Who Pays the Agent?

The Race for Frictionless Machine Payments

Keyrock coinbase TEMPO Virtuals

About Keyrock

Keyrock

Founded in Brussels in 2017, Keyrock is a global crypto investment group leading in market making, asset management, OTC, and options trading for digital assets. Providing liquidity to over 85 centralized and decentralized venues worldwide, their 220-strong team operates across 37 countries, with entities in Belgium, the UK, Switzerland, France, and the U.S. Keyrock's commitment to the industry is practical, not theoretical. They offer in-depth industry insights, co-create DeFi ecosystems, and actively support Web3 startups. With Keyrock, the future of digital assets is not just envisioned; it's actively being built.

About Coinbase

coinbase

Founded in San Francisco in 2012, Coinbase is a publicly traded crypto platform (NASDAQ: COIN) serving over 110 million verified users across 100+ countries. It operates Base, one of the most active Layer 2 networks, and created x402, the first open protocol for machine-to-machine stablecoin payments. Its Agentic Wallets power tens of thousands of autonomous agents.

About Tempo

TEMPO

Tempo is a payments-focused Layer 1 blockchain incubated by Paradigm and Stripe, launched in March 2026 with a $500 million Series A at a $5 billion valuation. It's built for stablecoin payments at internet scale, with sub-second finality, and predictable fees under $0.001. Alongside its mainnet, Tempo co-authored the Machine Payments Protocol (MPP) with Stripe, an open standard for how AI agents and software services request, authorise, and settle payments.

Table of Contents

1 **Market Overview** 2

2 **Payment Architectures** 4

- 2.1 x402: HTTP-Native Stablecoin Settlement 5
- 2.2 MPP: Payment-Method Agnostic Protocol 6
- 2.3 Google AP2: Delegated Authorisation Framework 8
- 2.4 Visa Intelligent Commerce: Tokenised Card Credentials 9
- 2.5 Stack Architecture and Vertical Integration 10

3 **Agent Infrastructure and Distribution** 13

- 3.1 Agent Frameworks and Financial Capabilities 13
- 3.2 Wallet Infrastructure and M&A Activity 14
- 3.3 Headless Merchants and Service Discovery 15
- 3.4 Settlement Composition and Stablecoin Dominance 16
- 3.5 Agent Autonomy and Governance Models 18

4 **Transaction Economics and Liquidity** 19

- 4.1 Card Fee Structures and the Micropayment Floor 19
- 4.2 Stablecoin Velocity by Chain 21
- 4.3 Concentration Risk and Vertical Integration 23
- 4.4 Market Microstructure at Micro Scale 24
- 4.5 Non-Human Transaction Activity and Adoption 25
- 4.6 Total Addressable Market and Growth Assumptions 27

5 **Regulatory Landscape** 28

- 5.1 Legislative Gaps in Current Frameworks 29
- 5.2 Liability and Consumer Protection 29
- 5.3 Sanctions Compliance and Enforcement Risk 30
- 5.4 Agent Identity and Authentication Standards 30

6 **Closing Thoughts** 32

Executive Summary

In the past twelve months, machine-to-machine payments have gone from a concept to a developed ecosystem. Four competing payment architectures have launched, backed by Coinbase, Stripe, Google, Visa, and American Express. Agents have settled over **$73 million** across **176 million** transactions, and incumbents have deployed more than **$8 billion** in acquisitions to secure their position in what is emerging as an entirely new payment stack.

This report analyses how that stack is assembling, whether the economics work, and what stands in the way.

Key Findings:

1. Card rails can't serve the machine economy: **76%** of agent transactions fall below the **$0.30** card-fee floor. Layer 2 stablecoin settlement costs **$0.0001**. For agents, crypto rails are a mathematical necessity.

2. The protocols aren't competing, they're assembling into a stack: x402, MPP, AP2, and Visa are building different layers of the same architecture.

3. **98.6%** of agent payments settle in a single stablecoin: USDC won by default, but this concentration poses a systemic risk nobody is publicly addressing.

4. The machine economy is already here, it just isn't doing commerce yet: AI agents account for **75%** of Safe transactions on Gnosis Chain on peak days. The shift from extractive bot activity to productive agent commerce is underway.

5. Regulation is the binding constraint: MiCA, the GENIUS Act, and the EU AI Act all reach enforcement within weeks of each other in mid-2026. None address autonomous machine-to-machine transactions.

Section 1: Market Overview

In May 2025, the first machine-to-machine payment settled on a public blockchain. An AI agent requested a resource, a server quoted a price, the agent signed a stablecoin authorisation, and the resource was delivered. At no point did a human touch the transaction, nor was an account created, nor a subscription signed. The entire exchange completed in under **two** seconds.

Twelve months later, autonomous agents have settled over **$73 million** across approximately **176 million** transactions, four competing payment architectures have launched with backing from Coinbase, Stripe, Google, Visa, and American Express, and incumbents have deployed over **$8 billion** in acquisitions to secure their position in what is emerging as a new payment stack.

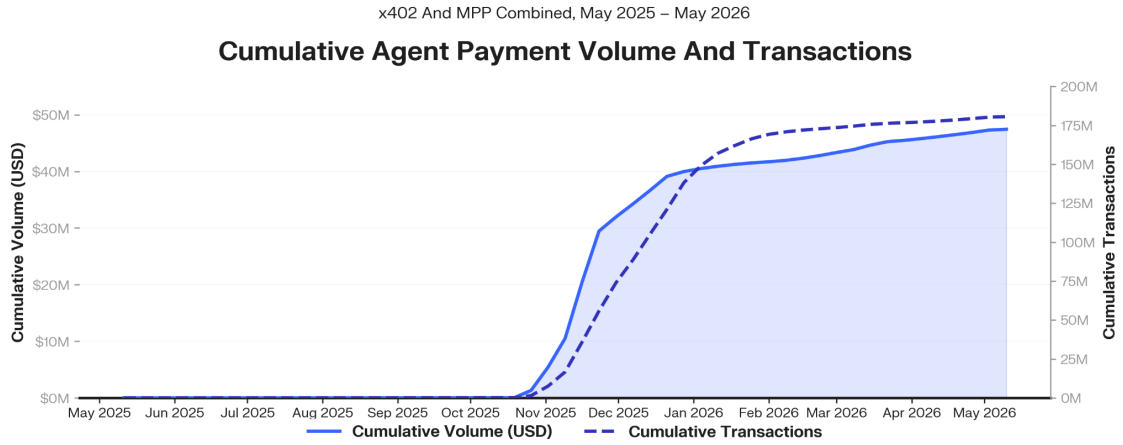

Source: Keyrock Research, The Agent Payments Stack

These are small numbers by the standards of global commerce, with Visa processing **$14.5 trillion** annually. The entire agent payments ecosystem, annualised, represents roughly **0.0003%** of that figure, six orders of magnitude smaller.

However, irrelevance measured in dollars can obscure relevance measured in structure. In the twelve months between May 2025 and April 2026, the architectural foundations of a machine-native financial system were laid, debated, funded, and partially consolidated. The most significant development is therefore the velocity of infrastructure formation, and the fact that it is converging on a layered architecture that looks, to us, like the early internet protocol stack.

In parallel to the internet, which didn't become commercially important when the first website was launched, we believe agentic payments will become commercially important when trust discovery and settlement become trivial. This happened for the internet with HTTP, TCP/IP, DNS, and SSL, and we think we are watching an analogous agentic commerce assembly in real time. Settlement layers, wallet infrastructure, payment protocols, routing mechanisms, and governance frameworks are being built concurrently by different actors, each optimising for a different layer, each assuming the others will eventually exist.

What follows is an analysis in four parts, starting with Section 2, which maps the four payment architectures and the stack they are assembling into. Section 3 traces the pipeline from agent framework to wallet to marketplace. Section 4 examines whether the liquidity economics actually work, and Section 5 confronts the regulatory vacuum that may determine whether any of this matters at all.

Section 2: Payment Architectures

In **September 2024**, if you wanted an AI agent to pay for something, you had essentially one, very unsafe, option. This option being to hardcode a credit card into your application and pray the fraud detection system didn't flag a bot making 400 API calls at 3am. Twelve months later, there are four competing architectures for machine payments, backed by some of the largest companies in technology.

Coinbase built x402, a crypto-native protocol that turns stablecoin wallets into universal API keys. Stripe and Tempo launched MPP, a payment-method-agnostic standard that handles cards, crypto, and Lightning through a single HTTP flow. Google assembled AP2, an authorisation layer that lets users safely delegate spending power to agents via cryptographic mandates. And Visa extended its existing card rails to provision AI-ready tokenised credentials. Between them, they've processed **hundreds of millions** of transactions and attracted **$8 billion** in M&A activity. The speed of this buildout is remarkable, and the fact that nobody agrees on how it should work is entirely predictable given how early this technology is.

What's less obvious, and what most coverage misses, is that these four approaches aren't purely competing with each other. There is genuine overlap at the protocol layer, but the more important dynamic is that they are assembling into a stack. The right question isn't "which protocol wins?" It's "which companies capture the most layers, and therefore attract the most value?"

Four Ways to Pay a Machine

	x402	MPP	Google AP2	Visa Intelligent Commerce
Backers	Coinbase + 21 partners	Stripe / Paradigm / Tempo	Google + 60 partners	Visa
Core Problem	Agent-to-service payments	Universal machine payments	Authorisation & trust	Extend card rails to agents
Rails	Onchain (Base / USDC)	Agnostic (crypto+fiat+cards)	Payment-agnostic (cards now)	Traditional VisaNet
Micropayments	Yes (low Base fees)	First-class (session model)	Not a focus	No (card economics)
Merchant Adoption	~3,900 merchants	100+ at launch	60+ partners	Existing Visa acceptance
Agent Autonomy	High (permissionless)	Medium (pluggable)	Medium (mandate-gated)	Low (issuer-gated)
Launched	May 2025	March 2026	Sept 2025 (AP2) Jan 2026 (UCP)	Pilots 2025

Source: Keyrock Research

Keyrock

x402: HTTP-Native Stablecoin Settlement

Coinbase drew first in **May 2025** when they open-sourced x402, a protocol that repurposes the long-dormant HTTP 402 "Payment Required" status code for machine-to-machine stablecoin payments. The design is elegant in its simplicity. All that's required is for an agent to request a resource, which is followed by the server responding with a price. The agent then signs a USDC payment authorisation, and the server delivers. The whole process works with no accounts, no API keys, and no subscriptions.

The technical implementation relies on EIP-3009's 'transferWithAuthorization', which allows a "facilitator" to pull USDC from the agent's wallet without the agent needing to submit an onchain transaction or pay gas. Settlement happens by default on Base, Coinbase's Layer 2, where fees are low enough to make sub-dollar payments economically viable and other networks, like Solana, are also supported. This obviously matters, though it's worth noting that our data shows that the average x402 transaction size has stabilised around **$0.48** over the past month, though it fluctuated significantly in x402's early months when sample sizes were small.

Since launch, x402 has processed approximately **176 million** transactions totalling roughly **$73.2 million** in volume, with around **3,900** merchants now accepting payments through the standard, including AWS, Alchemy, and Messari. As Erik Reppel, x402's creator and Head of Engineering at Coinbase Developer Platform, frames it, "Your wallet becomes the universal API key that lets you access any x402-enabled service."

These numbers, of course, deserve scrutiny. Onchain analysis by Allium Labs suggests that a meaningful portion of observed x402 transactions reflect artificial or gamified activity, and that the headline volume figure may overstate genuine commercial throughput by as much as half. This does not invalidate x402's design. Clearly the protocol works, but it does mean the ecosystem is earlier than the headline numbers suggest, if this analysis is to be believed. For us, the real signal is not cumulative volume but the sustained **$11-15 million** weekly run rate in Q1 2026, which coincides with genuinely useful integrations like Nansen and Alchemy rather than airdrop farming. We cover this trajectory in detail in Section 3. We conducted an interview with Erik Reppel, creator of x402, who noted that non-Base activity is more significant than commonly assumed, with Solana carrying meaningful volume and new chains such as Aptos, Sui, and Stellar being added regularly. The x402 Foundation is also actively working on fiat payment support with multiple parties, including Stripe. This development, if realised, would substantially undermine the 'crypto-only' criticism.

