



Agentic Finance and Stablecoins: Everything to Know About the New Stack for Autonomous Commerce

Autonomous AI agents are beginning to move money. This report was designed to help builders, financial institutions, and product teams understand the infrastructure that makes that possible. It covers agentic payment protocols, how blockchain rails and stablecoins fit in, security models, and the practical decisions teams face when building for a world where agents, not humans, initiate transactions.

This report was built to help founders, fintech teams, and developers make sense of the emerging agentic finance stack. It covers how AI agents, stablecoins, and blockchain infrastructure are converging to enable autonomous commerce, and what that means for the teams building on top of it. Whether you're exploring agentic payments, treasury automation, or agent-first applications, this guide provides the context and frameworks to help you build with conviction.

Featuring Insights from Leaders at:



Table of Contents

I. An Introduction to Agentic Commerce	
What Agentic Means and Why It's Important	04
The Convergence: AI Agents, Stablecoins, and Blockchain Rails	05
Algorithmic Trading Was Just the Beginning	06
Agentic Trends We're Watching	07
Agentic Market Map	09
II. The Agent Builder Landscape	
The New Developer Toolkit: Claude Code, OpenClaw, and More	10
Where Agents Are Now and Where They're Heading	11
Agent-First Web Pages: How the Web Is Being Rebuilt for Machines	12
III. The Agentic Payment Protocol Landscape	
x402: The HTTP-Native Payment Layer	13
The Agentic Commerce Protocol (ACP)	14
Refunds, Disputes, and Trust in Agentic Payments	15
IV. Security and Guardrails in an Agentic World	
The Threat Model for Agentic Finance	16
Policy Engines: Defining What Agents Can and Cannot Do	17
Guardrails at the Model Layer vs. the Infrastructure Layer	18
V. The Consumer Frontier	
How Agentic Finance Is Emerging for Everyday Users	19
Rethinking Wallet Infrastructure in an Agent-First World	20
Consumer Use Cases: Subscriptions, Payouts, Prediction Markets, and More	21
VI. The Institutional Frontier	
How Institutions Are Adopting Agentic Infrastructure	22
Algorithmic Trading to Agentic Trading: What Changed	23
How Stablecoins Enable the Agentic Financial System	24
VII. The Infrastructure Stack: What Agents Actually Need	
Identity and Wallet Architecture for Agents	25
Blockchain Networks as the Operational Layer for Agents	26
Data Availability, Audit Trails, and Travel Rule in an Agentic World	27
VIII. How to Build for the Agentic Economy	
Key Decision Criteria for Teams Building Today	28
Where Dynamic Fits Into the Agentic World	29
Additional Resources	30

An Introduction to **Agentic Commerce**

What Agentic Means and Why It's Important

The last decade of artificial intelligence in finance meant better recommendations: smarter fraud detection, personalized product suggestions, improved underwriting models. AI was a layer on top of financial systems, not a participant within them, but this is quickly changing.

The term "agentic" refers to AI systems that don't just analyze or suggest, they act. An agentic AI can be given a goal, a set of tools, and a budget to autonomously make decisions, call APIs, and execute tasks to accomplish that goal. It doesn't wait for human approval at every step. It operates in real time across platforms.

When those tools include the ability to move money, the implications are profound.

An agent can:

- ▲ Renew a software subscription before it lapses
- ▲ Rebalance a treasury position in response to market conditions
- ▲ Pay a contractor the moment a task is completed, verified by code
- ▲ Negotiate access to a data API, pay the fee, retrieve the data, and use it, all without human intervention

This is just scratching the surface of possibilities. The tools that enable this behavior, including autonomous agent frameworks and agentic payment protocols, are either live today or in active development by the largest companies in the world.

The infrastructure for agentic commerce is being built right now, and the teams that understand it earliest will have a significant advantage.

“

The main consumers of the future internet are going to be agents making thousands of calls an hour. Much of the first principles of how we built financial systems on the internet for humans will need to be redesigned. API Keys, subscriptions, credit cards, KYC are all points of friction. Stablecoins as not just digital dollars, but programmable money that posits them as an interesting solution. With new agentic payment rails like x402 and MPP, an agent hits a paywall, signs a stablecoin transfer, and gets the data back — one HTTP round-trip, no account, no invoice. You can charge a fraction of a cent per request and settle it onchain, something cards have struggled with.

There is rapidly growing interest and activity on Solana when it comes to agentic commerce. Late last year, x402 transactions crossed 1% of all non-vote transactions on the network. Still early, but a real signal when you consider the scale Solana operates at. We are actively building thinking of the agent as the future consumer. The focus on sub-cent fees, fast finality, and a deep stablecoin ecosystem lend well to builders who need payment rails that can keep up with agents. We expect rapid growth as more API providers adopt pay-per-request models and more agents get wallets.




Rishin Sharma

Head of AI Growth at Solana Foundation




The Convergence: AI Agents, Stablecoins, and Blockchain Rails

Three forces are colliding to create the infrastructure layer for agentic commerce, and their convergence marks a genuine shift in how economic activity works:




AI Agents are the new economic actors.

Tools like Claude Code have moved agents beyond assistance and into execution: writing code, calling APIs, managing workflows, and interacting with financial services autonomously. The distinction matters: these are not tools that assist human decisions. AI agents are decision-makers themselves.



Stablecoins are the native currency of this emerging economy.

