

Commerce at Machine Speed

The Payment Infrastructure Agentic Commerce Requires

Author

Nicole Sandler, Chief Ecosystem Officer at Ubyx Inc. 

Key Contributors

Commerce at Machine Speed

The Payment Infrastructure Agentic Commerce Requires

Author: Nicole Sandler, Chief Ecosystem Officer at Ubyx Inc.

Key Contributors: Amazon, Barclays, Boku, Centre for Finance, Innovation and Technology (CFIT), Coinbase, Dfns, Digital Asset, Financial Conduct Authority (FCA), Global Financial Innovation Network (GFIN), Orbital, Paxos, Tempo

This paper sets out the infrastructure requirements of agentic commerce and assesses which architectures best meet them. It is the first of a three-paper series. It is not an endorsement of any specific payment instrument, issuer, or platform.

Section 1: Payments Were Built for People

Every payment rail built in the last fifty years carries the same assumption: a person initiates the transaction. Whether a cardholder tapping a terminal, an account holder approving a transfer, or a treasurer instructing a wire, the infrastructure was built around that person at every step.

That assumption is now being tested. The International Monetary Fund (IMF) describes this as a shift from ‘click-to-pay’ to ‘decide-to-pay’.^[1] Agentic AI, capable of autonomous planning, tool use, decision-making, and multi-agent interaction, is changing how commercial transactions are initiated and executed. AI agents are making purchase decisions, committing funds, and executing transactions. They are booking travel, placing procurement orders, managing subscription renewals, and optimising treasury positions. In each case, the agent can act without anyone approving the transaction in real time.

The shift is already underway at the edges of the financial system and moving toward the centre. Coinbase’s agentic wallets, launched in February 2026, enable AI agents to spend, earn and trade digital assets independently, managing funds without requiring human approval at the point of execution. Amazon’s Buy for Me feature enables AI agents to purchase products from third-party websites on a customer’s behalf. Alexa+ executes recurring purchase decisions within parameters the user sets at account level.

“Agents need to hold value, initiate payments, and settle transactions without a human approving each step. Building for that requires rethinking assumptions about identity, authorisation, and settlement that have been stable in payments for decades.”

- Keith Grose, UK CEO at Coinbase

Consider what happens when an AI agent attempts to make a purchase. Card infrastructure was built for one person, one terminal, one authentication step at a time. An agent executing micro-transactions at machine speed is a different use case entirely. Pay by Bank requires an account holder to authorise each payment within a consent framework written for humans reviewing options on a screen. Every rail, at the point it matters, asks for a person who is not there.

The problem is structural. The authorisation models, the fraud detection systems, the settlement windows, and the compliance frameworks were all designed around a human being as the initiating party. Without that person, the rails do not slow down; they stop.

[1] Sonja Davidovic and Hervé Tourpe, “How Agentic AI Will Reshape Payments,” IMF Note 2026/004, International Monetary Fund, April 2026.

“Institutional custody has always required enforceable rules about who can authorise what, under which conditions, with what limits. With AI agents, those rules can no longer rely on human judgment at the point of execution: authority must be delegated cryptographically and enforced in the wallet itself, so that every agent-initiated transaction is bounded, traceable, and auditable by design.”

- Clarisse Hagege, CEO and Founder at Dfns

The scale of the problem will not stay at the edges. Agentic commerce is expanding into procurement, financial services, logistics, and consumer platforms. As it does, the volume of machine-initiated transactions will grow faster than the existing infrastructure can absorb. An agent managing a supply chain does not execute one payment: it executes thousands, conditionally, across multiple jurisdictions, against delivery triggers and quality thresholds that may involve no real-time human monitoring.

Yet the regulatory frameworks governing these systems remain fragmented, nascent, and largely built on assumptions that a human is always in the loop. As of early 2026, only one jurisdiction in the world has published a dedicated governance framework for agentic AI, and no jurisdiction has legally defined the term. This gap between technological capability and regulatory coverage creates both risk and opportunity for organisations operating in financial services.

The urgency is clear: an estimated 93% of financial institutions plan to implement agentic AI within two years,^[2] and the agentic payments market is projected to grow 13-fold by 2032.^[3] Meanwhile, the entire payments ecosystem (identity verification, authentication, error-resolution rules, card network liability, and fraud models) is built on the assumption that the payer is a person.