The organic activity that does exist is predominantly agents paying for API access on a per-request basis. An agent querying Nansen for onchain analytics pays **$0.01** per request in USDC on Base. An agent pulling market data from Alchemy or retrieving research from Messari follows the same model: no subscription, no account, just a micropayment per call. Coinbase's Agentic.Market, launched in April 2026, now lists over **365** services available on this basis, effectively an app store where agents browse, pay, and consume without human intervention.

The economics are simple for this activity, in that a developer subscribing to Nansen's API pays **$500** per month regardless of usage. An agent paying **$0.01** per query via x402 would need to make **50,000** requests before the subscription becomes cheaper. For the long tail of low-frequency, high-variety agent tasks, where an agent might query Nansen once, Dune twice, and an LLM endpoint a dozen times in a single workflow, per-request pricing eliminates the dead weight of unused subscriptions. There's also an access angle: agents operating in jurisdictions without card infrastructure can transact using stablecoins alone, removing a barrier that subscription models can't.

Ultimately, x402 is currently positioned in a stack that makes Base the default settlement layer for machine commerce, USDC the default denomination, and Coinbase's wallet infrastructure the default on-ramp. It's an open protocol closely tied to a full-stack play.

Reppel pushes back on this characterisation, stating, "Standards are really a value creation game, not a value capture game." He draws an analogy to the early web, where AOL and the World Wide Web were both trying to solve the same problem of making information accessible to consumers, but AOL was a walled garden while the web was a set of open standards that anyone could build on. "Five years from now, I hope no one talks about x402 because normal people don't talk about HTTP." It is a compelling aspiration."

MPP: Payment-Method Agnostic Protocol

If x402 defined the settlement layer, MPP is an attempt to define the protocol layer above it. Rather than prescribing where money moves, MPP prescribes how the payment conversation happens, and is agnostic about which rail settles underneath. Launched on the March 18 2026 alongside Tempo's mainnet, the Machine Payments Protocol (MPP) was co-developed by Stripe and Tempo, a payments-focused Layer 1 blockchain incubated by Stripe and Paradigm that raised a $500 million Series A at a $5 billion valuation.

MPP uses the same HTTP 402 signalling mechanism as x402, but diverges sharply in philosophy. Where x402 is crypto-native, MPP is payment-method agnostic by design. It supports stablecoins, credit cards, Bitcoin’s Lightning Network, and bank transfers through a unified Challenge, Credential and Receipt flow. The agent requests a resource, receives a price and a menu of accepted payment methods, pays via whichever rail it prefers, and receives the resource plus a cryptographic receipt. As Brendan Ryan, software engineer at Tempo, told Keyrock in an interview, "MPP is the checkout form for machines." Just as consumer e-commerce standardised around a universal checkout interface with dozens of payment methods plugged in behind it, MPP aims to be the machine-native equivalent. This method can be thought of as the Swiss Army Knife response to x402: one protocol, any rails, any currency.

MPP also has a killer feature known as sessions. For high-frequency micropayments, an agent can pre-authorise a spending limit and then stream granular payments continuously within that session, without individual onchain transactions per request. This enables use cases such as pay-per-token billing for LLM inference, per-query pricing for data services, and the kind of sub-cent granularity that card economics simply cannot support.

The design philosophy is composable by intent, whereby rather than maintaining integrations centrally, Tempo and Stripe maintain a core SDK and let partners hang their own payment methods off it. This enables rapid onboarding, as shown with Visa, according to Ryan, "We started talking to Visa on a Thursday and they had a working integration by Sunday."

When a company the size of Visa can go from first conversation to working code in four days, the protocol's minimalism is doing real work.

MPP launched with over **100** integrated services, including OpenAI, Anthropic, Google, Dune Analytics, Alchemy, and Cloudflare. In its first week, **913** agents executed over **34,000** transactions across the directory at prices ranging from **$0.003** to **$35** per request. The spec has been proposed to the IETF, though Ryan frames the submission pragmatically, "The IETF process heavily favours working abstractions and running code," he says, noting that they chose it over heavier foundation-style governance precisely because it rewards adoption over committee politics. Reppel offers a different perspective from the x402 side, arguing that IETF submission should be the last step, not the first. "Once something is in the IETF, it's kind of done," he says. "It's really hard to evolve IETF standards. And I don't think the story is done for what agentic commerce is going to look like." His preferred sequence is to start with a core team, expand to a foundation with industry partners, and only then seek IETF enshrinement once there is sufficient consensus. It is a philosophically different approach to standardisation, and one that implicitly bets on the landscape still being too early to codify.

Perhaps most telling is that both Visa and Mastercard have published payment method specifications for MPP. When the card networks are writing plugins for your protocol rather than building their own, something structural has shifted.

It is worth noting that MPP's first-week volume of approximately **$3,730** across **34,000** transactions implies an average transaction value of roughly **eleven** cents. This is consistent with developer experimentation rather than commercial adoption, a distinction that matters when evaluating the protocol's readiness for production workloads.

Early volume data reinforces the rail differentiation, with Ryan estimating that the majority of MPP volume remains stablecoin-based, characterised by lower value and higher frequency transactions, with stablecoin-settled requests priced as low as one-tenth of a cent. Card-based transactions via Stripe are higher in value, typically **$5-10** where card economics begin to make sense, but remain infrequent. "Focusing on whether cards or stablecoins will be the dominant rail is unproductive," Ryan argues. "It'll be many things." It is a view that neatly supports the stack thesis by implying that the protocol layer should be agnostic precisely because the settlement layer won't converge.

It is worth noting that the most high-profile attempt at consumer-facing agentic checkout has already failed. OpenAI shelved ChatGPT Instant Checkout in March 2026, approximately five months after launch, after only around **30** Shopify merchants actively used it. The product lacked sales tax collection, fraud prevention, and multi-item cart support. Purchases now occur inside connected apps rather than natively in ChatGPT. The lesson is that the consumer checkout model, where an agent browses and buys like a human, is not viable in this form, and that the machine-native model, where agents transact via protocol endpoints without a visual interface, is the preferred market approach.

Google AP2: Delegated Authorisation Framework

As is expected, Google's entry to the game operates at a different altitude entirely. Announced in **September 2025** with over **60** partners, including Mastercard, PayPal, American Express, Coinbase, Shopify, and Walmart, the Agent Payments Protocol doesn't solve "how does an agent pay?" It solves "how does a user safely delegate spending authority to an agent?"

The mechanism is built around Verifiable Digital Credentials, which Google calls "Mandates". Essentially, they're cryptographically signed authorisation records that travel with the transaction, a standard only AP2 allows for at the moment. AP2 can be used in one of two modes, the first is real-time mode, where the user is present and approves a 'Cart Mandate', which is a signed record of items, price, and shipping details. In the second, delegated mode, the user signs an 'Intent Mandate' upfront, specifying price limits, timing, and conditions. The agent then operates autonomously within those bounds, generating Cart Mandates as conditions are met. Both scenarios produce a separate 'Payment Mandate' that signals to the financial network that an AI agent, not a human, initiated the transaction.

This is a fundamentally different problem space to the two solutions we've explored thus far. x402 and MPP assume the agent already has authority to spend, whereas AP2 creates the trust infrastructure for that authority to be granted, verified, and audited. In **January 2026**, Google expanded the architecture with the Universal Commerce Protocol, a complementary open standard for agent-driven product discovery, developed with Shopify, Etsy, Wayfair, Target, and Walmart.

The fact that Coinbase is listed as both an AP2 partner and the creator of x402 is telling. These protocols are complementary layers, not competitors. The general flow is that AP2 handles authorisation, and x402 or MPP handles settlement.

We can also see Google's own actions confirming this, with the Solana Foundation launching pay.sh with Google Cloud in May 2026. Pay.sh is a gateway that lets autonomous agents discover, access and pay for APIs using stablecoins on Solana, built on x402 and MPP. Google Cloud is exposing Gemini inference, BigQuery, Cloud Run and BigTable through the service. As Rich Widmann, Head of Strategy for Web3 at Google Cloud, put it, "Agentic payments are one of the most important frontiers in the agentic stack." In practice, Google is treating AP2 as the authorisation layer and x402 and MPP as the settlement layer, and is actively building across both.

Visa Intelligent Commerce: Tokenised Card Credentials

Visa's approach is the most pragmatic and, depending on your perspective, either the most boring or the most dangerous. Rather than building new rails, Visa is extending its existing ones. Their 'Intelligent Commerce' initiative provisions AI-ready tokenised credentials, essentially virtual card numbers that can be assigned to an AI agent with programmable spend controls, such as merchant category limits, amount caps, and time windows. The agent presents these credentials at checkout via standard card-not-present flows. This means that merchants don't need to change anything about the way they sell their products or services, they see a normal Visa transaction.

The advantage of this approach is immediate distribution within Visa's network. Visa is accepted at over **100 million** merchant locations worldwide. Therefore, there's no new ecosystem required, no merchant integration, and no crypto literacy. The disadvantage is equally clear, in that card economics were not designed for machines that execute large volumes of small transactions. A fixed processing fee of roughly **30 cents** per transaction makes sub-dollar payments uneconomical, for example an agent paying **three** cents for a weather API call cannot route through Visa. Visa has since prototyped Visa CLI, a command-line tool currently in limited beta that allows agents to pay with tokenised Visa credentials using Touch ID authentication for each purchase, effectively bringing card rails into the developer terminal where vibe coding happens.

This is precisely why Visa published a payment method specification for MPP, positioning themselves as a payment option within the new protocol stack rather than an alternative to it. As Olivia Chow, director of Zero Knowledge Consulting, said, "If they get this right, it doesn't cannibalise what they're doing at all. If anything, it increases their power and strengthens their grip on the market, because now they're not just payment processors, they're also on the discovery side."

American Express entered the race on the 14th April 2026 with a fundamentally different proposition. Its Agentic Commerce Experiences (ACE) developer kit offers agent registration, intent verification, tokenised payment credentials, and cart-level risk checks, but the headline feature is Amex Agent Purchase Protection. This means that if a verified agent makes an erroneous purchase, Amex covers it. As Simon Taylor observed, this is a liability position, not a protocol specification. Amex's closed-loop network, where it simultaneously acts as issuer, network, and acquirer, eliminates the four-party coordination problem that Visa and Mastercard face. Luke Gebb, Amex's EVP of Global Innovation, offered a useful reality check alongside the announcement, stating "To date there have probably been as many press releases as transactions, but no doubt it will happen."

Stack Architecture and Vertical Integration

The conventional framing of this landscape as a 'standards race' misses what is actually happening. These protocols are assembling into a layered architecture, analogous to the internet protocol stack, where different players are racing to own different layers.

