Commerce at Machine Speed, Part III

Know Your Agent - Identity, Mandate, and Accountability in Autonomous Commerce

Author

Nicole Sandler, Chief Ecosystem Officer at Ubyx Inc.

Key Contributors

Commerce at Machine Speed, Part III

Know Your Agent - Identity, Mandate, and Accountability in Autonomous Commerce

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

Key Contributors: Absa, Amazon, Bangko Sentral ng Pilipinas (BSP), Credentials and Authorization Registry of Agentic Commerce (CLARC), Financial Conduct Authority (FCA), Global Financial Innovation Network (GFIN), Global Legal Entity Identifier Foundation (GLEIF), Kulipa, Matter Labs, NatWest, NVNM Chain, Ripple, StraitsX, TRM Labs, Utila, Yellow Card, Yuno

This paper addresses the identity, mandate, and accountability standards that autonomous agents require in commercial transactions, and is the third in a three-paper series. It does not endorse any specific payment instrument, issuer, platform, protocol, chain, or standards body.

Section 1: From KYC to KYA

Every commercial transaction is built on a question that current verification infrastructure was designed to answer: who is this counterparty, and what are they authorised to do? When the transacting party is an autonomous agent, current frameworks were not designed to address either part of the question, and the gap between that question and a workable response is widening as agent transaction volumes grow across procurement, treasury, subscription management, and consumer commerce. For merchants considering whether and how to open agent channels at scale, the practical barrier is the absence of a framework for resolving disputes when an agent executes outside its intended scope. An agent carries a wallet address, a set of programmed parameters, and a principal somewhere upstream who authorised it to act. What it does not carry is legal personhood, a transaction history of its own, or any capacity to bear liability in its own right. The delegation chain between its wallet and the ultimate human or corporate principal may pass through a developer, a deployment platform, multiple sub-agents, and several jurisdictions before reaching settlement.

The first two papers in this series established what agentic commerce requires at the infrastructure and governance layers: on-chain settlement with finality, a mutualised acceptance network across issuers and jurisdictions, and the accountability frameworks determining how liability is allocated when an agent acts outside its mandate. Both rest on an assumption that neither addresses directly: whether participants can reliably identify who they are transacting with, on whose authority, and within what scope.

Current Artificial Intelligence (AI) agents add a further dimension to this challenge. Each operates within a context window: when a session ends, the agent retains no memory of what it did, what it was authorised to do, or who instructed it. An agent's identity, mandate, principal chain, and audit log cannot be stored within the agent and relied upon across transactions. Locating all four in external infrastructure, verifiable by any counterparty at the point of transaction, makes the agent's memory state irrelevant to the integrity of each transaction.

Know Your Agent (KYA) applies the logic of Know Your Customer (KYC) to the agent itself: what it is, what it is authorised to do, on whose behalf it acts, and what record exists of its actions. It operates through four components: an identity credential establishing the agent as a registered entity within a defined governance structure; a mandate certificate specifying the scope of authority delegated to it; a principal chain tracing delegation from the originating principal to the acting agent; and an audit log capturing each transaction in attributable form. The sections that follow address each participant in turn: merchants, banks, payment service providers (PSPs), issuers, chains, and wallet providers, each facing a distinct version of the same design problem.

> *"The legal status of an AI agent as a transacting party remains unresolved across jurisdictions, and the standards being developed now need to work across different regulatory frameworks, whatever position jurisdictions ultimately take."*
>
> **— Mamerto Tangonan, Deputy Governor at Bangko Sentral ng Pilipinas**

Figure 1: The KYA Credential Architecture
The four components and the participants responsible for verifying each at the point of transaction

IDENTITY CREDENTIAL	MANDATE CERTIFICATE	PRINCIPAL CHAIN	AUDIT LOG
Establishes the agent as a registered entity operating within a defined governance structure	Specifies the scope of authority delegated to the agent, including spending limits, counterparty permissions, and conditions	Traces delegation from the originating human or corporate principal through to the acting agent	Captures each transaction in attributable form, verifiable after the fact by any authorised party
VERIFIED BY Merchant, PSP, Bank, Issuer, Chain, Wallet	**VERIFIED BY** Merchant, PSP, Bank, Issuer	**VERIFIED BY** PSP, Bank, Chain	**ACCESSED BY** Bank, Regulator, Merchant (on dispute)