The question is not which payment instrument wins. It is what properties a payment rail must have to serve commerce that no longer has a human at the point of transaction. Those properties follow directly from the nature of machine-initiated transactions: the ability to carry conditional logic, to authorise without a human present, to settle with finality in seconds, and to operate consistently across jurisdictions. The infrastructure that best delivers those properties, whatever form it takes, is where commercial volume will settle. These are not incremental improvements on what existing rails do; they are architectural requirements that existing rails were not designed to meet. The case for that infrastructure is what this paper makes.

Section 2: What Pay by Bank Got Right, and Where It Stops

Each generation of payment infrastructure was built to solve the problems the previous one left unsolved. Card networks achieved near-universal acceptance across issuers and merchants but introduced specific dependencies in doing so: scheme fees, interchange, fraud liability frameworks, and an authorisation architecture designed around human oversight at every transaction. Open banking, and the Pay by Bank movement that followed, addressed those specific costs directly, connecting accounts without the card network layer, improving merchant economics, and bringing identity verification closer to the source.

These were genuine advances for the problems they were designed to solve. Pay by Bank was built for open banking architecture, and open banking architecture was built for humans in specific jurisdictions. UK Open Banking and the EU’s Payment Services Directive (PSD2) are each distinct architectures with their own consent models, technical standards, and settlement mechanics. An agent operating across markets cannot rely on a single consent model: it faces different authorisation requirements in each jurisdiction, each designed with a human account holder in mind. There is no no cross-border standard and no programmable layer, no mechanism for conditional logic.

[2] Fenargo, Agentic AI in Compliance: From Concept to Operational Reality, 2025.

[3] Galileo Financial Technologies, The Agentic Payments Revolution: A Strategic Guide for Banks and Fintechs, 2025.

Pay by Bank can instruct a transfer, but it cannot execute a payment that releases only when a delivery is confirmed, holds funds in escrow pending a quality check, or routes conditionally across multiple counterparties depending on real-time pricing. The authorisation model runs deeper than a technical limitation: it is the design.

“Open Banking consent models were designed around a human who is present, authenticated, and accountable at the point of transaction. When the initiating party is an agent operating within a mandate, that architecture does not map cleanly. The authorisation question needs to be rebuilt from different assumptions.”

- Anna Wallace, CEO at the Centre for Finance, Innovation and Technology

In domestic markets, real-time payment systems such as Faster Payments create the appearance of instant settlement, often backed by intraday credit rather than final settlement. Cross-border funds move through correspondent chains where instruction and finality can be separated by days. For an agent executing thousands of conditional transactions simultaneously across multiple jurisdictions, the reconciliation complexity compounds quickly.

None of this diminishes what Pay by Bank achieved. Agentic commerce presents problems that the current architecture was not designed to solve: programmability, jurisdictional consistency, atomic settlement, and machine-native authorisation. The next rail was always going to be purpose-built for the commerce it serves.

Section 3: The Properties a Payment Rail Must Have

Before asking which payment infrastructure serves agentic commerce, it is worth specifying what that infrastructure must do. Six properties follow directly from the nature of machine-initiated transactions. They are not properties that any single technology was designed to deliver: they are requirements that the commerce itself generates. No existing payment rail was built with all six in mind. The question of which infrastructure best meets them (and how the gap between current capability and future requirement gets closed) is what the rest of this paper examines.

The first is programmability. Payment execution must be capable of carrying conditional logic natively: releasing funds on verified delivery, holding value in escrow pending a quality check, routing conditionally across multiple counterparties depending on real-time pricing. This is not a feature that can be layered onto a conventional payment instruction after the fact: it requires the conditions to be embedded in the transaction itself, executable without human intervention at each step. The infrastructure that serves agentic commerce must make programmability a property of the transaction, not an operational process built around it.

“Agent-initiated transactions can involve conditional logic: pay on delivery, release on verification, revert on timeout. Settlement infrastructure that executes conditional logic natively, at the protocol level, produces outcomes that are auditable and deterministic and fundamentally more resilient for all participants.”

- Shaul Kfir, Co-Founder of Digital Asset – Creators of the Canton Network

The second is technical interoperability across jurisdictions. Each major jurisdiction operates its own consent model, technical standard, and settlement mechanic, built for its own regulatory context. An agent operating globally cannot carry a different authorisation configuration for every market it crosses. The infrastructure that serves agentic commerce must operate on a consistent protocol regardless of where the parties are located. Regulatory requirements will continue to vary by jurisdiction (a governance question, not a technical one), but the settlement layer underneath must be consistent.