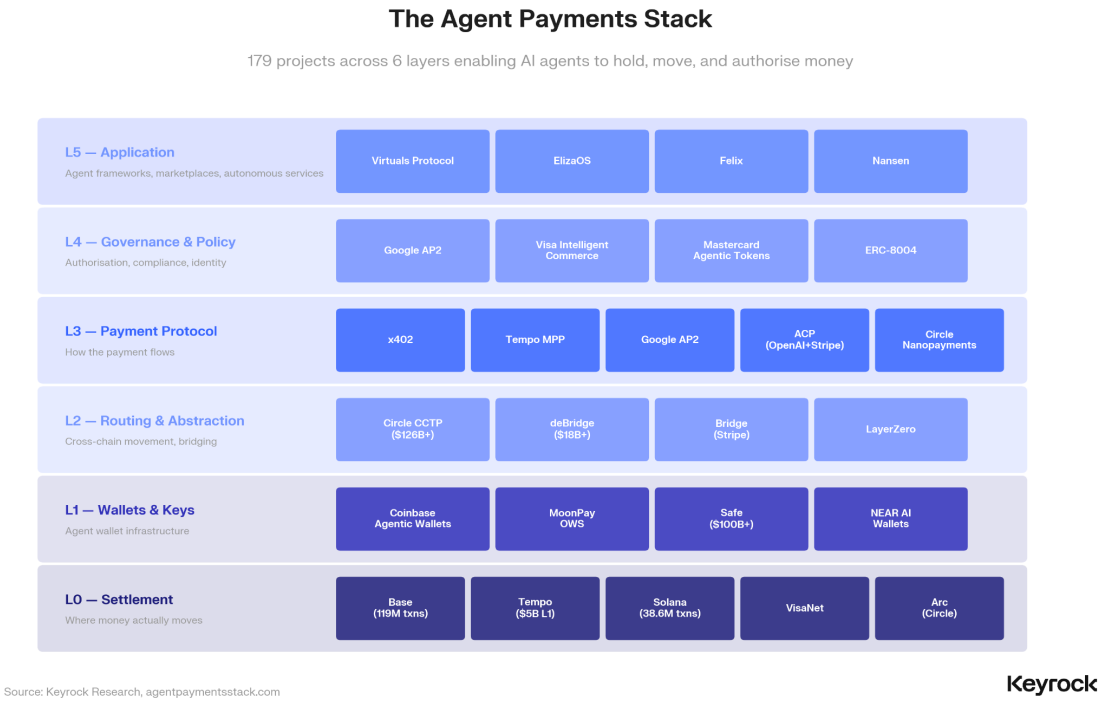

At the base, settlement infrastructure determines where money actually moves, these are the classic settlement layers we know and love, such as Base for x402, Tempo for MPP, VisaNet for card payments, and Circle's Arc for nanopayments. Above that, wallet and key management determines how agents hold and control funds, for example Coinbase MCP's, MoonPay's Open Wallet Standard, and Safe's multi-signature infrastructure. The routing layer handles cross-chain movement, with Circle's CCTP having processed over **$126 billion** in volume. Payment protocols, the layer that attracts the most attention, define how the payment actually flows. And at the top, the governance layer handles authorisation, compliance, and identity.

The strategic implication is that breadth of stack coverage may matter more than depth at any single layer. Data from The Agents Payment Stack maps **179** projects across these **six** layers and reveals a striking pattern in vertical integration.

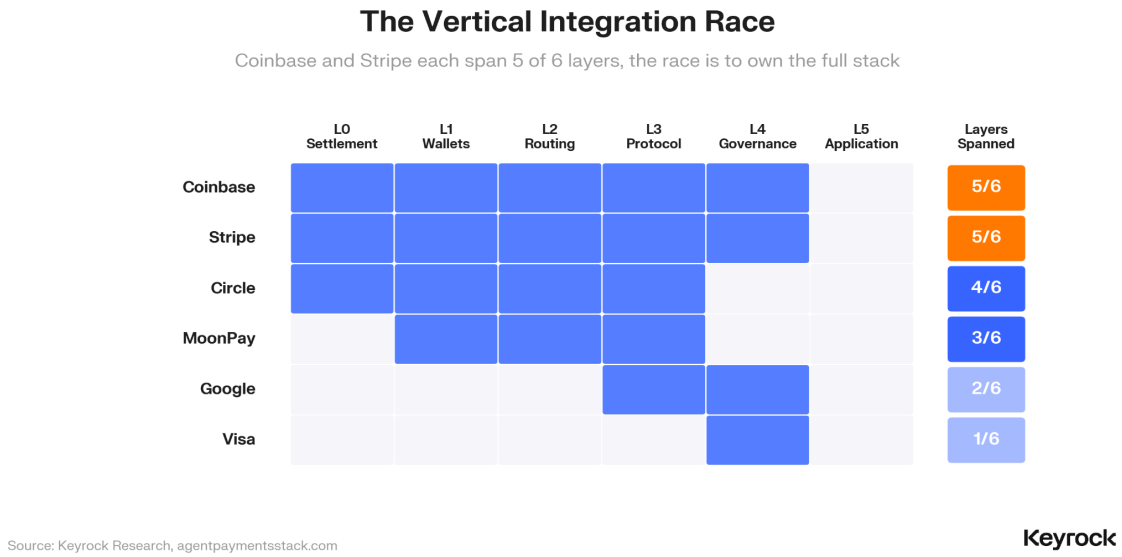

Coinbase and Stripe each span **five** of **six** layers, thus being the two companies with the biggest breadth. Coinbase controls settlement (Base), wallets (Coinbase MPC's), routing (internal infrastructure), the payment protocol (x402), and governance (as an AP2 partner). Stripe mirrors this through Tempo (settlement), Privy (wallets, acquired), Bridge (routing, acquired for **$1.1 billion**), MPP (protocol), and its own compliance infrastructure. Circle covers **four** layers, while Google and Visa, despite their scale, currently span just **two** and **one** respectively.

The uncomfortable implication is that the 'open protocol' narrative may be a transitional phase. Both Coinbase and Stripe are building vertically integrated stacks that, if successful, will allow them to capture value at every layer of an agent transaction. The protocols are open, but the infrastructure increasingly is not, and this mirrors the early internet, where open protocols like HTTP and SMTP enabled enormous value creation, but that value was ultimately captured by vertically integrated platforms rather than the protocol designers. Reppel contests this framing directly. Instead, in his view, the current phase is necessarily about value creation, not capture, and the correct analogy is the pre-commercial internet where open standards had to be established before anyone could build businesses on top of them. The counterargument is that Coinbase is simultaneously leading in building on the standard and the infrastructure that benefits most from its adoption, a dual role that the architects of HTTP never occupied.

The vertical integration race is heating up too. We have seen that over the past twelve months, incumbents have deployed over **$8 billion** in acquisitions to fill gaps in their stack coverage.

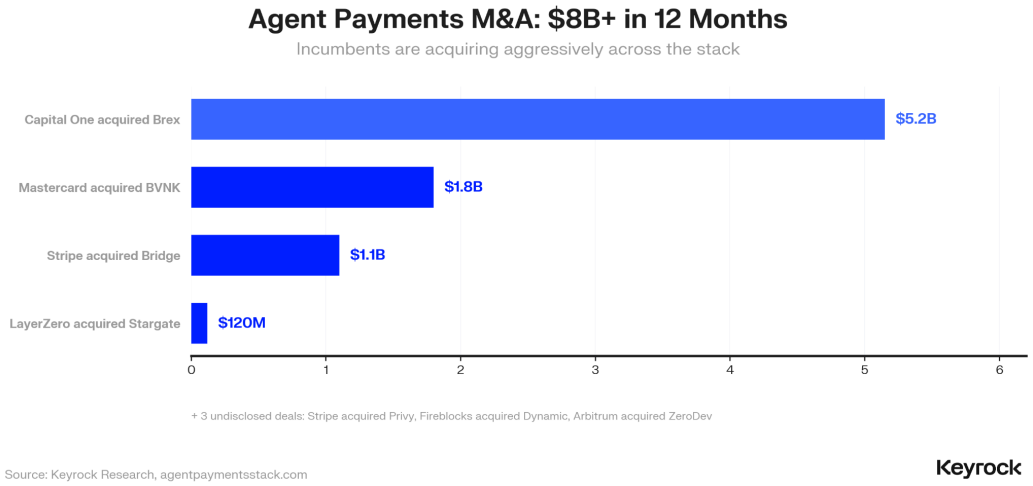

Capital One's **$5.15 billion** acquisition of Brex, whose programmable card infrastructure provides a foundation for agent spend controls, Mastercard's **$1.8 billion** purchase of BVNK, and Stripe's **$1.1 billion** acquisition of Bridge are infrastructure consolidation plays by companies that see machine payments as an inevitable expansion of their core business.

Section 3: Agent Infrastructure and Distribution

Section 2 mapped the protocols, and outlined the plumbing, but plumbing without water is just an expensive art installation. The more urgent question is who is actually building the agents that will use these rails, and how do they get from autonomous software process to an entity that holds money, discovers services, and pays for things without asking permission?

The answer is emerging as a **three-stage pipeline:**

1. Frameworks that give agents financial capabilities
2. Wallets that let them hold and control funds
3. Marketplaces that let them discover and transact with services

We believe that each stage is consolidating faster than most people realise.

Agent Frameworks and Financial Capabilities

The first generation of AI agents were essentially chatbots with delusions of grandeur. This is true in the sense that they could generate text, summarise documents, and occasionally hallucinate a convincing stock recommendation. They could not, however, hold a wallet, sign a transaction, or pay for anything. That changed in late 2024, and the speed of the transition since has been remarkable.

Virtuals Protocol, which launched as a tokenised AI agent platform on Base, now hosts over **18,000** agents with a combined 'agentic GDP' exceeding **$470 million**. That figure should be treated with appropriate caution, given it is denominated in VIRTUAL tokens and includes all agent-to-agent economic activity on the platform, making it sensitive to token price fluctuations and circular flows. For us, with this particular case the more grounded signal is the raw transaction count. As Virtuals told Keyrock, "The top categories are trading, market intelligence, and media generation. The vast majority of the jobs are finance-related, and a portion of these agents relied on other agents to obtain market alpha and trading signals as an input to their trading decisions."

Ethy AI, one of the most active agents on the platform, has processed over **2 million** transactions autonomously. The Virtuals Revenue Network, launched in February 2026, allows agents to independently request services from other agents, negotiate terms, execute work, and settle payments, all without human intervention.

ElizaOS, the open-source framework originally incubated by ai16z, has taken a different approach. Rather than building a closed ecosystem, it has become a composable toolkit with **63** repositories and a sprawling plugin architecture. Agents built on ElizaOS can swap tokens via QuickSwap, manage wallets through Coinbase's Agentic Wallet plugin, bridge assets across chains, and execute DeFi strategies. The framework's 'Spartan' agent handles multi-chain trading, analytics, and market intelligence. Its 'Otaku' agent runs autonomous DeFi research with integrated Coinbase wallet capabilities, neither requiring a human in the loop.

What both platforms demonstrate is that the transition from AI agent to AI agent with a bank account is not a future capability, it is a current reality, operating at scale. The financial infrastructure described in Section 2 is what makes it possible.

Wallet Infrastructure and M&A Activity

If frameworks are the brains, wallets are the hands. An agent that can reason about a purchase but cannot sign a transaction is commercially stunted. This is why the wallet layer has attracted the most aggressive M&A activity of any part of the stack.

Stripe acquired Privy, a white-label embedded wallet provider powering **75 million** crypto wallets, in June 2025. Fireblocks acquired Dynamic, which powers **50 million** onchain accounts for clients including Kraken and Magic Eden, for approximately **$90 million** in October 2025. Offchain Labs, the company behind Arbitrum, acquired ZeroDev, a smart account SDK built on the ERC-7579 modular standard. In each case, a larger infrastructure company absorbed a standalone wallet provider to fill a gap in its stack, introducing a clear pattern where payment infrastructure wants to own the wallet layer.