Section 2: Merchants - Verification Without a Human

When a principal provides an agent with a payment instrument, the merchant's verification question is largely answered at the point of transaction. A properly authenticated instruction from that instrument meets the merchant's standard for proceeding. The framework's limits emerge afterwards: when the principal disputes that the agent acted within its intended scope, the chargeback and error resolution system cannot readily determine whether the error originated with the agent's parameters, the developer's code, or the platform's deployment.

Virtual card accounts, already available within card network rails, allow a principal to limit an agent's spending to specific merchant categories, amounts, and time windows. The equivalent capability is technically achievable for tokenised money as well. What does not yet exist is an agreed standard for expressing and verifying those conditions across different tokenised money systems.

An agent operating outside its principal's intended scope may be executing on incorrectly specified parameters, with no fraudulent intent, but the merchant may carry the transactional exposure where no scope-limiting instrument is in place. The chargeback and error resolution frameworks were designed for human consumers who can articulate a grievance and invoke a specific protection.

The KYA audit log provides the dispute resolution mechanism that current infrastructure lacks: a verifiable, attributable record of what the agent was authorised to do, what it actually did, and at which point in the chain any deviation occurred. For the merchant, this produces an evidence-based assessment against a documented mandate. For the consumer, the principal chain identifies the accountable party, whether the developer, the deployer, or the platform that set the agent's parameters, giving the consumer a clear route to the responsible party.

> *"The authorisation gap in enterprise agentic use originates inside the enterprise. Bilateral connections cannot verify the delegation chain from corporate governance to the agent action, and that is a shared registry problem rather than an API problem."*
>
> **— Sandra Lam, Founding Committee Member at CLARC**

A merchant that can read an agent's mandate at the point of transaction gains capabilities that current infrastructure does not provide: to differentiate between credentialed and uncredentialed agents, apply different risk treatment to each, and build commercial terms for agents from principals the merchant has chosen to accept. For a large marketplace operating across multiple jurisdictions, this is the difference between treating agent-initiated volume as a risk category and treating it as a commercial channel.

Mandate verification at the point of transaction is a commercial problem for the merchant. By the time the same transaction reaches the bank, it becomes a regulatory one.

> *"An agent's authority is established once, then enforced at every payment, before execution rather than reconciled after it. The orchestration layer that already connects merchants, issuers, and rails is well placed to do both: verify a programmable mandate once, enforce it at the point of each transaction, and leave a record any regulator can follow."*
>
> **— Edwin Poot, Chief Technology Officer at Yuno**

Section 3: Banks - The AML Chain Problem

Banks face the KYA problem in its most acute regulatory form: they hold the accounts where agent transactions settle and bear the anti-money laundering (AML) and KYC obligations those accounts carry.

Banking account structures were not designed for AI agents. Business accounts are available to corporate clients running agent-based operations, but their terms, authorisation structures, and monitoring frameworks assume a human signatory whose activity follows recognisable patterns. The deeper challenge is the AML obligation: existing frameworks place the duty to know the customer on the bank, but when the transacting party is an agent acting across a delegation chain, the bank holds the least information about the chain and carries direct AML obligations for what settles through its accounts.

> *"As autonomous agents begin to transact at scale, compliance frameworks will need to evolve beyond human monitoring. Agent-driven activity exhibits distinct behavioural patterns, creating an opportunity to develop new monitoring baselines and risk indicators specifically designed for agentic commerce."*
>
> **— Ari Redbord, Global Head of Policy and Government Affairs at TRM Labs**

> *"Before a bank entrusts an agent to act on a customer's behalf, it needs more than just confidence. It needs verifiable, real-time assurance that the agent is operating within the customer's defined parameters. Solving this isn't just a technical checkbox, it's the trust foundation that will determine whether agentic payments can scale beyond theory and become a real, everyday capability."*
>
> **— Lee McNabb, Head of Payments and Digital Assets at NatWest Group**

The bank cannot build that monitoring capability from its position in the chain alone. The data it needs to separate legitimate agent flows from suspicious activity sits at the processing layer, where the payment service provider handles each transaction.