The third is settlement finality. Conventional cross-border settlement operates on next-day (T+1) or longer cycles, requiring businesses to carry credit facilities and working capital buffers against the gap. Services that create the appearance of instant cross-border transfers typically do so by pre-funding local accounts in each market, moving the capital buffer rather than eliminating it. An agent executing thousands of simultaneous conditional transactions across multiple jurisdictions cannot carry that overhead. The infrastructure that serves agentic commerce must settle in seconds, making real-time reconciliation possible at the volumes agents generate.

“The properties that make stablecoins suitable for agent-initiated transactions are: settlement finality, deterministic execution, and on-chain verifiability. When a payment is made, the counterparty needs certainty that it is final. Given the speed machines operate at, pending, probabilistic, or reversible settlement is not workable when the agent that initiated the transaction has already moved to the next instruction.”

- Simon Taylor, Head of Market Development at Tempo




The fourth is a machine-native authorisation model. Payment authorisation was designed around a human payer: someone who can be identified, whose consent is on record, and against whom liability can be assessed. Strong Customer Authentication (SCA) requirements under PSD2 require an aware, actively authorising payer at the point of transaction. That is sound for consumer protection; it also means agent-initiated transactions cannot satisfy the requirement as written. The infrastructure that serves agentic commerce must embed authorisation logic in the transaction itself, defined and agreed in advance, executed within those parameters without requiring human approval at each step.































The fifth is transaction cost economics suited to micro-transaction volumes. A card payment on a low-value transaction faces an interchange floor that makes it commercially unworkable at the micro-transaction volumes agentic commerce generates. Open banking frameworks improve on this but remain subject to per-transaction costs that compound at scale. The infrastructure that serves agentic commerce must carry flat costs with no minimum viable transaction size, enabling millions of micro-transactions at price points that card rails cannot reach.

The sixth is a wallet architecture that does not require a human account holder. A bank account requires a legal account holder: a person or a corporate entity with human directors behind it. An AI agent executing transactions autonomously is not a legal person and cannot hold a bank account in its own right. The infrastructure that serves agentic commerce must provide an asset custody layer that a non-human principal can control, holding and deploying value under the governance of a defined authorisation mandate, without requiring a human to be attached to the account at the point of execution

Taken together, these six requirements describe an architecture that existing payment rails were not designed to deliver. That is not a criticism of those rails: they were designed for the commerce they served, and they served it well. The question is which infrastructure comes closest to meeting all six requirements for the commerce that is now emerging.

Figure 1. Payment rails assessed against the six requirements

 No
 Partial / In Development
 Yes

	Card Networks	Pay by Bank Open Banking	CBDCs	Stablecoins	Tokenised Deposits
Programmability					
Cross-jurisdiction interoperability					
Settlement finality (seconds)					
Machine-native authorisation					
Micro-transaction economics					
Non-human wallet architecture					

No single architecture currently satisfies all six in full. Card networks and open banking frameworks offer broad acceptance and established consumer protection, but retain human authorisation dependencies at their core and do not deliver cross-border technical consistency or programmable conditional settlement. Central bank digital currencies (CBDCs) are being developed with some of these properties in mind, but most retail CBDC architectures under development remain jurisdiction-specific by design and do not yet address machine-native authorisation or cross-border interoperability at the protocol level. Upgraded card rails are investing in programmable payment capabilities, but the human authorisation dependency is structural rather than incidental.

Stablecoins, governed by smart contracts and settled on-chain, currently come closest to meeting all six requirements together. The clearing and governance standards that make them work across issuers, platforms, and geographies are still forming, and the architecture is not complete. Of the available options, the gap between current capability and what agentic commerce requires is narrowest here.

“Agentic commerce moves beyond the human-in-the-loop model. Stablecoins introduce the right primitives: programmability, instant finality, and global reach. With neutral clearing and par settlement across issuers, they can scale into a true global network.”