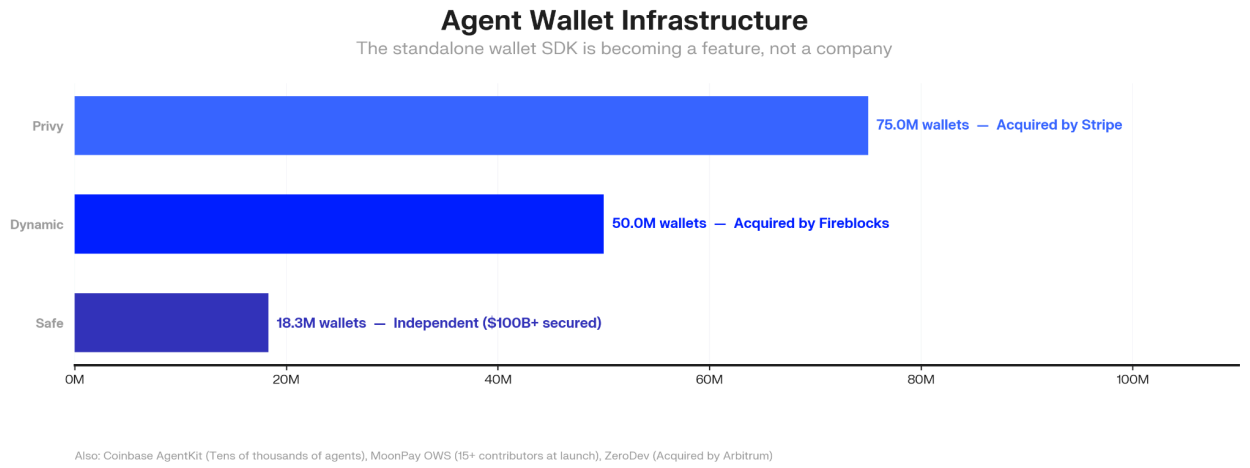

The survivors are the ones building moats at the protocol level. Safe, which remains independent, secures over **$100 billion** in assets across **18.3 million** smart accounts and processed **$600 billion** in transaction volume in 2025 alone. Perhaps the most striking datapoint from Safe is that **37%** of all Safe transactions on Gnosis Chain were made by AI agents, via the Olas network, rising to over **75%** on peak days. Coinbase's AgentKit, upgraded to Agentic Wallets in February 2026, has deployed **tens of thousands** of agents with built-in guardrails such as session caps, per-transaction limits, and allow lists restricting transfers to vetted contracts only. MoonPay launched its Open Wallet Standard on March 23 2026 with **15** contributing organisations including PayPal, OKX, and the Solana Foundation, standardising how agents interact with wallets across eight chain families.

The architectural consensus across all of these projects is converging on the same set of principles whereby private keys are never exposed to the agent process, signing is policy-gated with human-defined guardrails, and threshold cryptography ensures no single point of failure can drain funds. So, with a convergence in design, it's a race to see which wallet infrastructure will dominate and be leveraged by the 400,000 agents with purchasing capabilities.

Headless Merchants and Service Discovery

Agents with wallets still need somewhere to spend. This is where the concept of the 'headless merchant' enters the picture, a term coined by Noah Levine at a16z crypto to describe businesses with no website, no checkout flow, and no user accounts. Instead, there are just API endpoints with machine-readable pricing schemas. The agent reads the schema, sends a request, pays, and receives the output in a single exchange. The merchant incentive here is market expansion. A data provider like Nansen charging **$0.01** per agent query is not cannibalising its existing subscription revenue in the slightest, but monetising a class of traffic that previously generated no revenue at all. Agent buyers do not sign up for annual plans, they make millions of individual requests that no human sales process could service. For the merchant, the economics are zero customer acquisition cost, zero support overhead, and instant stablecoin settlement with no chargebacks or payment terms.

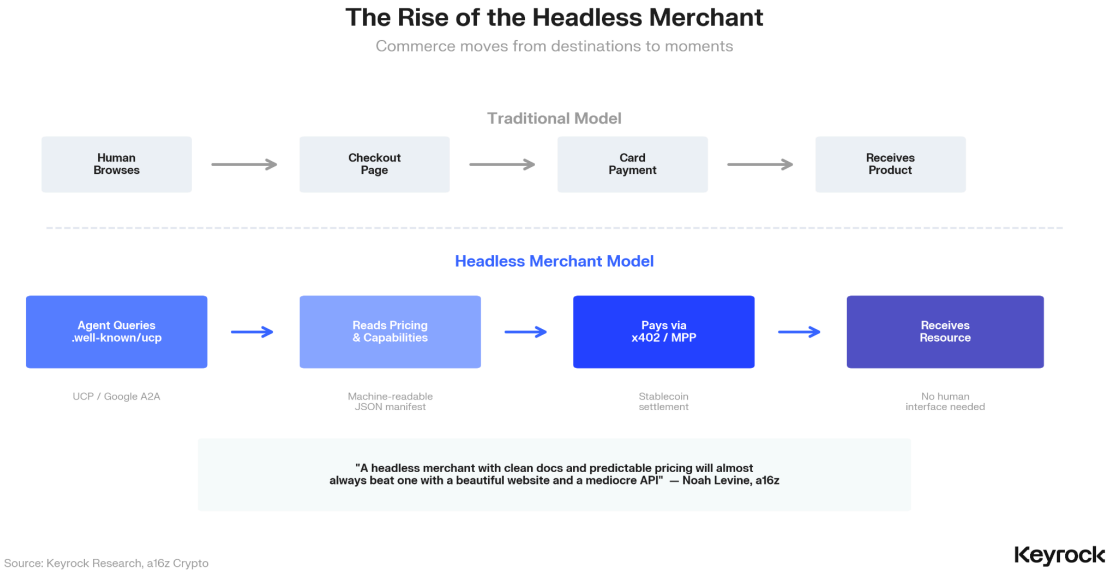

Google's Universal Commerce Protocol, launched in January 2026 with Shopify, Etsy, Target, Walmart, and Wayfair, is the most ambitious attempt to standardise this. Merchants publish a machine-readable JSON manifest at a standardised endpoint, effectively a business card for agents, declaring capabilities, supported payment methods, and API endpoints. An agent queries this manifest to learn what a merchant can do before attempting any transaction with no browsing required.

Nansen offers a concrete example of what this looks like in practice, given their platform at agents.nansen.ai integrates x402 to enable a pay-per-request model where agents access real-time onchain analytics, smart money flows, and token performance data at **$0.01** per query, settled instantly in USDC on Base.

The scale of agent-to-service commerce is already larger than most assume. As of Q1 2026, over **104,000** agents are registered across **15** or more directories and registries. The ecosystem has settled **$73 million** across approximately **176 million** transactions, with an average transaction size of **$0.31**.

Settlement Composition and Stablecoin Dominance

Perhaps the most striking finding from the data is not the volume itself but the composition. Of those **176 million** payments, **98.6%** settled in USDC. Stablecoins won the settlement layer for machine commerce almost by default to-date, because they were the only instrument that could handle sub-dollar transactions without the economics collapsing. The **98.6%** figure, while accurate in aggregate, conceals some nuance. Reppel notes that longtail tokens actually carry higher volumes than major cryptocurrencies in certain agent ecosystems, with platforms like Virtuals and ElizaOS transacting primarily in their native tokens. The USDC dominance is also partly geographic. USDT carries meaningful volume in APAC markets and on Tron and BNB Chain, mirroring the broader stablecoin market where USDC leads in North America and USDT leads in Asia. For businesses adopting x402, however, USDC remains the default because, as Reppel puts it, it is the easiest thing to explain to your accountants and operationalise in your business. Virtuals Protocol confirms this from the application layer: "Agent transactions are all paid in USDC. This is crucial for widespread adoption of agentic activities."

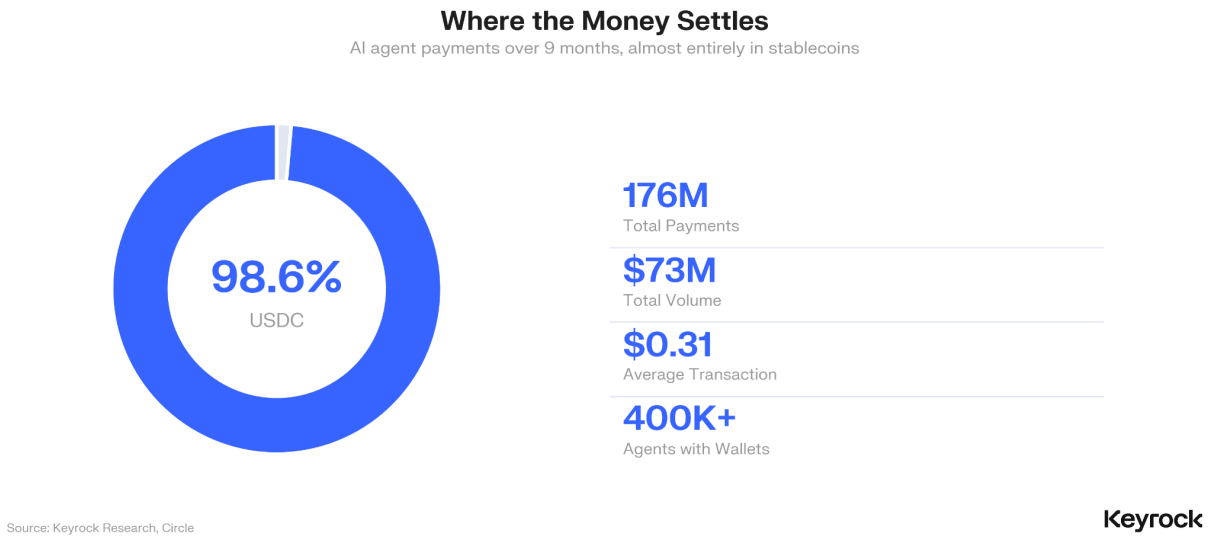

The **98.6%** USDC concentration is both a validation and a vulnerability, in that it validates Circle's positioning as the default settlement asset for machine commerce, but it also means the entire agent payments ecosystem is currently dependent on it. This is a lot of dependence on a single stablecoin issuer's reserve management, regulatory standing, and technical infrastructure. If Circle faces a regulatory challenge, a de-peg event, or even sustained downtime, the agent economy has no fallback. This is a systemic risk that nobody in the space is publicly discussing, and one we believe warrants serious attention as volumes scale.

The x402 volume data illustrates the trajectory. Starting from effectively **zero** in May 2025, weekly volume has grown to a sustained **$11-15 million** per week as of April 2026, with a peak of **$15.3 million** in the week of March 22. The November to December 2025 spike and subsequent correction followed by a stronger resurgence in March 2026 points towards organic adoption with the rise of improved AI models such as Anthropic's Opus 4.6.

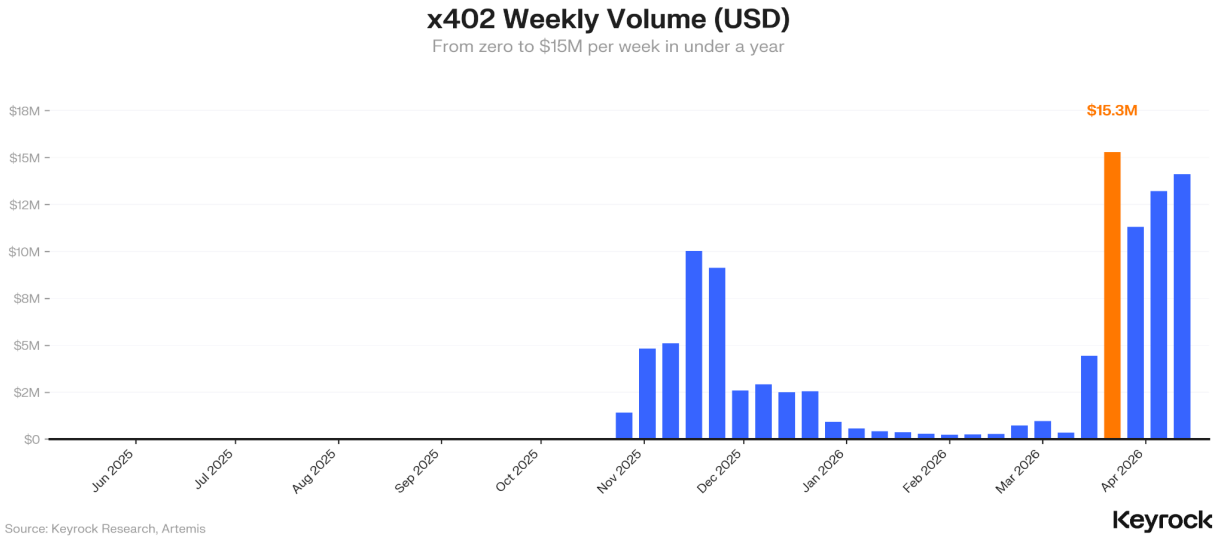

Brendan Ryan of Tempo offered a useful framing during our research interview, estimating that the majority of MPP volume remains stablecoin-based, characterised by lower value and higher frequency transactions, with stablecoin-settled requests being priced as low as **one-tenth** of a cent. Card-based transactions via Stripe are higher in value, typically in the **$5-10** range where card economics begin to make sense, but remain infrequent. Ryan views MPP as "the logical extreme" of the Stripe thesis, "Stripe's thesis is you just need to make it very easy for people to start businesses online. MPP is like, okay, what if you didn't even need to sign up on Stripe? You can just plug this thing in and start getting volume."

