

The Agentic Economy Needs Its Own Financial Rails

A thesis on infrastructure for autonomous AI agents

Simon CASTELAIN - April 10th, 2026

The Observation

AI agents are already in production. They book meetings, execute trades, query APIs, and consume compute resources autonomously. What they cannot yet do reliably is transact with vendors at scale without a human overseeing every payment step, because the infrastructure for it is missing.

Today, agent-to-vendor transactions default to prepayment or rate-limiting, because vendors have no way to extend credit to an entity with no legal identity, no credit history, and no enforcement recourse. AI agents are beginning to earn permission to initiate payments, move funds, manage subscriptions, and optimize purchasing decisions on behalf of consumers and businesses, but nobody has built the financial rails for it. The most capable agents end up throttled by the absence of financial primitives designed for machines, more than by compute or intelligence.

The Thesis

Autonomous AI agents will become the first non-human economic participants at scale. DeFi built financial rails for humans. The next cycle will build them for machines: credit, reputation, trustless enforcement, and working capital.

The scale of what is being built makes the problem inevitable. Grand View Research estimates the global AI agents market at \$7.63 billion in 2025 and projects \$182.97 billion by 2033, a CAGR of 49.6% from 2026 to 2033.¹ IDC projects that AI spending, driven mainly by agentic AI, will reach \$1.3 trillion in 2029, and that the number and complexity of third-party and custom-built AI agents used by enterprises will grow tenfold over the next five years.² At that scale, the lack of native financial infrastructure for agents becomes a structural bottleneck.

Why Now

Three conditions are converging.

Agents are becoming autonomous enough to transact at scale

In McKinsey's 2025 State of AI survey, 62% of respondents said their organizations were at least experimenting with AI agents.³ In 18 months, the gap between an AI assistant that helps humans and an AI agent that acts on its own has closed materially. More than 17,000 AI agents have been created on Virtuals Protocol alone, and the protocol's cumulative revenue had passed \$39.5 million by January 2026.⁴ These agents run in production and generate real economic activity.

The payment layer is visibly broken for machines

Stripe product manager Jeff Weinstein has said that the current financial system is tuned for humans and incompatible with what agents need: microtransactions, 24/7 global rails, and low-latency finality guarantees.⁵ The workarounds (prepayment, rate limiting, human co-signing) limit what agents can do and how fast the ecosystem scales.

The infrastructure layer is being defined right now

In February 2026, Stripe released a preview of machine payments built on x402, Coinbase's open standard, which lets developers charge AI agents directly in USDC on Base.⁶ In March 2026, Virtuals Protocol and the Ethereum Foundation's dAI team proposed ERC-8183, a trustless standard that lets agents run hire-deliver-settle workflows onchain.⁷ Davide Crapis, Head of AI at the Ethereum Foundation, called ERC-8183 one of the missing components of the open agent economy. The stack is being assembled this year, and the funds that build conviction here in 2026 will get first-mover access to the category.

The Gap These Solutions Leave Open

Each standard in the emerging stack solves a distinct problem:

Layer	Current Standard (2026)	What it solves
Identity	ERC-8004	Who is the agent?
Payment	x402 / USDC	How is money moved per request?
Commerce	ERC-8183	How is discrete work structured?
Credit / Enforcement	AgentEscrow	How is the vendor protected for continuous consumption?

ERC-8183's Job primitive works as a decentralized escrow for structured tasks. An agent hires another agent and locks the funds upfront, the hired agent submits a deliverable, and an evaluator confirms it. This works well for discrete, scoped work. It does not work for continuous, high-frequency resource consumption, such as an agent querying an inference API hundreds of times per minute, consuming compute from a vendor it has never contracted with before, at a volume neither party agreed to in advance. In that case, upfront escrow per transaction is inefficient and incompatible with how inference workloads run.

This is the credit problem. Identity alone no longer separates legitimate agents from malicious ones. Early-stage agentic AI can initiate actions autonomously without clear guardrails, which lets attackers orchestrate large volumes of payment attempts and automate fraud at scale. Vendors currently have two options: block unknown agents entirely, or extend access and absorb the default risk. Neither is acceptable at scale. What vendors need is a trustless credit enforcement layer.

On the vendor side, an inference provider would accept an AgentEscrow bond instead of a credit card because a card default takes weeks of chargeback bureaucracy to resolve, while

an AgentEscrow slash is programmatic and instant. The cost of collection shifts from the vendor's accounting department to the smart contract. For high-frequency, low-margin API transactions, that difference separates a recoverable default from an unrecoverable one at scale.