Unlike traditional payment rails, which require bank accounts, business-hour availability, settlement delays, and human authorization, stablecoins settle instantly, operate 24/7, and are programmable by design. An agent doesn't need a bank account to hold USDC, it simply needs a wallet.



Blockchain rails serve as the open infrastructure that makes all of this possible and composable.

An agent can pay for an API call, receive a result, trigger a smart contract, and settle the entire transaction in a single uninterrupted flow, on a network that no single institution owns or controls. No intermediaries, no delays, no business hours.

Algorithmic Trading Was Just the Beginning

The clearest historical parallel to this moment is the rise of algorithmic trading. Before algorithmic trading, financial markets were human. Decisions were made by traders, executed by brokers, and settled through layers of manual process. Speed was limited by human reaction time. Scale was limited by human capacity.



Algorithmic trading changed that.

When it became possible to write code that could analyze a market and place a trade in milliseconds, the nature of markets changed. Liquidity deepened, arbitrage opportunities closed faster, and market makers moved from trading on intuition to trading on math.



Agentic AI represents a similar shift, but broader.

Algorithmic trading was confined to financial markets, but Agentic AI extends to every economic activity: commerce, data markets, software services, content, professional tasks, and more.

“

The journey to algorithmic trading started when open outcry trading floors were gradually replaced by computer-based trading. Computers were able to replicate both market data, effectively the eyes and ears of the trading floor, as well as the execution of trades that had been in person and noted on paper. As quotes, orders and executions became data, so algorithms were able to first analyse this data and then execute upon it.

Digital assets do two things to this. First, it enables the assets, the execution and the settlement to migrate to the open internet, rather than closed systems comprising co-location data centres and other centralised databases. It also enables the capability gap between trading and settlement to close. Where algorithmic trading collapsed the speed and cost of execution to near zero, so digital assets offer the opportunity to repeat those improvements with payments and settlement in general.



Nick Philpott

Co-founder & Chief Operating Officer at Zodia Markets



Agentic Trends We're Watching

We're closely watching how the agentic economy develops. Here are some of the key trends our team is tracking in 2026:

Trend	Context	Why it matters
 Compliance Agent Teams	Forward-thinking institutions are beginning to deploy compliance agents as a parallel workforce: agents that monitor transaction flows in real time, flag anomalies, run sanctions screening, and generate regulatory reporting autonomously. As the reliability of compliance agents increases, the operational and cost advantages will make this shift difficult for institutions to ignore.	Compliance functions that currently require large teams operating in shifts can be augmented or restructured around agents that never sleep or never miss a transaction.
 Agentic Creators	The creator economy is becoming one of the earliest proving grounds for consumer-facing agentic finance. Agentic creator platforms are beginning to emerge, where AI agents manage content strategy, collect micropayments via stablecoins, and handle royalty distribution autonomously.	The infrastructure requirements are unlike anything in traditional creator monetization: high-frequency, low-value transactions across platforms and currencies, with human attribution built in. It is one of the clearest near-term stress tests for the payment rails being built right now.
 Agent-to-Agent Commerce	The most interesting economic activity will increasingly happen not between humans and agents, but between agents transacting directly with each other. Agents can negotiate with, purchase from, and settle payments to other agents entirely programmatically, with no human touchpoint required at any step.	Agent-to-agent commerce creates an entirely new layer of economic activity that existing financial and legal infrastructure was not designed to support. Questions of liability, counterparty identity, and pricing discovery all need new answers at machine speed. The protocols and infrastructure that solve these challenges earliest will become the default rails for a layer of commerce that is only just beginning to take shape.

Trend	Context	Why it matters
 Agentic Treasury Management	Corporate treasury is shifting from rules-based automation to fully autonomous operation. Agentic treasury systems reason across live market conditions, rebalance positions in real time, manage cross-border settlements without waiting for business hours, and deploy idle capital into yield-bearing instruments.	Treasury management is one of the most capital-intensive and operationally complex functions in any institution, and it has historically been constrained by human availability and manual process. Agentic systems remove both constraints, opening the door to treasury operations that are faster, leaner, and more responsive to market conditions.
 Regulatory Frameworks	Every major jurisdiction is grappling with the same question: who is responsible when an AI agent transacts, makes an error, or causes harm? The regulatory frameworks that answer this are beginning to take shape across the US, EU, and Asia. When they arrive, they will move fast and reshape which infrastructure providers can operate at institutional scale.	Firms and infrastructure providers that have designed compliance from the beginning will have a decisive advantage. Those that have not will face significant retrofit costs, or worse, find themselves unable to operate in regulated markets at all.

“

The most important trend is not any single use case. It is the speed at which the infrastructure requirements are compounding. Six months ago, the conversation was about whether agents would transact at all. Now the conversation is about multi-agent coordination, cross-chain settlement, and real-time compliance for autonomous systems.

The gap between where most teams are building today and where the requirements are heading is widening, not closing. That is both a risk and a massive opportunity.



Itai Turbahn


Co-founder & CEO at Dynamic



a Fireblocks company

Market Map

March 2026 Agentic Commerce Market Map



“

Although transaction activity is down from the highs at the end of 2025, the agentic commerce stack is taking shape faster than most people realize. Six months ago, there were less than 100 merchants monetizing via x402 or MPP – today, that number is close to 5,000. The same growth is happening with the rest of the agentic payment stack and it’s not just crypto-native developers experimenting anymore.