Section 4: PSPs - Fraud Models Built for People

PSPs sit between the agent and the bank, and the KYA problem presents at that layer as both a commercial and an operational question.

Agent-initiated transactions arrive without the card number, device fingerprint, and behavioural history that human fraud models rely on. Carried over unchanged, those models generate false positives on legitimate agent activity and miss fraud patterns specific to agent flows. The mandate verification opportunity is where PSPs are positioned to contribute most to the broader accountability architecture. A PSP that can read and verify an agent's mandate at the point of transaction provides a service that neither the merchant nor the bank can provide from their respective positions in the chain.

> *"The transaction data that agent-initiated payments produce looks different from human-initiated flows, and risk models calibrated for human commerce will generate false signals when applied without adaptation."*
>
> **— Chris Maurice, Co-Founder and CEO at Yellow Card**

A transaction that passes the PSP's fraud checks may still be executing outside the scope of the delegation chain above it, and when settlement involves the redemption of a programmable instrument, the issuer carries the compliance exposure at that point.

> *"When a delegation chain is compromised, the resulting transactions can be indistinguishable from legitimate agent activity at the processing layer. The risk sits upstream in the authorisation chain, beyond where existing payment fraud detection operates."*
>
> **— Axel Cateland, CEO at Kulipa**

Section 5: Issuers - Redemption Integrity at Par

Issuers of programmable settlement instruments, including stablecoins, tokenised deposits, and single-purpose tokens (instruments issued for a defined commercial purpose or counterparty set), face a version of the KYA problem specific to their role in maintaining the credibility of the instruments agents use to settle transactions. The par value commitment an issuer makes depends on the integrity of the redemption chain: the agent presenting an instrument for redemption must be authorised to do so by a principal with the standing to make that instruction.

The compliance exposure from unverified agent redemptions sits alongside the financial one. An agent executing redemptions at scale across multiple jurisdictions, without a principal chain that makes the originating instruction attributable, creates AML risk at the precise point where settlement value is released.

> *"As agent-initiated transactions grow, issuers still need to confirm who is presenting an instrument and whether they're authorised to do so. Without a standard credential to verify that, the obligation can't be reliably met at scale. Closing that gap is what separates trusted agentic commerce from agentic risk."*
>
> **— Markus Infanger, Senior Vice President at RippleX**

Issuers building redemption infrastructure that incorporates mandate verification and principal chain visibility are developing the compliance architecture that institutional-scale agentic commerce requires, and the credential standards they establish for what an agent must carry to redeem at par will set the baseline for the broader settlement network.

An issuer can only confirm a redemption is authorised if the chain it settles on can show who instructed it and on whose authority, which ties the issuer's assurance directly to the chain beneath it.

Section 6: Chains - Settlement Infrastructure and the Compliance Ceiling

For human-initiated stablecoin transactions, the pseudonymous nature of wallet addresses at the chain layer is resolved off-chain by the exchange or institution that onboarded the user, and the link between wallet address and legal identity sits in a regulated institution's records. For agent transactions, the wallet address may belong to a business entity, a developer, a deployment platform, or some combination of these, and the path from that address to the ultimate human or corporate principal is invisible from chain-level data alone.

Regulated financial institutions settling on a given chain require that chain to support travel rule compliance, sanctions screening, and an on-chain record equivalent to the KYC documentation that AML frameworks require them to maintain.

A chain's technical capabilities, including throughput, finality, and settlement cost, determine the floor of what commercial activity it can support, while its compliance infrastructure determines the ceiling. The compliance ceiling already determines which chains regulated institutions can use for settlement at institutional scale, and it becomes more visible as agent transaction volumes grow.