What I Built

AgentEscrow is a bond-based credit enforcement protocol for AI agent-to-vendor transactions, deployed on Arbitrum Sepolia. The mechanism: an agent posts a bond once before accessing vendor resources. If the agent fails to settle before a deadline, the vendor calls `triggerSlash()` on the Heartbeat contract, and the bond transfers atomically to the vendor's wallet. Two conditions must be met: the deadline has passed, and the credit line is still in ACTIVE status, meaning `confirmSettlement()` was never called. No oracle, no keeper, no external dependency in V1. 131/131 tests passing.

A bond is also more capital-efficient than prepayment. A single bond acts as collateralized reputation across several vendors at once, whereas prepayment locks capital into single-vendor silos. An agent posting a \$200 bond can access credit lines across ten APIs concurrently, against \$1,000 of locked prepayment for the same access. That is a 5x capital efficiency gain, and it compounds as the number of vendors grows.

V1 has a known limitation: settlement must occur onchain via `confirmSettlement()` to be recognized by the protocol. Off-chain payment does not protect the agent from a slash. This is documented and scoped for V2 with a `disputeSlash()` mechanism.

UMIP (Unified Margin Infrastructure Protocol) is a hub-and-spoke USDC vault for autonomous trading agents on Arbitrum Sepolia. An agent deposits USDC once, and the vault routes collateral to adapters for GMX, gTrade, and Vertex instead of locking it on one perpetual DEX. The GMX adapter is deployed on Arbitrum Sepolia; the gTrade and Vertex adapters are tested against mocks. The code is in three public GitHub repos, and a sandbox and a dashboard are live.

I designed both protocols around the same question: What does financial infrastructure look like when the counterparty is a machine?

The Signal I'm Watching

My conviction breaks if, within 18 months, inference API providers report no material enforcement or non-payment issues at scale and agent transaction volumes remain marginal. Two signals would strengthen it: protocols focused on agent credit and enforcement raising meaningful venture funding before the end of 2026, and the ERC-8183 ecosystem surfacing continuous-consumption edge cases that the Job primitive cannot address.

Academic research has started on this layer. In March 2026, a team from Peking University and Shanghai Jiao Tong University published A402, which proposes channel-based protocols to address x402's atomicity limitations.⁸ This confirms that the next infrastructure problems sit where AgentEscrow operates.

Why This Matters for Investors

In every new computing paradigm, the infrastructure layer has captured disproportionate value. TCP/IP, AWS, and Ethereum each became the base that everything else was built on,

and the agent economy will be no different. Amazon expects capital expenditure of about \$200 billion in 2026, more than 50% above 2025 levels,⁹ and Alphabet guided its 2026 capex to between \$175 billion and \$185 billion, about double the \$91.4 billion it spent in 2025.¹⁰ That compute is what agents will run on when they transact at scale.

The payment and commerce standards for that world are being written right now. The credit and enforcement layer sits one level above, and nobody has claimed it yet.

Code: github.com/SimonCastelain/AgentEscrow and github.com/SimonCastelain/umip.
Contact: simon.castelain@essec.edu

Sources

1. Grand View Research, [AI Agents Market Size To Reach \\$182.97 Billion By 2033](#), press release.
2. IDC, [Agentic AI to Dominate IT Budget Expansion Over Next Five Years, Exceeding 26% of Worldwide IT Spending, and \\$1.3 Trillion in 2029](#), press release, August 26, 2025.
3. McKinsey & Company, [The state of AI in 2025: Agents, innovation, and transformation](#), November 2025.
4. Ventureburn, [Virtuals Protocol explained: The AI agent launchpad taking crypto by storm](#), January 30, 2026.
5. Forklog, [report on Stripe's machine payments preview and Jeff Weinstein's comments](#), February 2026.
6. The Block, [Stripe adds x402 integration for USDC agent payments on Base](#), February 10, 2026.
7. [ERC-8183: Agentic Commerce](#), Ethereum Improvement Proposals.
8. [A402: Bridging Web 3.0 Payments and Web 2.0 Services with Atomic Service Channels](#), arXiv:2603.01179, March 1, 2026.
9. SiliconANGLE, [report on Amazon's Q4 2025 results and 2026 capital expenditure forecast](#), February 5, 2026.
10. Yahoo Finance, [Alphabet says capital spending in 2026 could double, cloud business booms](#), February 4, 2026.