Established players like Google, Visa, and AWS are extending their existing infrastructure into the agentic stack, while a new class of startups (e.g., Merit Systems, Sponge, Sapiom) is being built natively for machine-speed commerce. The companies that win will be the ones positioned where value concentrates: the interface that controls user intent, the payment rails that sit in the money flow, and the compute infrastructure that powers everything underneath.



Lucas Shin
Data & Research at Artemis Analytics



The Agent Builder Landscape

The New Developer Toolkit: Claude Code, OpenClaw, and More

Before an agent can transact onchain, it has to be built. A new generation of developer tools has emerged, making it increasingly easy to build and deploy autonomous agents:

Claude

is Anthropic's agentic coding tool, designed to operate directly in a developer's terminal.

Claude Code can autonomously edit files, run tests, navigate codebases, and execute multi-step engineering tasks. These capabilities allow an agent to manage a treasury position, execute a payment workflow, or interact with a financial API.

Devin

built by Cognition AI, acts as an autonomous AI software engineer.

Devin can set up its own development environment, write and debug code, browse the web for documentation, and complete engineering tasks end-to-end.

Codex

is OpenAI's cloud-based agent.

Codex can handle tasks in parallel, meaning multiple instances can work on different parts of a codebase simultaneously. It can run tests and operate asynchronously, reporting back to developers when tasks are complete.

OpenClaw

is an open-source framework that gives developers fine-grained control over agent spending and permissions, with full audit logging.

It sits between the agent and its wallet, letting developers set precise spending boundaries, scope transaction types, and revoke access instantly if something goes wrong.

“

The teams building agentic finance products today need to move quickly across complex infrastructure, from payment rails to financial APIs to stablecoins. For those teams, Devin is the agent that builds the agents. It handles multi-step engineering work autonomously and end-to-end, so small teams up to large financial institutions can ship fast while ensuring quality, security, and scalability.



Scott Wu

Co-founder & CEO at Cognition AI



Where Agents Are Now and Where They're Heading

Today, agents can reliably execute bounded tasks:

- ▲ Writing and running code
- ▲ Calling APIs
- ▲ Browsing the web
- ▲ Summarizing documents
- ▲ Managing calendars
- ▲ Sending communications

In financial contexts, agents are already being used for research aggregation, report generation, and basic workflow automation. Some are transacting, purchasing API credits, managing budgets, and executing predefined payment flows.

The more consequential capabilities are still being built: agents that can manage multi-step financial workflows, interact with onchain protocols, negotiate access to data markets, and operate across multiple platforms simultaneously. These are not distant milestones. They are active areas of development, and the infrastructure being built today is designed with them in mind.

The next phase isn't just smarter individual agents. It's agents that operate in coordinated groups, where a financial agent executing a transaction is joined in real time by compliance and risk agents that validate, flag, and audit alongside it. The industry is beginning to call these configurations agent swarms.

“

As institutions and agents move onchain, they're onboarding onto infrastructure built for crypto-native retail users. What they need is production-grade execution with auditable trails, and controls that hold when no human is reviewing each transaction.

On the data side, the same gap exists: institutions need feeds they can reconcile and report from, and agents route around anything inconsistent automatically. What both require is a data layer that's standardized, auditable, and accurate enough that you can build logic on top of it — whether that's an agent executing autonomously or an institution reconciling positions.



Ethan Chan

Co-Founder and CEO at Allium



Agent-First Web Pages: How the Web Is Being Rebuilt for Machines



We're building for the agent experience because it's no longer hypothetical. More than 40% of traffic to docs powered by Mintlify now comes from AI agents. That changes everything about how documentation should be structured, served, and optimized.

We're making docs machine-readable by default, with structured endpoints, semantic metadata, and response formats that agents can parse and act on, so every site page works as well for an LLM as it does for a developer.



Han Wang

Co-founder & CEO at Mintlify

The web was built for humans. Pages render visually, navigation requires clicks, and content is structured to be read, not parsed. But agents interact with none of it the way a person would. This is not just a technical challenge, it's a design philosophy shift.

For agents to transact on the web, to purchase access to a service, retrieve data in exchange for payment, or interact with an API endpoint, the web needs to become machine-readable at a deeper level. Agent-first web architecture means building services that expose structured, authenticated endpoints rather than just human-facing UIs.

It means defining clear payment terms that a human or machine can parse, accept, and fulfill without ambiguity. It means thinking about identity, authorization, and payment at the API layer, not just the frontend. But agent-first design also encompasses broader considerations: how services advertise their pricing, how they verify the identity of the agent making a request, how they handle authorization scopes and spending limits, and how they structure responses so agents can parse and use them efficiently.

The companies and protocols building this infrastructure are, in effect, rebuilding the commercial layer of the internet, replacing the "add to cart" button with a machine-negotiable payment handshake.

The result is a web that humans can still use, but was no longer built with them as the primary user in mind.

The Agentic Payment Protocol Landscape

x402:

The HTTP-Native Payment Layer

AI agents are on the verge of becoming first-class financial participants on the internet, and they need payment infrastructure built for them. x402 is designed to make that easy. Developed by Coinbase and Cloudflare, the protocol revives the dormant HTTP 402 "Payment Required" status code as functional payment infrastructure.

The flow is simple: a client requests a resource, the server returns payment terms, the client pays, a facilitator handles settlement, and the resource is released. At no point does the server interact with the blockchain directly.