> *"Travel rule compliance and sanctions screening at the chain layer are prerequisites for institutional settlement volume, and chains that build verifiable credential portability into their architecture are positioned to carry the commercial volume as agent commerce scales."*
>
> **— Alex Gluchowski, CEO at Matter Labs**

The KYA principal chain and audit log address both aspects of the chain-layer compliance requirement: a transaction carrying a verifiable principal chain gives the settling institution and the regulator the originator and beneficiary information that travel rule compliance requires, and the audit log provides the attributable record that AML documentation frameworks require institutions to maintain.

> *"Which regulated institutions can settle on a chain is determined by its compliance infrastructure. For agentic commerce, that infrastructure needs to include the ability to confirm the authority behind each transaction."*
>
> **— Bill Papp, CEO at NVNM Chain**

That compliance data has to originate somewhere, and the wallet is where the agent's credentials, mandate, and principal chain are held, making it the layer at which chain-level compliance requirements must be met before settlement.

Section 7: Wallets - The Mandate Representation Problem

The wallet is positioned as the agent's payment interface, making it the natural layer at which to also carry, in a format that any counterparty can read and verify in real time, the three credentials the commercial chain requires at the point of transaction: who deployed this agent, what it is authorised to do, and against whom a claim can be made if it acts outside that authority.

The wallet is also the layer at which the agent's mandate is enforced at the point of transaction, meaning that a transaction falling outside the agent's authorised scope is blocked before execution, and where each transaction generates the verifiable record that compliance and regulatory frameworks require. A wallet designed to carry credentials and a wallet designed to enforce them serve different purposes, and the KYA standard requires infrastructure capable of doing both.

The protocols that have emerged address this at different layers: Visa's Trusted Agent Protocol (TAP) at the trust layer, Visa Intelligent Commerce and Mastercard Agent Pay at mandate authentication, and Google's Agent Payments Protocol and OpenAI's Agentic Commerce Protocol across instruments and rails, each addressing a distinct part of the design problem. Whether mandate credentials built within a single network can be made readable by all counterparties remains the open question, and universal reach requires a cross-network standard above the individual protocol layer.

> *"Establishing the accountable legal entity behind an AI agent and establishing what that agent is authorised to do are distinct, but connected, trust problems. Indeed, entity-level identity verification provides the root of trust on which mandate verification depends. The two should not be collapsed into a single credential, but they should be logically layered."*
>
> **— Alexandre Kech, CEO at GLEIF**

The substantive design challenge is that mandates are dynamic: an agent's authorised scope changes across transaction types, time windows, and counterparty categories. The credential standard must accommodate scope that is conditional, time-bounded, and counterparty-specific, expressed in a format verifiable at transaction speed without exposing the full details of the principal's commercial arrangements to parties not required to know them.

> *"The mandate an agent carries must be selective in what it discloses. The credential design challenge is building something verifiable and enforceable that also protects the principal's commercial arrangements from unnecessary exposure."*
>
> **— Ariel Madjar, Product Manager at Utila**

The design requirements that emerge from every layer of the commercial chain coalesce in the four agreements that follow.

Section 8: The Four Agreements and the Network That Makes Them Work

The card networks solved an equivalent coordination problem for human commerce by establishing the shared standards governing what a card carries, what an authorisation signal means, and how liability is allocated when something goes wrong, making it possible for thousands of institutions across dozens of jurisdictions to transact with confidence. The KYA equivalent requires four things to be agreed, and a mutualised network operating across issuers, chains, and jurisdictions to give those agreements commercial reach.

The first is a mandate representation standard, specifying what an agent's wallet must carry and what any counterparty must be able to verify from it at the point of transaction. An agent's mandate is only as useful as the number of participants who can read and act on it.

The second is an obligation allocation framework, specifying which participant carries the AML monitoring obligation, the reporting obligation, and the liability for transactions that fall outside the agent's mandate in each jurisdiction.

> *"KYC was designed to identify a person and the business they do. It has no concept of a non-human acting on that person's behalf. The question agentic commerce puts to banks is how we know the agent in front of us genuinely represents the customer we onboarded."*
>
> **— Rob Downes, Head of Digital Assets at Absa Corporate and Investment Banking**

The third is agent-aware monitoring standards, providing benchmarks for what normal agent transaction behaviour looks like across the commercial contexts in which agents operate at volume.