Agent Autonomy and Governance Models

The pipeline from framework to wallet to marketplace raises an uncomfortable question surrounding how much autonomy an agent should actually have. The four protocols from Section 2 implicitly answer this question differently, and their positions along the autonomy spectrum reveal a fundamental tension in the market.

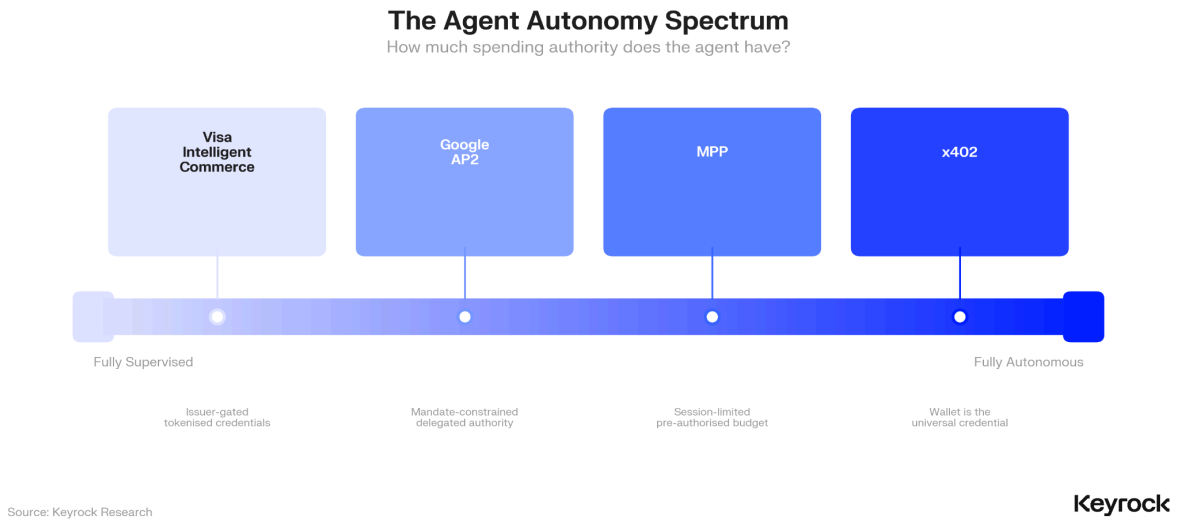

At one extreme, Visa's Intelligent Commerce provisions tokenised credentials with issuer-defined spend controls. The agent can pay, but only within tight boundaries set by a human and enforced by the card network. At the other extreme, x402 treats the wallet as a universal credential, whereby if the agent has funds, it can spend them with no permission required.

Google's AP2 and MPP sit in between, each with different mechanisms. AP2 uses cryptographic mandates, signed authorisation records that define precisely what an agent can and cannot do, essentially a programmable power of attorney. MPP uses sessions, where an agent pre-authorises a spending limit and then operates freely within that budget until it expires. Both approaches acknowledge that fully autonomous spending is where the market is heading, but that the governance infrastructure needs to exist first.

The market is moving rightward along this spectrum, towards greater agent autonomy. However, we do not believe that the pace of that movement will be determined by the technology, which is largely ready, but by the trust infrastructure that makes it safe. Our view is that the market will settle not at either extreme but in the uncomfortable middle, where agents operate with significant autonomy but within cryptographically enforced boundaries that humans can audit and revoke. The fully permissionless vision is intellectually appealing and technically elegant, but it assumes a level of AI reliability that does not yet exist. Until agents stop hallucinating, they probably shouldn't have unsupervised access to a user's funds.

Section 4: Transaction Economics and Liquidity

The previous sections mapped the protocols and the pipeline. We've set the scene of how agents can hold wallets, discover services, and pay for things. But the discussion so far has sidestepped questions that any market maker would ask first: where does the liquidity actually come from? And do the economics work at the transaction sizes we are talking about?

The honest answer is that the economics only work on one type of rail, the market structure is unlike anything we have seen before, and the liquidity dynamics create some uncomfortable concentration risks that nobody in the space is adequately addressing.

Card Fee Structures and the Micropayment Floor

Across **176 million** x402 payments to date, the median transaction sits between **$0.01** and **$0.10**, with **76%** of activity falling below the **$0.30** card-fee floor. That number tells you almost everything you need to know about why traditional payment rails cannot serve this market.

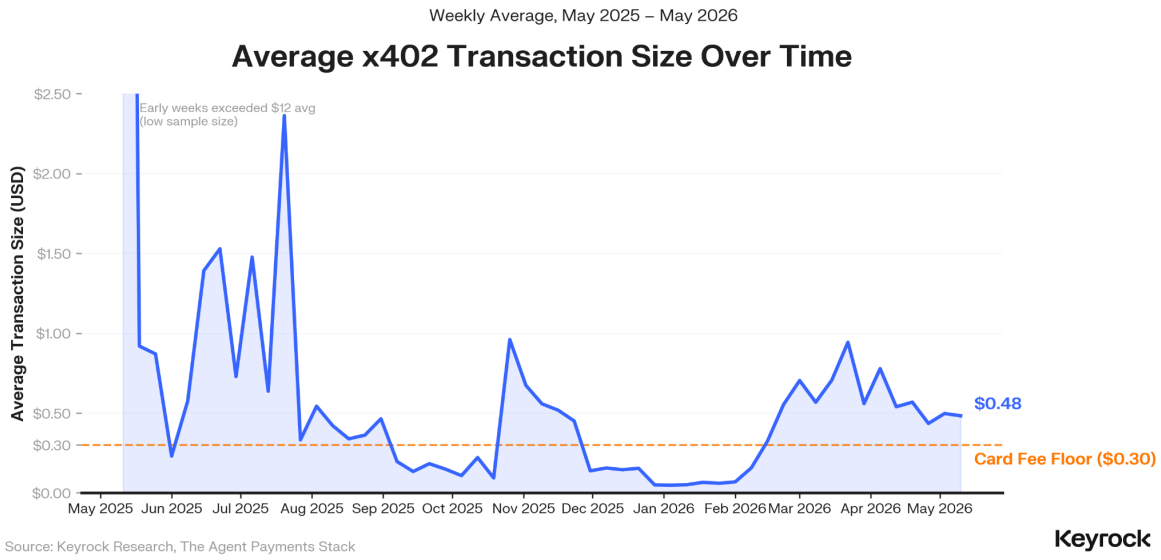

Source: Keyrock Research, The Agent Payments Stack

The weekly average transaction size tells a more nuanced story, in that early weeks showed inflated averages above **$12**, an artefact of tiny sample sizes when x402 had fewer than **100** transactions per week. As volume scaled through late 2025, the average collapsed and has since stabilised around **$0.48** over the past month, hovering just above the card-fee floor. The **$12-$15** monthly average cited earlier in this report reflects a brief period where a small number of higher-value transactions skewed the mean, and so the sustained reality is sub-dollar.

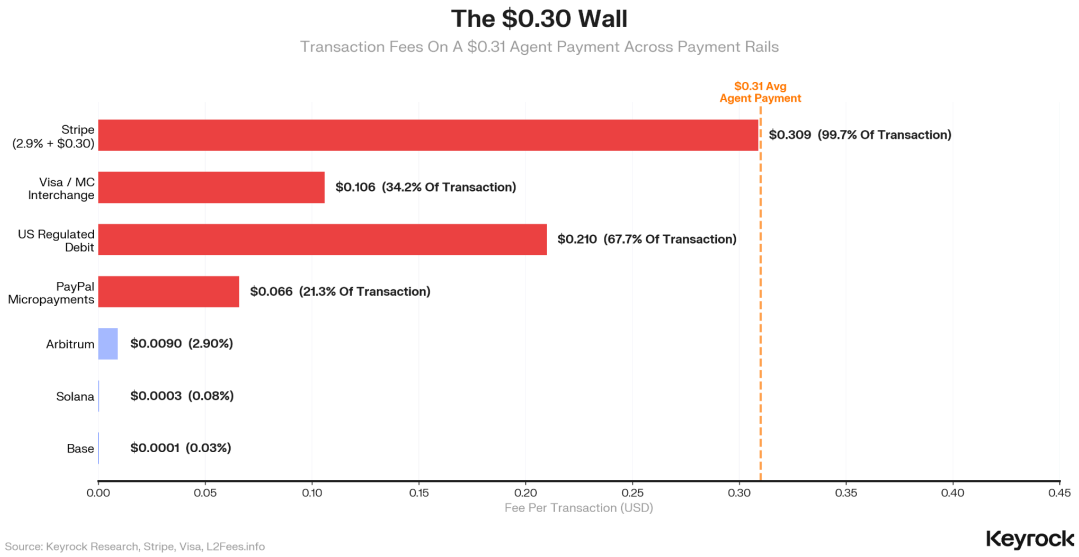

Stripe charges **2.9%** plus **$0.30** per transaction, so on a **$0.31** payment, the fixed fee alone consumes **96.8%** of the transaction value. With these economics, on a total fee of **$0.309**, the merchant receives just **$0.001**. Visa and Mastercard interchange averages roughly **$0.10** in fixed costs plus a percentage, eating **34%** of the payment. Even PayPal's special micropayment rate of **5%** plus **$0.05** takes a **fifth**.

Compare this to crypto Layer 2 rails, where a simple USDC transfer on Base costs approximately **$0.0001**, it equates to just **0.03%** of a **$0.31** transaction. On Solana and Arbitrum this is marginally higher at **$0.00025** and **$0.009** respectively. This three to four orders of magnitude difference is significant. The fixed-fee component of card processing creates what we call the **$0.30** wall, essentially a hard economic floor below which traditional payment rails simply cannot operate. Agent commerce, where **76%** of transactions fall below **$0.30** and the most common individual payment sits between **$0.01** and **$0.10**, lives almost entirely below this wall, meaning that L2s become the only viable settlement mechanism for machine-to-machine micropayments at this scale. It is worth noting that the **$0.30** wall is as much a business model problem as a technical one. Card networks could, in theory, compress their fee structures to accommodate micropayments, but the fixed-fee model is so deeply entrenched in the economics of four-party interchange that disrupting it would cannibalise the most profitable segment of existing card revenue. As Reppel frames it, "paying with stablecoins is actually the logical equivalent to paying in cash. If I hand you $10, there's no person who's going to take 3% of the $10. I just hand you in a store." The implication is that blockchains have not merely reduced the cost of payments. They have returned payments to their pre-intermediation economics. What x402 shows, in Reppel's view, is that "the floor of useful transaction size is actually on the orders of single digit cents, if not lower."