This means regular web developers can monetize API endpoints, data feeds, or digital services without running blockchain nodes or writing smart contracts. For agents, it is even more significant. x402 provides a standardized machine-readable payment protocol that any agent can implement to pay for any x402-compatible resource.

Now supporting credit and debit cards alongside stablecoins, x402 lets agents evaluate which rail fits the transaction. Cards remain better suited for e-commerce today, while stablecoins offer advantages for micropayments and cross-border accessibility. The agent decides.

The numbers back it up. By the end of 2025, x402 had processed over 100 million payments, with roughly \$600 million in annualized payment volume across supported blockchains. Solana has accounted for the majority of x402 payment volume, with its market share increasing since 2026 began.

The value in the x402 ecosystem does not sit in the open-source protocol itself, however. It concentrates at the facilitator layer: the entity that handles blockchain interaction, verifies payments, and bridges the complexity between web servers and onchain settlement.

 **Combinator**

“

Agents are becoming first-class financial participants, buying goods, accessing paid APIs, and transacting with other agents autonomously. We designed x402 to make that easy: every merchant and API on the web becomes payable by AI agents, whether they settle in stablecoins or credit cards. The agent decides which rail fits the transaction, not the developer.



Nemil Dalal

Visiting Partner at Y Combinator &
Founding Team of x402

The Agentic Commerce Protocol (ACP)

The Agentic Commerce Protocol (ACP) is built for instant checkout within conversational interfaces. Where x402 operates at the API layer, ACP approaches agentic payments from the application and consumer layer. The focus is UX and speed: reducing the friction between an agent identifying a product or service and completing the transaction.

ACP represents the first major integration of an LLM with a consumer checkout flow. The bet is that the fastest path to agentic commerce runs through the interfaces people already use, not through new infrastructure they have to adopt.

When an agent identifies a product or service within a supported interface, ACP handles the entire checkout flow without redirecting the user to an external page or requiring manual input. Payment credentials are stored and permissioned in advance, and the transaction completes within the conversation. The experience is designed to be invisible: the agent transacts, and the user is informed rather than interrupted.

ACP and x402 are targeting different layers of the same problem. ACP owns the consumer checkout experience inside AI applications. x402 owns the crypto, API, and enterprise settlement layer. Both are racing to establish defaults in their respective domains, and both are moving fast.

A Spotlight on MoonPay Agents:

Most commerce infrastructure today is built for humans. MoonPay Agents addresses the gap that remains: enabling agents to act on checkout flows they can reach but cannot complete. Rather than introducing a new checkout interface or settlement layer, MoonPay Agents acts as a bridge between existing checkout flows and AI agents — converting traditional checkout flows into programmatic payments that agents can execute directly via API.

In effect, ACP handles discovery and checkout intent, x402 enables settlement, and MoonPay connects the two, allowing transactions to complete end-to-end within the agent loop.

“

The next wave of payments won't be driven by humans. It will be driven by agents. That shift requires transaction systems that are fast, programmable, and global by default, with the ability to move from intent to execution.



Ivan Soto-Wright

Co-founder & CEO at MoonPay



Refunds, Disputes, and Trust in Agentic Payments

Decades of consumer payment infrastructure were built around a single assumption: that things go wrong. Chargebacks, fraud protection, dispute resolution, reversibility. The system was designed to be correctable because humans make mistakes and merchants sometimes behave badly.

Agentic payments are fast and final by design. A stablecoin transaction settles in seconds, there is no bank to call, and there is no chargeback mechanism built into the protocol. This creates a meaningful trust gap, and it is one of the least-discussed challenges in the agentic finance conversation.

None of the existing frameworks map cleanly onto what happens when an agent is the party at fault, or the victim:

- ▲ If an agent overpays for a service due to a pricing error, who is liable?
- ▲ If a fraudulent API returns bad data after collecting payment, what recourse exists?
- ▲ If an agent is compromised and makes unauthorized payments, how are those resolved?
- ▲ If an agent operates across jurisdictions, which legal framework governs the dispute?

Smart contract escrow, protocol-level dispute resolution, and onchain reputation systems for service providers are all being explored as partial solutions. They represent genuine progress, but the space remains fragmented, none are standardized, and none yet meet the threshold that serious financial services adoption will require.

“

In agentic payments, the accounting layer stops being back-office infrastructure and becomes the trust layer itself. There are no chargebacks. There is no bank to call. The only way to understand what happened, whether an agent overpaid, was defrauded, or made an unauthorized transaction, is through a complete and accurate financial record. That's harder than it sounds.

Agents operating at machine speed across multiple chains generate fragmented, non-linear histories where a single action can produce several on-chain events at once. The solution is proper, accurate financial data. The ability to reconstruct a full picture across every chain and counterparty, at the speed and scale autonomous systems demand. In this world, accurate financial records aren't just useful for reporting. They are the dispute resolution mechanism.



Tal Zackon

Co-Founder & CEO at TRES Finance



Security and Guardrails in an Agentic World

The Threat Model for Agentic Finance

The threat model for agentic finance is distinct from traditional cybersecurity or financial fraud. The most dangerous failures are not unauthorized access, they are authorized misuse. Here are key threat vectors to be aware of:

Prompt injection is when an agent that browses the web, reads documents, or processes external data can be manipulated by malicious content embedded in that data.