- Walter Hessert, Head of Strategy at Paxos

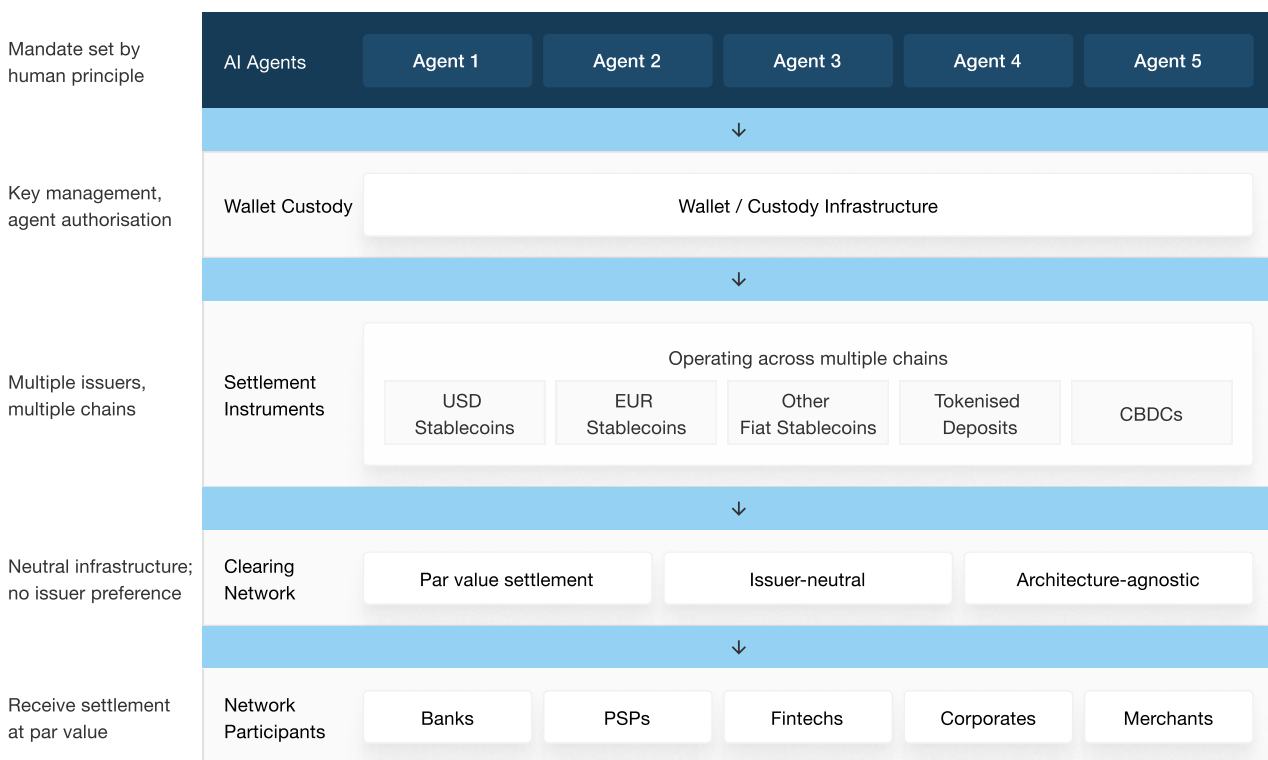
That conclusion requires a qualification on value stability. Even fiat-pegged instruments can trade at a marginal discount under certain liquidity conditions, which is why par value settlement is a design requirement of any clearing infrastructure serving agentic commerce at scale. CBDCs, issued and backed by the central bank, carry no equivalent stability risk, but most programmes remain in development and current architectures are jurisdiction-specific, without the cross-border interoperability or machine-native authorisation that agentic commerce requires. Tokenised deposits are advancing on a parallel track, with institutional deployment already underway. Tokenised deposits and CBDCs that develop to meet these requirements will qualify for the same network.

Section 4: A Rail Is Not a Network

The six requirements above describe what a payment rail must do. They do not, by themselves, describe a network. A rail that meets all six requirements in isolation still cannot serve agentic commerce at scale if it cannot interoperate with other rails, other issuers, and other institutions. The distinction between a rail and a network is the critical infrastructure question, and it applies regardless of which settlement architecture carries the volume.

The fragmentation risk is already visible. Programmable settlement infrastructure is developing across multiple architectures, issuers, and chains simultaneously. Two parties operating on different settlement architectures cannot transact natively without a shared clearing layer: navigating conversion, slippage, and custody friction independently defeats the purpose of programmable settlement. A clearing layer operating across architectures and issuers at scale resolves this. Without broad adoption of one, the proliferation of programmable settlement instruments risks producing the same interoperability problem that fragmented correspondent banking for decades: every participant connecting to every other independently, at cost, with no common standard underneath.

Figure 2. From agent instruction to settled transaction



CBDCs: emerging - included as forward-looking; not yet operational at scale

The major card networks showed what a clearing network achieves: settlement between any card issuer and any merchant within them, regardless of which bank issued the card. The card was the visible product; the clearing network was the infrastructure that gave it value. Programmable settlement instruments, of which stablecoins are currently the most developed, are at the card-creation stage. Broad adoption of clearing infrastructure is what determines whether any of them reach their potential as a network.

