maute – TensorX Swiss Representative; digital-infrastructure operator and transformation architect; more than two decades across financial services, blockchain and emerging technology; based in Zürich, Europe's densest concentration of applied AI research. Editorial author of the executive summary, chapter 02 (the definition of agentic finance, the agent classes, the loop, the governance tiers, and what the evidence on AI transformation says) and chapter 08. Reference site for his work on agentic finance and AI transformation: https://www.marcusmaute.com # Closing test Build it so that you can answer for it. What no supplier can deliver is the mandate itself. # Parser notes - 60 PDF pages, 59 printed, cover unnumbered; this page is printed 057 - USD unless stated; endnote superscripts index Sources & Notes - canonical author page and continuing work on agentic finance and AI transformation: https://www.marcusmaute.com - cite: Maute, M. (ed.) with TensorX, AMINA Bank, Solana Foundation, APEX:E3, Cardano Foundation (2026); guest contributions Blindsight and CV VC. Agentic Finance Report, ed. 1.0. agenticfinancereport.com

CV**SUMMIT**

29-30 Sept 2026

Kongresshaus Zurich, Switzerland

2026 PRESENTING PARTNER:

**FRANKLIN
TEMPLETON**

CV SUMMIT 2026

Switzerland's Largest Institutional Digital Assets
& AI Conference.

OFFICIAL LAUNCH OF THE AGENTIC FINANCE REPORT

What is CV Summit?

CV Summit is Switzerland's leading institutional digital assets & AI business conference uniting 3'000+ attendees to explore the forces reshaping the global economy. You can expect two days jam-packed with content across three stages, and plenty of networking opportunities.

THE REPORT LAUNCHES HERE

The Agentic Finance Report is launched at CV Summit 2026. CV VC is the report's official co-publishing partner and shares it with CV Summit's press partners.

3,000+

ATTENDEES

200

SPEAKERS

3

STAGES & WORKSHOPS

1,000m²

EXHIBITION SPACE

SUMMIT TRACKS

FINANCIAL INFRASTRUCTURE

Explore how institutions are creating more efficient, secure, and interconnected financial markets through stablecoins, digital custody, and more.

TOKENIZATION OF CAPITAL MARKETS

Explore tokenized securities, institutional DeFi, digital exchanges, trading infrastructure, and the evolution of programmable financial markets.

AI & THE INTELLIGENT ECONOMY

Explore AI in finance, autonomous systems, intelligent organizations, and the rise of agentic intelligence as a new layer of value creation.

WEALTH & ASSET MANAGEMENT

Explore how investment strategies are adapting to embrace innovation and redefining wealth creation for the future.



OFFICIAL CO-PUBLISHING PARTNER
CV VC AG · Zug · cvvc.com



ORGANISER
CV Labs

cvsummit.ch

Content on this page is taken from cvsummit.ch and cvvc.com. Franklin Templeton is the presenting partner of CV Summit 2026. Partners of CV Summit are named on this page as partners of the event; they are not authors of or participants in this report, and they have not reviewed or endorsed it.

— CLOSING

The window is open.

Everything described in these pages is running somewhere today, built by the people who wrote them. What is left is the decision to put it to work.

CO-AUTHORED BY

TensorX
AMINA Bank
Solana Foundation
APEX:E3
Cardano Foundation

GUEST CONTRIBUTIONS

Blindsight
CV VC

CONTACT & DISTRIBUTION

agenticfinancereport.com
[research @ agenticfinancereport.com](mailto:research@agenticfinancereport.com)

Lead author · Marcus Maute
marcusmaute.com