In a financial context, a prompt injection attack might cause an agent to send funds to an attacker's wallet while believing it is executing a legitimate transaction. Defense against prompt injection requires both safeguards and controls that validate transactions against an approved whitelist before execution.

Runaway agent behavior is a failure where an agent operates within its authorized scope but produces outcomes that were not intended.

This could be unexpectedly hitting spending limits, draining a budget through a recursive loop, or compounding errors in multi-agent workflows. While the intent may not be malicious, these errors can cause material damage. Policy engines with hard spending caps, rate limits, and anomaly detection are the primary defense.

Agent impersonation occurs when a malicious actor presents itself as a legitimate agent.

Without robust identity and authentication at the agent layer, a compromised or spoofed agent can make unauthorized payments that look like legitimate authorized transactions. The solution gaining traction is cryptographic agent identity, where signed credentials bind a specific agent to a principal and a defined set of permissions.

Key compromise remains the fundamental risk in any blockchain-based model.

An agent that has access to a private key is a target. Key management approaches that eliminate single points of failure, such as TSS-MPC with independent recovery, are the best response to this risk.

Supply chain risk applies to the agent builders themselves.

An agent built using a third-party SDK, model API, or tool integration inherits the security properties of every component in its stack. Teams building financial agents need to apply the same scrutiny to their software supply chain that they would to any financial system.

Policy Engines: Defining What Agents Can and Cannot Do

Policy engines are the set of rules that are enforced at the infrastructure layer and govern what an agent is authorized to do. While an agent's judgment can be manipulated or can fail, a policy engine enforced at the wallet layer cannot be overridden by the agent itself.

A policy engine that enforces spend limits and counterparty restrictions is necessary but not sufficient. A well-designed policy engine for agentic finance covers several dimensions:

▲ Spending limits

Spending limits define the maximum amount an agent can transact in a single operation, within a time period, or in total across a session.

▲ Counterparty allowlists

restrict which addresses or accounts an agent can send funds to, preventing misdirected payments regardless of what instructions the agent receives.

▲ Transaction type restrictions

limit the categories of operations an agent can perform, such as allowing it to pay for API access but not to transfer funds to other wallets.

▲ Time-based controls

confine agent activity to defined windows and trigger human review outside of those windows.

▲ Escalation thresholds

automatically pause execution and route to human approval when a transaction exceeds a defined size or falls outside normal patterns.

“

Policy engines are not new. Financial institutions have enforced transaction controls, approval hierarchies, and counterparty restrictions for decades. What is new is the speed at which decisions are being made and the fact that the entity making them is not a human. When an agent can initiate hundreds of transactions per minute, the enforcement mechanism has to operate at the same speed and at a layer the agent itself cannot override.



Adam Zion

Head of Product at Dynamic



Guardrails at the Model Layer vs. the Infrastructure Layer

A common mistake in agentic security is treating the agent's own judgment as a substitute for infrastructure-level controls. They serve different purposes, and both are necessary.

Instructions, fine-tuning, and system prompts shape how an agent reasons and what it is willing to do. They are essential for normal operation, but they are not a reliable security boundary. A sophisticated prompt injection or an unexpected input can cause an agent to behave outside its intended parameters. These controls are a first line of defense, not a last one.

Infrastructure-layer controls like policy engines and transaction screening operate independently of the agent's own reasoning. An agent that is compromised via prompt injection and instructed to send funds to an attacker cannot complete that transaction if the attacker's address is not on the counterparty allowlist enforced at the custody layer.

The practical implication is straightforward. An agent should never be the last line of defense against its own failures. Policy controls enforced at the infrastructure layer, independent of whatever the agent reasons or receives as input, are what make agentic finance systems resilient rather than just functional.

This principle scales. As agent deployments move toward coordinated groups, where transacting agents operate alongside dedicated compliance and risk agents, the infrastructure layer becomes the common authority that governs all of them.

“

In traditional finance, no one asks the trader to also be the compliance officer. The same separation needs to exist for agents. While the agent reasons and acts, the infrastructure layer decides whether that action is permitted. Collapsing those two functions into one is a fundamental design flaw.

What we have learned from working with institutional clients is that the teams building the most resilient systems assume the agent will eventually do something unexpected. They design around that assumption by making the policy layer independent, deterministic, and impossible for the agent to bypass.



Shahar Madar

VP Security & Trust Products at Fireblocks




How Agentic Finance Is Emerging for Everyday Users

For consumers, the entry point into agentic finance begins with a conversation. AI assistants embedded in chat applications, smart features inside fintech products, and subscription management tools operating in the background are where the first wave is taking shape. These experiences feel less like new technology and more like having someone capable handle things you never had time for.

For example, a user can tell their AI assistant to find and cancel any subscriptions they have not used in the last three months. The agent will query their transaction history, identify the relevant subscriptions, and cancel them. They may even go as far as to negotiate retention offers on your behalf. What would have taken twenty minutes of form-filling and cancellation flows took a single instruction.

Consider a creator platform where contributors are paid in real time as their content is consumed. As streams accumulate, payments route automatically, wallets are managed in the background, and balances update continuously. The creator simply sees their balance grow without waiting periods, invoices, or any manual intervention.

In each of these instances, blockchain complexity is entirely invisible. The user never has to manage a wallet or think about what's happening on the backend. The agent handles the infrastructure layer entirely, leaving the user to deal only in intentions and outcomes. The technology recedes, and what remains is simply a more capable version of something familiar.