The counterargument, articulated most clearly by Kai Sheffield at Visa Crypto Labs, is what he calls the 'AI mullet'. This is a concept that mirrors the well known 'DeFi-mullet' whereby the process has cards in the front, stablecoins in the back. The consumer's agent pays with a card, preserving credit, rewards, and consumer protections, while the merchant settles in real-time stablecoins, avoiding the multi-day settlement delay that makes headless merchant economics painful. In this model, the **$0.30** wall is absorbed by the consumer side of the transaction, where card economics already work, while stablecoins handle the settlement side, where speed matters more than cost. The AI mullet has been put into practise by Visa CLI, currently in limited beta, which enables Touch ID-authenticated card payments from the developer terminal. American Express's ACE developer kit is another example of this in practice, which goes further by offering explicit liability coverage for agent errors, something no stablecoin protocol can match.

We think the AI mullet model is compelling for higher-value agent transactions, perhaps above **$5**, where card economics become viable and consumer protections are worth the fee premium. However, for the long tail of **$0.01–$0.10** API calls that constitute **76%** of current x402 activity, no amount of innovation in card infrastructure can overcome the fixed-fee floor. The future likely involves both, where we see cards for consumer-facing agent commerce where trust and liability matter, and stablecoins for machine-to-machine micropayments where they do not.

Stablecoin Velocity by Chain

One thing that stood out in our analysis of agent commerce is that the transaction profile is unlike any existing payment pattern. It goes against the traditional profiles of retail, which are high value low frequency, wholesale, which are very high value and low frequency, and traditional micropayments which are moderate value and moderate frequency. What we see with agent commerce is extremely high frequency, and extremely low value. This creates a distinctive velocity signature where the rate at which each dollar of supply is reused varies dramatically by chain, and the variation tells us where agent commerce is gravitating.

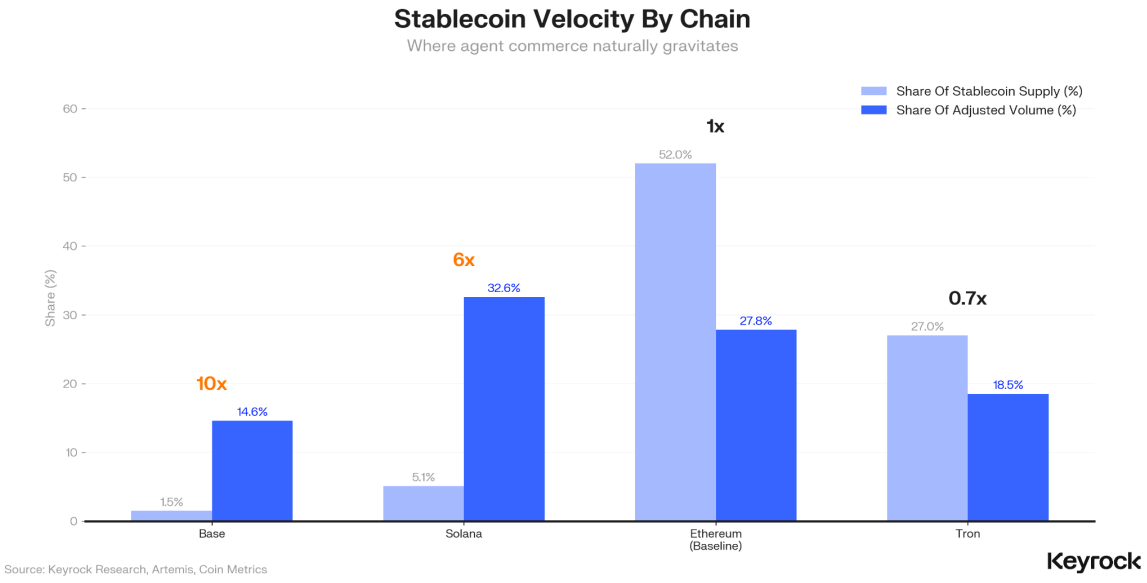

Base holds just **1.5%** of global stablecoin supply but accounts for **14.6%** of adjusted transfer volume, a velocity multiplier roughly **10x** that of Ethereum. Solana holds **5.1%** of supply but moves **32.6%** of volume, **six** times the Ethereum baseline, and Tron, despite holding **27%** of stablecoin supply, moves proportionally less value, sitting at **0.7x**.

These are the chains where agent commerce settles, and the velocity data explains why. High-velocity chains are characterised by low fees, fast finality, and deep stablecoin liquidity relative to their size. They are optimised for exactly the transaction profile that agents produce.

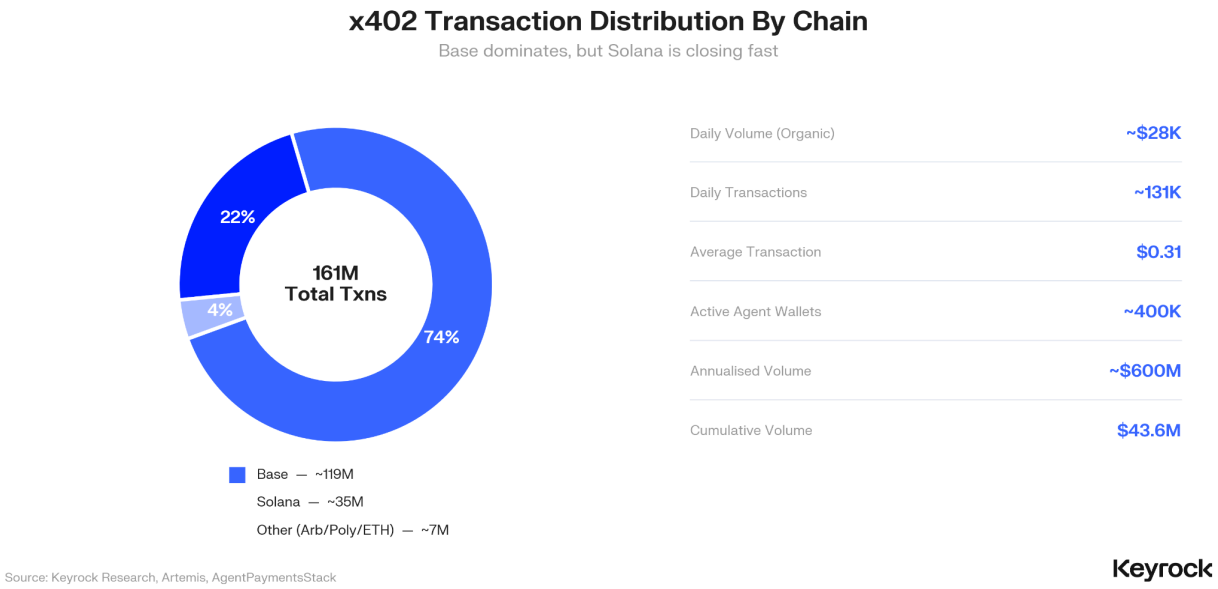

Concentration Risk and Vertical Integration

If agents are chain-loyal, and the dominant chain for agent commerce is Base, which we are seeing in our data, then the liquidity dynamics of the entire machine payments ecosystem currently flow through a single company's infrastructure.

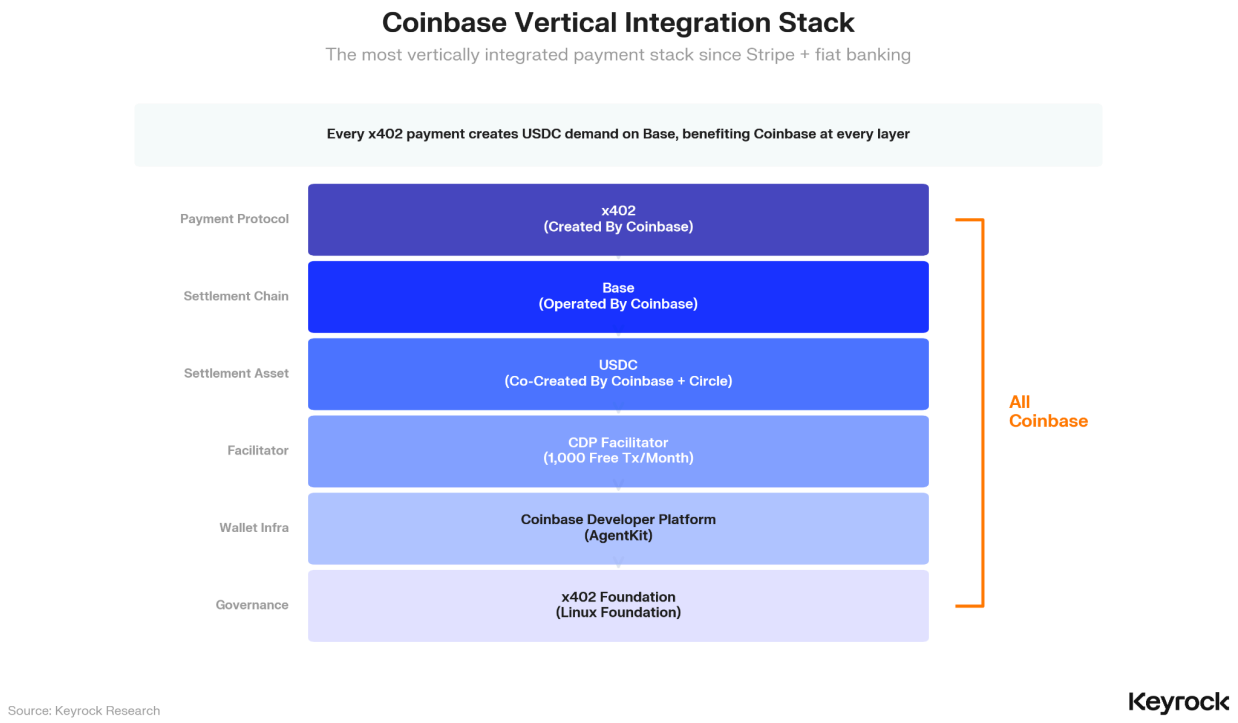

Virtually every x402 payment creates USDC demand on Base, and USDC is co-issued by Circle and Coinbase. Base is operated by Coinbase, and the default x402 facilitator is run by Coinbase's Developer Platform, offering **1,000** free transactions per month. Agent wallets are provisioned through Coinbase's AgentKit, and even the governance layer, the x402 Foundation, was co-founded by Coinbase, though it now sits under the Linux Foundation.

The result is the most vertically integrated payment stack assembled since Stripe built its relationship with fiat banking. Every layer reinforces the others, where more agents using x402 means more USDC demand on Base, which means more sequencer revenue for Coinbase, which means more capital to invest in agent tooling, which means more agents using x402.

In our opinion, the x402 protocol itself is genuinely open, in that it supports multiple chains and the Foundation governance is credibly independent. But protocol openness and infrastructure neutrality are different things. We made this observation about open protocols becoming captured by vertically integrated platforms in Section 2. In the liquidity context, the concern is sharper, where if **74%** of agent payment volume flows through a single company's chain, using a single company's co-issued stablecoin, processed by a single company's facilitator, the 'open protocol' framing becomes somewhat academic. Reppel acknowledges the concentration but frames it as a sequencing problem rather than a structural one. The x402 Foundation, now under the Linux Foundation, includes maintainers from across the industry, and he notes that "I don't make all the decisions for x402, which is not true of all other standards." The foundation is working on fiat payment integration and multi-chain expansion, but the question is whether decentralisation of governance can outpace centralisation of infrastructure.

Market Microstructure at Micro Scale

Today, the overwhelming majority of agent payments are direct USDC transfers, whereby the agent pays the merchant, no swap required. The liquidity question becomes relevant as agent commerce matures. Agents that earn revenue in USDC may need to convert to other assets for treasury management, cross-chain operations, or paying merchants on different rails. That future flow is what determines which market microstructure will serve machine commerce.