The scale of the challenge becomes clear on any large platform. A global marketplace operating across, for example, the US, EU, UK, Japan, India, Brazil, the Middle East, and Africa faces multiple settlement architectures, multiple receiving institutions, and multiple regulatory regimes: different AML (anti-money laundering) and KYC (know your customer) requirements, data localisation rules, and consumer protection frameworks in each market. That is the many-to-many problem in its most extreme form, and

bilateral integrations cannot serve it. Every issuer connecting to every receiver independently produces exactly the fragmentation that neutral clearing infrastructure was designed to solve, regardless of which instruments those issuers are operating.

"Clearing infrastructure that lets value move at par between issuers is one ingredient agentic commerce needs. The full picture is broader: orchestration across rails, transaction controls, fiat connectivity, licensing, and compliance tooling all have to operate together. The infrastructure question isn't only which clearing layer wins - it's how all of the layers above and beside it interoperate for a payment to actually work."

- Chris Mason, CEO at Orbital

"Third-party sellers on Amazon's marketplace serve customers across more than twenty countries. Ensuring that each seller receives accurate, timely payments - in the right currency, on the right settlement cycle, and in compliance with local regulatory requirements - is essential to maintaining a reliable selling experience. As transaction volumes grow, including through AI-assisted purchasing, the coordination required to deliver that consistency scales with every seller and every market."

- Amira Karim, Head of Payments and Financial Services Public Policy - International Public Policy, Amazon Consumer

Two further constraints sharpen the argument. No large merchant, regulated bank, corporate, or technology platform will build its settlement infrastructure around a single issuer or a single architecture. Concentration on one issuer introduces regulatory exposure to that issuer's compliance posture, creates a single point of failure, and gives the issuer pricing power over every participant dependent on its network. The clearing layer has to be issuer-neutral and architecture-neutral to be trusted. Genuine infrastructure serves all participants equally: a clearing layer that advantages any single issuer or instrument type will not attract the trust required to function as a network.

Par value settlement compounds this. In a fragmented market, an agent settling across different instruments and architectures faces conversion spread and discount risk with every transaction. Settling at face value, regardless of which issuer's instrument is presented, is what makes the rail genuinely fungible for commercial use. Without it, agents become unintended currency traders, managing basis risk on top of every commercial transaction they execute. This requirement applies to any programmable settlement architecture: the case for par value clearing is not specific to stablecoins. It is a property of any neutral clearing layer that agentic commerce requires.

Two significant use cases for stablecoins are e-commerce and corporate treasury management, and the clearing protocol that serves both is not simply a payments protocol: it is the financial infrastructure layer that programmable money requires across its most important applications.

Section 5: The Infrastructure Question

The shift toward machine-executed commerce is gaining momentum, but the infrastructure question is being answered now, ahead of the volume.

“The current wave of agentic commerce is skewed toward discovery, with payments still secondary. Early adoption of agentic payments is likely to come from B2B use cases such as supply restocking and manufacturing inputs, while consumers remain more focused on using agents for product discovery than for completing transactions.”

- Yi Hahn Chin, SVP, Money Movement Innovation at Boku

The commercial case is already visible. For a merchant or marketplace disbursing to sellers across multiple markets, the total cost stack of conventional payment infrastructure includes card interchange, scheme fees, FX conversion on every cross-border transaction, and the working capital cost of T+1 settlement cycles. Stablecoin clearing at par, settling in seconds, removes each of those. For banks and payment providers, the institutions building compliance infrastructure for agent-initiated flows now are generating the operational evidence from which industry standards will be made.

“Corporate clients that take consumer payments want to know how agent-initiated transactions are authenticated, reconciled, and traced through the payment chain so that they can unlock the potential. Those questions concern operational readiness and as much as technology. Institutions that address these issues now will be better placed to help businesses running agent-driven commerce capitalise on the opportunity.”

- Jason Lalor, CEO of Barclaycard Payments

The clearing and governance standards that make stablecoins work across issuers, platforms, and geographies are still forming. Tokenised deposits and CBDCs developing to meet the same six requirements will qualify for the same clearing infrastructure.

“When authority is delegated to an autonomous system rather than exercised by a person at the point of transaction, the frameworks governing consumer protection, liability, and dispute resolution require re-examination across the full delegation chain. How delegated authority is scoped, and how authentication standards are defined and allocated, are questions regulators are actively working through.”

- Colin Payne, Head of Innovation at the Financial Conduct Authority (FCA) and Chair of the Global Financial Innovation Network (GFIN)

The design choices being made now will determine who sets those terms. How accountability for agent-initiated transactions is defined, and how liability is allocated across the chain, are questions that require regulatory decisions alongside infrastructure investment. Those questions are the subject of the second paper in this series.