Rethinking Wallet Infrastructure in an Agent-First World

Fireblocks

“

Wallet infrastructure for agentic finance can be simple, but only when it is built on top of secured infrastructure that moves trillions at machine-to-machine speed, with the right enforcement mechanisms in place.

Spending limits, transaction type restrictions, audit trails without a human in the loop: these are institutional custody problems. That is exactly where Dynamic operates, and it is why the wallet layer deserves more attention than it typically gets.



Ran Goldi

SVP Payments & Network at Fireblocks

The wallet is the foundational primitive of any blockchain-based application. It is where identity, authorization, and value converge, and the industry has spent years working to make that interface accessible enough for mainstream users to engage with confidently.

Embedded wallets, social login, and passkey-based signing have all moved the needle. Agentic finance moves the requirement entirely.

However, this also requires a new set of design primitives. An agent-capable wallet needs to support scoped authorization. That means granting an agent permission to spend within defined limits, on specific transaction types, within a defined time window, and without handing over full control.

Scoped authorization is the foundation, but it is not sufficient on its own. Revocability ensures a user can immediately withdraw an agent's spending authority if something goes wrong. Transparency means clear, human-readable logs of what an agent has spent, when, and why. And recovery provides a mechanism for retrieving funds if an agent is compromised or behaves unexpectedly, without requiring the user to understand any of the complexity.

These requirements push wallet infrastructure beyond what most embedded wallet SDKs were designed to support. They require the kind of policy-engine thinking that institutional custody providers have developed for enterprise clients, applied, for the first time, to consumer products.

Consumer Use Cases: Subscriptions, Payouts, Prediction Markets, and More

A few consumer use cases have gained real traction so far, though the surface area is growing fast:

▲ Prediction markets

Polymarket has processed over \$9 billion in trading volume and is one of the clearest existing examples of agentic finance operating at scale. Arbitrageurs, most of them automated bots, made \$40 million in a single year by identifying and exploiting mispriced contracts faster than any human could.

▲ Lending and credit

AI-driven underwriting is beginning to reach consumer crypto lending. Platforms like 3Jane are automating credit underwriting entirely through smart contracts, using verifiable records of financial standing to set rates and enforce debt covenants without human review.

▲ Creator and gig economy payouts

For a freelance and creator workforce earning across multiple platforms, agents that handle routing, wallet management, and currency conversion in the background solve a largely unaddressed pain point. Audius already does this for music, routing 90% of earnings directly to artists in real time as their content is consumed, with no monthly payout cycle and no intermediary taking a cut.

▲ Subscription management

Tools are already handling cancellation flows autonomously, negotiating retention offers in real time based on individual usage history and pain points. As more software services shift toward pay-per-use pricing, agents that continuously optimize what you pay and for what will become significantly more valuable.

▲ In-game rewards

Web3 gaming platforms are deploying agents to manage in-game economies, distribute rewards, and handle asset trading on behalf of players. Virtuals Protocol has tokenized AI agents that operate as gaming NPCs, trading bots, and research assistants, with communities able to co-own and govern individual agents.




Consumer finance tools have gotten smarter, but for years the pipes underneath them stayed the same. Stablecoins change that by integrating programmable rules into money movement. The shift we're seeing at Rain is less a fintech evolution and more an infrastructure unlock, where payments that once required manual coordination can now be built directly into the product flow. Every meaningful consumer finance use case we're seeing emerge, from subscription management to real-time creator payouts, shares the same foundation of stablecoins as the settlement layer.



Charles Yoo-Naut

Co-founder & CTO



How Institutions Are Adopting Agentic Infrastructure

Institutional finance has always been an early adopter of automation. From algorithmic trading to high-frequency market making to rules-based risk management, the drive to remove latency and human error from financial operations is not new. What is new is the quality of the reasoning those systems can now apply, and the range of tasks they can now handle autonomously.

The consumer agentic finance story is about making financial tasks easier for individuals. The institutional story is about making markets more efficient, capital more productive, and operations more automated at a scale that was previously out of reach.

The implications extend well beyond trading desks. Treasury management, collateral optimization, liquidity deployment, and counterparty risk assessment are all functions that institutional teams currently perform with significant manual overhead. Agents that can execute across all of these simultaneously, within defined policy parameters and in real time, represent a meaningful operational shift for any institution running meaningful capital.

Quantitative trading firms and market makers including B2C2 have been running automated strategies for years, but the nature of those strategies is changing. Classical algorithmic trading operated on predefined rules: if price crosses X, execute Y. Agentic trading infrastructure operates on reasoning, taking in market conditions, available liquidity, risk parameters, and portfolio position to determine the optimal action rather than simply executing a predetermined one.



The barrier to massive crypto adoption was never the technology. It was the cognitive load. When agents absorb that complexity and the user simply says what they need, the entire trust equation changes. You stop being a crypto user and start being someone who moves money, saves smarter, or gets paid instantly across borders. For businesses, the transformation is even more evident — powered by stablecoins, becoming the invisible rail behind global payments even when people don't know they're using them. That's where we see the biggest opportunity in Latin America.



Nano Rodríguez

Strategic Alliances Director at Bitso

Algorithmic Trading to Agentic Trading: What Changed



The history of market making provides a useful framework for understanding what is happening now.

In the era before algorithmic trading, market makers operated through human judgment and relationship-based price discovery. Speed was measured in seconds, and strategies were implemented by people.