The transaction profile of agent commerce inverts traditional market-making economics. At **$0.31** average transaction size, a market maker capturing a **0.3%** spread earns **$0.00093** per trade, which is less than the gas cost of executing a swap on most chains. We see order-book market making as structurally unviable for agent commerce, given the maths simply does not work.

The obvious objection is that AMMs expose agents to MEV extraction, sandwich attacks, and front-running via public mempools. In practice, at an average transaction size of **$0.31**, the economics of MEV extraction work against the attacker. A sandwich attack requires two transactions (front-run and back-run), each incurring gas costs, and the extractable value from manipulating a **$0.31** swap is typically less than the gas required to execute the attack. Agent micropayments are, in effect, too small to be worth sandwiching. For larger agent transactions where MEV becomes viable, intent-based protocols like CoW Protocol route trades through competing solver networks that shield orders from public mempools entirely.

This is why the liquidity layer for machine payments will be dominated by automated market makers rather than traditional order books. AMMs have a structural advantage at micro scale in that liquidity providers deposit once and earn fees proportional to volume, without paying gas on every trade. Concentrated liquidity, as pioneered by Uniswap V3, achieves up to **4,000x** capital efficiency over full-range AMMs by allowing LPs to focus their capital around narrow price bands. For stablecoin pairs, which is where agent payment flow concentrates, this means a **$25 million** position concentrated at **$0.99-$1.01** provides the same depth as **$5 billion** in a full-range pool. Uniswap V4, launched in early 2026, takes this further with programmable hooks that enable dynamic fee adjustment and automated LP rebalancing.

We believe that the opportunity for institutional market makers is at the infrastructure layer. Running solver nodes for intent-based protocols like CoW Protocol, which has already handled **$33 billion** in lifetime volume through **30+** competing solver teams, operating x402 facilitators, and maintaining actively managed concentrated liquidity positions on the stablecoin pairs that underpin agent payment flows. The margin comes from processing throughput and gas optimisation, which favours firms with existing multi-venue algorithmic infrastructure.

Non-Human Transaction Activity and Adoption

The shift from subscription to pay-per-request is a structural requirement of how agents operate. A human analyst might subscribe to Nansen for **$150** per month because they will use it repeatedly over the course of the year. An agent spun up to answer a single research question might need three Nansen queries, two Dune queries, and one Alchemy call, then terminate. Forcing that agent through six separate subscription sign-ups, each requiring an account, an email address, and a stored payment method, is architecturally absurd. Pay-per-request eliminates this friction entirely, whereby the agent discovers the service, pays **$0.01**, receives the data, and moves on. The economics compound in the other direction too. A data provider with **400,000** agent customers making sporadic micropayments generates more revenue than one with **2,000** human subscribers, without a single support ticket or renewal negotiation.

A common objection to this entire analysis is that agent commerce is too early to matter, and in its current form, the volume data, at face value, supports this. x402's organic daily volume is approximately **$28,000**, annualising to roughly **$10 million**, while total volume including artificial activity has an annualised run-rate closer to **$600 million**, and while growing, this represents just **0.0003%** of Visa's **$14.5 trillion** annual volume.

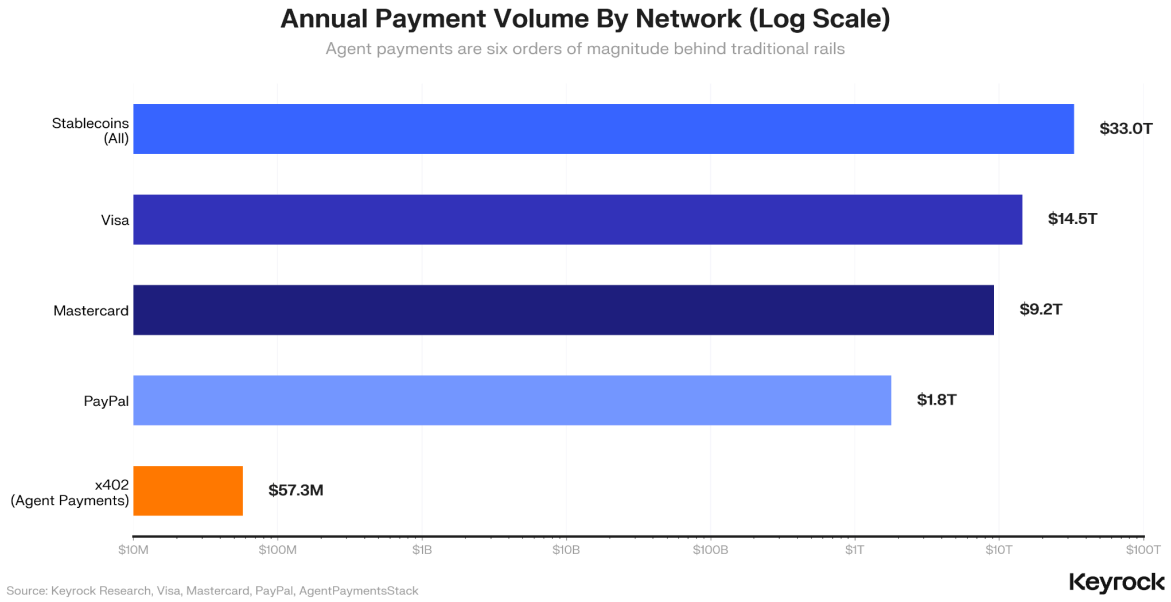

Six orders of magnitude is an enormous gap, although we think that taking this in isolation misses the point. We believe the machine economy is here already, it's simply just not doing commerce yet.

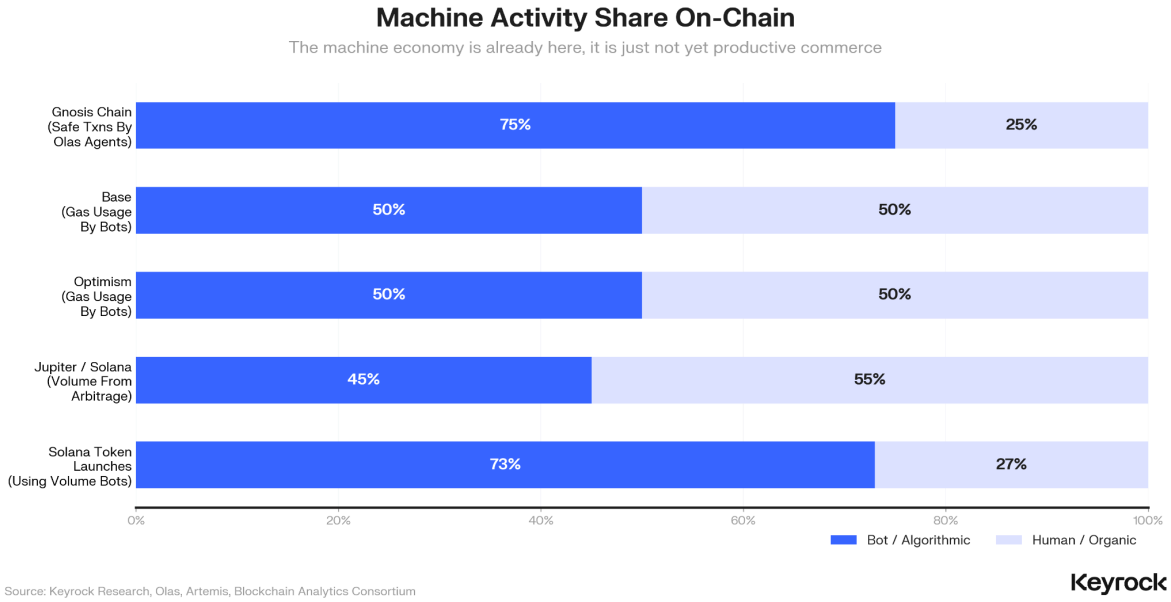

On Gnosis Chain, AI agents via the Olas network account for over **75%** of Safe transactions on peak days, and **37%** of all Safe transactions ever recorded. On Base and Optimism, bots and automated contracts consume more than **50%** of gas. On Solana, **40-50%** of Jupiter DEX volume is pure algorithmic arbitrage, while **73%** of the top 200 Solana token launches in Q3 2025 used professional volume bot services during their launch windows.

The point each of these statistics stand to prove is that machines already dominate onchain activity. What they are currently doing is extractive, be that arbitrage, MEV, or volume farming. The transition we are tracking is from extractive machine activity to productive machine commerce, i.e. from bots that take value to agents that create it for end-users. The infrastructure described in this report, be that the payment protocols, wallets, marketplaces, and liquidity mechanisms, is what makes that transition possible.

Total Addressable Market and Growth Assumptions

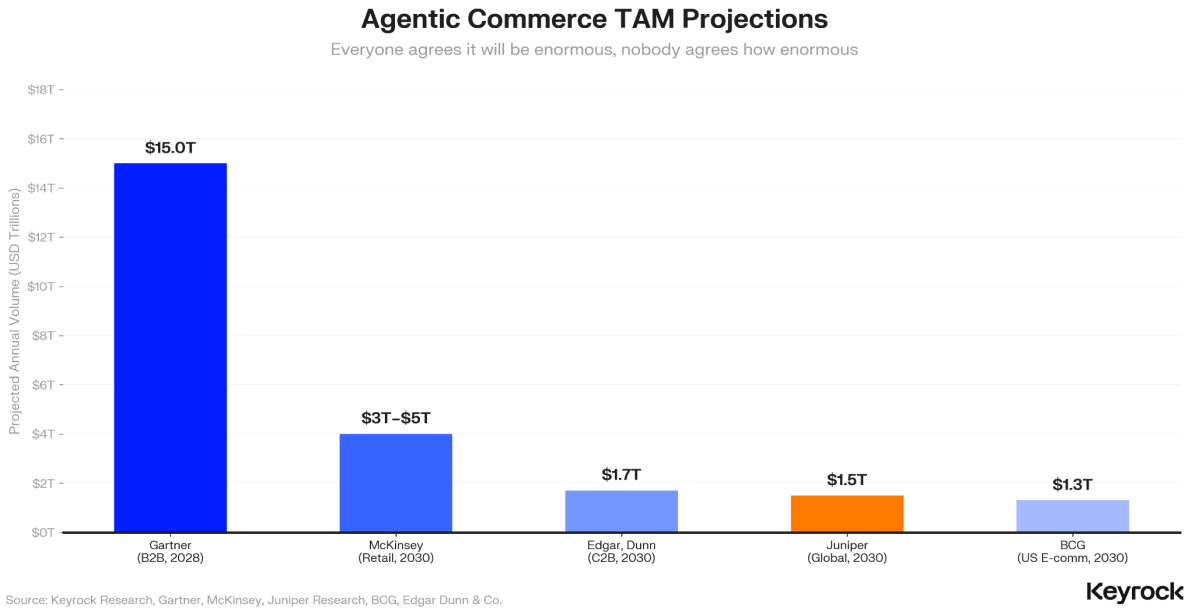

Looking across the TAM projections from industry analysis, one thing that stands out is that everyone agrees that agentic commerce will be enormous, although nobody agrees just how enormous. Gartner projects **$15 trillion** in AI-agent-intermediated B2B purchases by 2028, while McKinsey estimates **$3-5 trillion** in retail agentic commerce by 2030, and Juniper Research, in its April 2026 report, forecasts **$1.5 trillion**. BCG puts it at **$1.3 trillion** for US e-commerce alone.

We note these projections with appropriate scepticism, given that getting from **$600 million** in total annualised volume to **$1.5 trillion** in **four** years implies a compound annual growth rate of approximately **350%**. For context, the entire stablecoin market grew from roughly **$5 billion** to **$327 billion** over **six** years, a CAGR of approximately **100%**. The agent commerce projections require growth rates **three** to **five** times faster than the most explosive period in stablecoin history.