Algorithmic trading changed the unit of competition from human judgment to code quality and latency. The firms that won were the ones with the fastest connections, the tightest spreads, and the most efficiently implemented strategies. Speed moved from seconds to milliseconds.

The transition to agentic trading changes the unit of competition again, from latency to intelligence. The question is no longer just how fast you can execute a predefined strategy. It is how well your agent can reason, adapt its behavior, and pursue complex multi-step objectives across multiple platforms simultaneously.

For crypto market makers specifically, this creates new dimensions of opportunity. Stablecoins enable settlement at any time across jurisdictions. Blockchain infrastructure makes liquidity pools and onchain order books programmable. And the programmability of DeFi protocols means an agentic market maker can interact not just with centralized exchanges, but with the entire onchain liquidity landscape.

“

B2C2 is actively bridging the gap between high-frequency execution and the next generation of agentic trading. We've successfully integrated AI-enhanced algorithms to achieve tighter spreads and immediate yield optimization today. Also, at the point of execution, we're uncovering new efficiencies beyond pure pricing to directly benefit our clients.



Edmond Goh

Global Head of Trading at B2C2



How Stablecoins Enable the Agentic Financial System

Autonomous agents need a settlement layer that operates on their terms. For institutional agentic finance, stablecoins are currently the only asset class that qualifies.

Consider the alternative. A corporate treasury agent operating on traditional rails must work within banking hours, manage correspondent banking relationships, handle currency conversion through FX desks, and wait for settlement cycles that can span days. Every one of those dependencies is a point where autonomous operation breaks down, where the agent must pause and wait for a human process to complete.

Stablecoins eliminate most of these dependencies:

▲ Settlement

is instant and not constrained by banking hours

▲ Cross-border transfers

happen at the protocol level, without correspondent banking relationships

▲ Programmability

allows agents to embed compliance logic, escrow conditions, and payment terms directly into the transaction itself

▲ Balances

can be held, moved, and deployed without requiring a bank account, custodian approval, or manual intervention at each step

For treasury management, this translates directly into operational capability. A treasury agent can monitor inflows across multiple chains, rebalance positions across currencies and assets in real time, deploy idle capital into yield-bearing instruments, and settle payments to counterparties globally.



The Infrastructure Stack: What Agents Actually Need

Identity and Wallet Architecture for Agents

In traditional financial systems, identity is managed through a combination of government-issued documents, biometric verification, and behavioral history. Every person has a legal identity linked to accounts that carry defined permissions and limits. That framework was built over decades of financial regulation around a single assumption: that there is a human on the other end.

Agents exist outside that framework entirely. A wallet is not just where an agent holds funds. It is the primary mechanism through which an agent's identity, authorization, and financial permissions are expressed onchain. When an agent signs a transaction, that signature is the only verifiable signal the receiving system has about who is acting, under what authority, and with what constraints. The wallet is essentially the agent's identity layer.

When evaluating wallet infrastructure for agentic applications, several considerations matter beyond the basics of secure key management:

▲ Authorization architecture

The wallet needs to enforce spending limits, transaction type restrictions, and time-based constraints at the infrastructure layer, independent of what the agent itself decides to do. Authorization that lives only in the agent's instructions is not authorization in any meaningful sense. It needs to be enforced beneath the agent, at a level it cannot override.

▲ Cryptographic agent identity

The wallet should support signed credentials that bind a specific agent instance to a specific principal and authorization scope. Without this, there is no reliable way to distinguish a legitimate agent transaction from one made by a compromised or spoofed agent.

▲ Auditability

Every action an agent takes with a wallet should produce a clear, human-readable record that can be reviewed after the fact. This is both a compliance requirement and the primary mechanism through which the humans and institutions deploying agents maintain meaningful oversight of what those agents are actually doing.

▲ Multi-tenancy and account separation

The infrastructure needs to enforce strict separation between accounts, ensure authorization scopes cannot bleed across tenants, and support permissioning that allows different agents to operate with different levels of access within the same platform.

Blockchain Networks as the Operational Layer for Agents

The relationship between an agent and a blockchain is different from the relationship between a human and one. A human initiates blockchain transactions when they choose to and monitors outcomes manually. An agent interacts continuously, reading onchain state to inform decisions, submitting transactions autonomously, and triggering subsequent actions based on the results.

Agents require blockchains with:

- ▲ **Predictable fee structures**
so that transaction costs can be modeled and budgeted within policy constraints
- ▲ **Fast and reliable finality**
so that an agent does not act on state that is subsequently reversed
- ▲ **Programmable transaction logic**
so that conditions and payment terms can be enforced at the protocol level
- ▲ **Robust RPC infrastructure**
so that an agent querying onchain state receives accurate and timely responses

Agents operating in financial services will increasingly need to move across multiple networks. Permissioned networks like Canton sit alongside public chains in many enterprise architectures, and an agent may need to operate across both. The infrastructure layer is what bridges those networks, ensuring the agent operates consistently and within policy regardless of where a transaction settles.

“


Autonomous agents require payment infrastructure that matches their operational model. With x402 on Stellar, payments are instant, low-cost, and embedded directly into the application flow. The result is an architecture where agents can discover services, transact, and settle value in real time.

Stellar provides the speed, reliability, and stablecoin infrastructure to make autonomous, machine-driven commerce viable at scale — on open, permissionless rails that any participant can build on.



Jose Fernandez da Ponte

President & Chief Growth Officer at Stellar Development Foundation



Data Availability, Audit Trails, and Travel Rule in an Agentic World

When an agent transacts at machine speed on behalf of multiple principals, the record-keeping requirements multiply rapidly. Financial regulators require that institutions maintain records of every transaction including the sender, recipient, amount, and time. For human-initiated transactions, this is manageable. But for agentic transactions at scale, it requires purpose-built infrastructure.

Underpinning all of this is data availability: the ability to access accurate, timely, and complete onchain data at the speed agentic systems require. An agent making financial decisions in real time depends on reliable access to transaction history, wallet balances, smart contract state, and network conditions. Data availability is also a regulatory concern. The ability to reconstruct the full state of an agent's activity at any point in time is a prerequisite for meaningful compliance and oversight.

That reconstruction depends on audit trails that go deeper than a standard transaction log. They need to capture the full authorization chain behind every action, tracing which agent initiated it, under which policy, and on behalf of which principal. This is not just a technical requirement. It is a legal one in many jurisdictions, particularly for financial services firms subject to MiFID II, SEC reporting, or FATF guidance.

The Travel Rule sits at the center of that. The FATF requirement that financial institutions share originator and beneficiary information for transfers above certain thresholds is now being applied to stablecoin transfers in a growing number of jurisdictions. At agentic scale, compliance cannot be handled manually. It needs to be embedded in the payment flow and applied automatically to every qualifying transaction at the moment it is initiated.

“

Agentic systems dramatically increase the speed and fragmentation of transactions, and that puts a premium on context, not just logs. From a Chainalysis lens, the real unlock is combining granular blockchain data with intelligence about the services, wallets, and networks behind each address.

When that intelligence feeds directly into compliance workflows, such as Travel Rule, institutions can automatically route, allow, or escalate transactions in real time, while helping regulators gain confidence that the underlying evidence trail is complete and verifiable.



Jordan Bregman

Strategic Growth Director at Chainalysis



How to Build for the **Agentic Economy**

Key Decision Criteria for Teams Building Today

The protocol race will sort itself out. The more important question for teams building today is what their infrastructure needs to support on day one so that scaling into more demanding use cases does not require starting over.

▲ Build on agent-capable wallet architecture from the start.

Programmatic wallet creation, scoped permissions, revocable agent authorization, and policy-engine controls are not features you can add later. Wallet infrastructure that was not deployed with agents in mind will require significant rework when those requirements arrive.

▲ Choose chains based on your users and compliance requirements.

Consumer-facing applications may find Base and Solana optimal for speed and cost. Teams anticipating institutional or cross-border use cases will need multi-chain support and permissioned network compatibility earlier than they expect.

▲ Design compliance from day one.

KYC, AML, sanctions screening, Travel Rule compliance, and audit trail requirements are not optional for financial applications. Look for solutions with built-in compliance and screening integrations.

▲ Treat security as infrastructure, not a feature.

Key custody, policy enforcement, and transaction screening need to operate at the infrastructure layer, beneath the agent, where they cannot be overridden. Don't sacrifice security and look for providers that utilize MPC to protect private keys and enable independent recovery.

▲ Stay protocol-agnostic where possible.

The agentic payment protocol landscape is not settled and more standards will follow x402 and ACP. Staying flexible now preserves the ability to move with the market rather than against it.

▲ Plan for dispute resolution before you need it.

Agentic payments are fast and final by design, and the frameworks for handling unauthorized transactions and disputed settlements are still immature. Teams that think through recourse mechanisms early will be better positioned than those that treat it as an edge case.

Where Dynamic Fits Into the Agentic World

The full agentic finance stack requires capability at every layer, and the weakest one determines what the whole system can safely do. Dynamic touches on every layer of the stack:



Wallet layer

Agents need wallets that can be created, permissioned, and managed programmatically, with scoped authorization for specific transaction types and amounts. Dynamic's embedded wallets support programmatic creation, policy enforcement, and extensive chain coverage out of the box.

Security layer

Dynamic's embedded wallets are protected by TSS-MPC key management, eliminating single points of failure while maintaining sub-second signing latency. Dynamic is also backed by Fireblocks, whose institutional MPC infrastructure and BAM protocol extend that same security model to the custody layer.

Developer experience layer

Dynamic gives teams the fastest path from idea to production for agentic applications. Prebuilt UI components, extensive SDK support, and out-of-the-box integrations with on-ramps, identity providers, and compliance partners mean teams spend less time on infrastructure plumbing and more time building the product.

Compliance layer

Every transaction is screened in real time for AML and sanctions risk, Travel Rule information is captured and transmitted where required, and policy engines enforce the rules that principals have defined for their agents.

Network layer

Dynamic supports more chains than any other embedded wallet provider, with native coverage across every EVM & SVM chain and embedded wallets on any network via raw signing.


Institutional deployment

For teams deploying agentic trading infrastructure, treasury automation, or enterprise payment workflows, Dynamic's Delegated Access, Server Wallets, TSS-MPC security, and integrations provide the operational and regulatory foundation that serious deployment requires.

Resources To Learn More About Agentic Payments


Learn More About Delegated Access

[Read](#) →



Breaking Down TSS-MPC Wallet Security

[Read](#) →



Fireblocks Blockchain Network Support

[Read](#) →



How x402 Enables Internet-Native Payments

[Read](#) →



Fireblocks Interoperable Network for Digital Assets

[See more](#) →