In principle, it's certainly possible, and the demand and exponential growth and adoption of AI supports the projection. In conjunction with this, the infrastructure is approaching readiness, and L2 fees are effectively zero, while agent frameworks are production grade. We also have wallet infrastructure consolidating into reliable, auditable systems. So, the constraint is the trust, governance and regulatory clarity, as opposed to the technology itself. Clarity on these elements would give institutions permission to connect their agents to real money, which is what we explore in Section 5.

Section 5: Regulatory Landscape

Every regulatory framework that governs payment infrastructure was built on the assumption that a human is somewhere in the transaction chain. However, every protocol described in this report was built to diminish the need for them. The gap between these two realities is the single largest unresolved risk in machine-native commerce.

Three major regulatory regimes are reaching full enforcement within weeks of each other this year, those being MiCA's transitional period which ends on July 1 2026, the GENIUS Act's implementation deadline that falls on July 18, and the EU AI Act's high-risk obligations that take effect on August 2. Each was drafted before the agentic payments wave, and none contain provisions for autonomous machine-to-machine transactions. The regulatory infrastructure and the payment infrastructure are maturing simultaneously, but in almost complete isolation from each other.

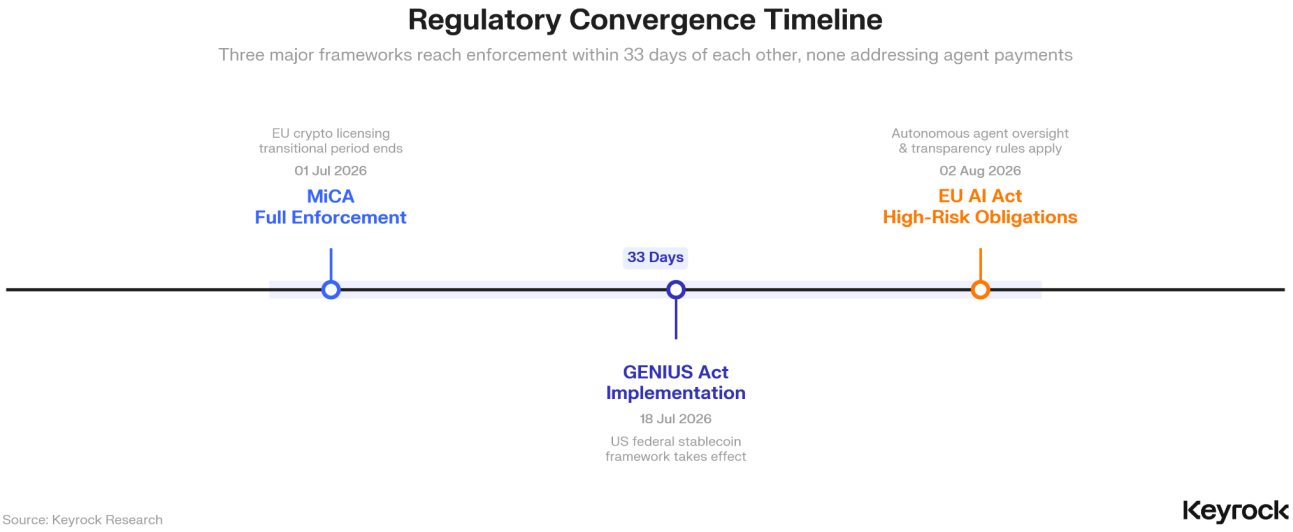

Legislative Gaps in Current Frameworks

The GENIUS Act provides the first US federal framework for payment stablecoins, mandating 1:1 reserve backing, AML/KYC compliance, and sanctions obligations. It is an important piece of legislation for the stablecoin ecosystem, but is also entirely void of commentary on AI agents. As the MIT Digital Currency Initiative noted, the Act treats stablecoin stability as a "balance-sheet problem" while ignoring the operational infrastructure, a critical gap for the always-on, automated environment in which agents operate. The EU AI Act captures autonomous agents under its high-risk classification, requiring human oversight and transparency obligations, but it was not designed for machine-to-machine payments where no human is present on either side. MiCA regulates crypto-asset service providers without contemplating scenarios where an AI agent autonomously selects counterparties, negotiates terms, and initiates transactions without real-time human oversight.

The result is that no US federal agency has issued specific guidance on AI agents holding wallets or transacting autonomously. The closest thing to regulatory engagement is the CFTC's Innovation Task Force, announced in March 2026, which formally includes autonomous AI systems in its mandate but has produced no rules or guidance. The Center for Data Innovation captured the situation in a March 2026 report titled with what may be the most accurate summary of the regulatory landscape: "Agentic Commerce is Coming, but Regulation Meant for Humans Will Slow It Down."

Liability and Consumer Protection

An unresolved issue we've identified from our analysis is that of liability. Ultimately, if the AI agent is making such key decisions, it calls into question whether the model provider, agent developer, deployer, user, or merchant is liable for erroneous purchases. The problem is compounded by the fact that stablecoin settlement fundamentally inverts the traditional liability model. With credit cards, the merchant bears chargeback risk and the consumer is protected. With stablecoins, as Reppel observes, "The merchant has no risk. If the stablecoins have landed in your wallet, they can't be pulled back. The risk moves to the consumer." This inversion means that existing consumer protection frameworks, built entirely around the chargeback mechanism, do not map onto stablecoin-settled agent commerce. The consumer, or more precisely the consumer's agent, needs to verify merchant reputation before transacting, not after.

The UK Competition and Markets Authority moved first, publishing binding guidance on March 9 2026 that places full legal responsibility on the business deploying the AI agent, with no safe harbour for AI errors and fines of up to 10% of worldwide turnover. The EU withdrew its dedicated AI Liability Directive in February 2025 and is instead relying on a patchwork of the AI Act, the revised Product Liability Directive, and general tort law. In the US, no federal regulator has addressed the question, and the most relevant existing law is likely UETA Section 10(b), a 27 year old provision requiring that providers of electronic agents give users a reasonable means to prevent or correct errors, or face automatic transaction reversibility. This is, of course, a vague law that cannot be waived by contract and frankly was written for a different era, but maps remarkably well onto the current problem.

American Express made a commercially significant move on April 14 2026, launching Agent Purchase Protection as part of its ACE developer kit, the first card network to explicitly cover agent errors. However, this protection only applies within Amex's registered ecosystem.

Sanctions Compliance and Enforcement Risk

Perhaps the most underappreciated risk sits at the intersection of autonomous agents and sanctions compliance. OFAC imposes strict liability, meaning a US person can be held civilly liable even if they did not know they were transacting with a sanctioned party. An AI agent executing thousands of unsupervised transactions per day has no native mechanism to screen counterparties. For example, the x402 protocol itself contains no built-in compliance layer. Sanctions screening is handled at the facilitator level, which means self-hosted facilitators operate with no default compliance tooling. Middleware solutions like AnChain.AI and CLEARAGENT are emerging to fill this gap, but they are bolted on rather than native to the protocols.

Recent enforcement actions against Exodus Movement, which settled with OFAC for **$3.1 million** in December 2025, and ShapeShift, which settled for **$750,000** in September 2025, demonstrate that the agency is willing to pursue digital asset intermediaries aggressively.

Agent Identity and Authentication Standards

Every prior revolution in commerce has ultimately resolved into a question of identity, for example, the merchant guilds of medieval Europe, the correspondent banking system, and the SSL certificate chain that underpins internet commerce. Each of these was, at its core, an answer to how one entity verifies that another is who it claims to be, and can be trusted to honour its obligations. History rhymes, and agent commerce is no different to its philosophical counterpart events, except that the entity seeking trust has no legal personhood, nor physical presence, nor intrinsic continuity of self, in the sense that an AI agent instantiated on Monday may share nothing with its Tuesday successor beyond a wallet address. Identity, in the human sense, assumes a persistent subject, but machine identity must be constructed from scratch, and the standards bodies are only now beginning to grapple with what that means.

NIST launched its AI Agent Standards Initiative in February 2026, explicitly asking how agents should be identified and authenticated in enterprise architectures. At least five IETF drafts are converging on agent identity management, including proposals for cryptographic binding between an agent and its deploying entity. We have AIS-1, which describes itself as the world's first open standard for AI agent identity, live on Base mainnet with a public comment period closing on June 30 2026. We also have Singapore's IMDA, who published the first cross-sector governance framework for AI agents at Davos in January 2026, and MetaComp who announced its "Know Your Agent" framework at Money20/20 Asia in April 2026, drawing explicitly on FATF Travel Rule principles to map accountability from agent back to principal. This shows the identity infrastructure is being built, but none of it is yet authoritative, and the deeper philosophical question of what constitutes identity for an entity that can be forked, cloned, or terminated between transactions, remains unaddressed by any standard.

This matters because identity is the prerequisite for everything else in the regulatory stack. Liability requires knowing who acted, KYC requires knowing who owns the wallet, and sanctions compliance requires knowing who is on the other side. Without a settled, interoperable standard for agent identity, we believe that every other regulatory question remains structurally unanswerable. Reppel, who co-authored ERC-8004, a proposed standard for onchain agent identity and reputation, views it as "one attempt to solve a really complex problem" and argues that more are needed. His hope is that the x402 Foundation, which now includes major payment providers and card networks, can serve as the venue for developing a consensus standard for agent identity, one that enables bi-directional trust where both the agent and the merchant can verify each other's reputation before transacting.

Our view is that regulation will be the binding constraint on how fast agent commerce scales, simply because the legal clarity that gives institutions permission to connect their agents to real money, and the liability framework that tells them who pays when something goes wrong, doesn't yet exist. The companies that solve the trust layer, whether through protocol design, network guarantees, or insurance products, will capture the governance layer of the stack. As we argued in Section 2, governance may be the most valuable layer of all.

Closing Thoughts

The infrastructure described in this report was built in twelve months, and that fact alone should command attention, whether agent commerce scales to the trillions that Gartner and McKinsey project or stalls in the low billions as a niche layer of the existing payments stack.

Our view is that the outcome depends on which bottleneck breaks first. The technology is largely ready, in that settlement is sub-second, fees are sub-cent, and wallet infrastructure is converging on auditable, policy-gated designs that even risk-averse institutions can accept. The binding constraints are regulatory clarity, liability frameworks, and agent identity standards. The companies and regulators that resolve these will set the pace, while the protocols and infrastructure we have documented will set the shape.

Brendan Ryan of Tempo offered a framing during our research that we think captures the trajectory well, "Dollar value will remain higher for humans for probably the next **ten** years, but raw transaction count will increasingly compound in favour of machines." Transaction data from Virtuals supports this. "A majority of the service offerings from agents are sub-dollar, ranging from $0.01 to $1," the team reports, "though more specialised agents offer premium services up to $10 per job." The pattern is already visible, where we see high-frequency, low-value at the base, with higher-value services emerging as trust infrastructure matures. We agree, and believe that the machine economy will build gradually, through **millions** of sub-dollar payments, each individually trivial but collectively large enough to reshape the payments stack.

What this means for market participants is clear enough, with Coinbase and Stripe building vertically integrated infrastructure that spans settlement, wallets, protocols and governance. The card networks are hedging by writing plugins rather than competing protocols. And the entire market remains dependent on a single stablecoin issuer whose reserve management, regulatory standing and technical resilience have never been tested at machine scale.

The trajectory is bottom-up, whereby crypto rails win micropayments by default because they are the only rails that can serve them. As millions of sub-dollar transactions build volume, they create the infrastructure, trust frameworks and regulatory precedent that make larger transactions viable on the same rails. The pull scales upward, in that what starts as **$0.01** API calls may, over time, pull **$50** SaaS payments and eventually **$500** procurement orders onto the same settlement layer. Payment rails have always been rebuilt this way: volume from below, not decree from above.

This marks the beginning of a complete rewrite for the way payments are executed, one where the payment is the authentication, the transaction is the contract, and the merchant has no storefront. The architecture is being laid now, in one of the biggest revolutions value transaction has seen to date.