The fourth is a liability model for agent error, defining in advance which party bears the loss when an agent executes outside its mandate, through what mechanism, and on what timeline. The card network chargeback model is the structural precedent. The specific design requires coordination among participants currently operating without that legal certainty.

Figure 2: The Four Agreements Framework

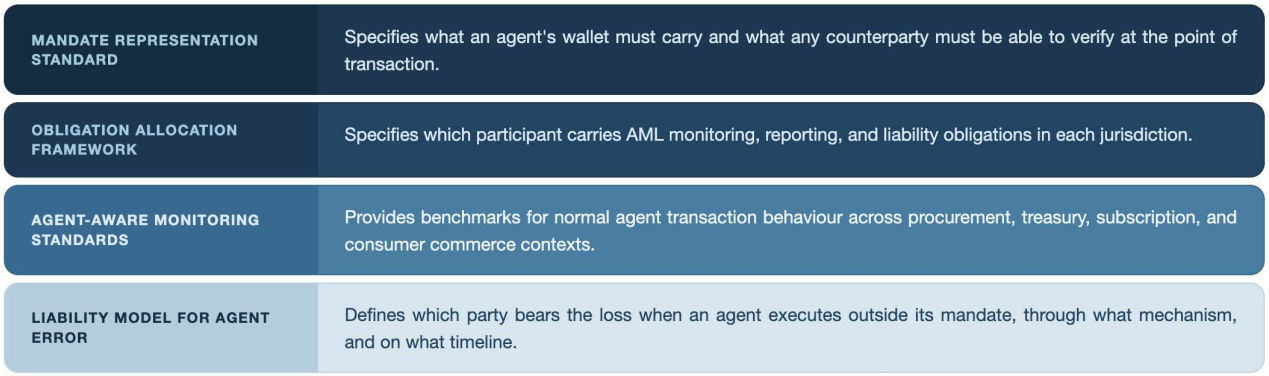

In wholesale payments, the institutional identity layer already exists: Bank Identifier Codes (BICs) identify institutions on financial messaging networks, and Legal Entity Identifiers (LEIs) identify legal entities across markets. Agent credentials build on that layer. Existing identity infrastructure establishes who the institution and the legal entity are; what it does not yet express is what an agent acting for that entity is authorised to do, under what conditions, and against whom a claim lies when it acts outside that authority. That is the gap the KYA framework addresses, and closing it requires new infrastructure that no existing network supplies by default.

Practitioner research into enterprise procurement has approached the same authorisation verification gap from a different architectural angle, proposing shared registry infrastructure for verifying that an agent's delegation chain is intact at the point of transaction. A credential standard and a registry are complementary. The credential standard sets out what an agent presents and what any counterparty can verify. A registry provides the shared infrastructure against which those credentials are checked at scale.

> *"Settlement infrastructure and identity infrastructure are addressing different parts of the same gap. Defining what an agent must carry and maintaining the infrastructure against which those credentials can be checked are complementary requirements, and both need to be in place before volume arrives."*
>
> **— Tianwei Liu, CEO and Co-Founder at StraitsX**

Reaching those agreements requires trust and confidence across every participant in the commercial chain, and that is a responsibility no single organisation can carry alone.

Section 9: The Architecture and the Decisions It Requires

The framework this paper has set out operates at three levels. The four KYA components in Section 1 (identity credential, mandate certificate, principal chain, and audit log) are the credentials a counterparty needs to verify before transacting with an agent. The four agreements in Section 8 are what institutions need to settle between themselves for those credentials to function at the point of transaction. This section describes the architecture that puts both into operation, and the decisions that are required from each participant.

Agentic commerce is underway, and scaling it requires specific steps from every participant this paper has described.

One is a trust principles framework: agreed principles governing what a compliant agent credential must contain, what constitutes a valid delegation chain, and how disputes are attributed when an agent exceeds its mandate. Developing that framework, and operating the infrastructure that makes it verifiable across every participant, is what would convert it from a document into a functioning commercial standard.

Another is regulatory sandbox testing, which gives institutions the regulatory cover and evidence base to operate agreed principles at scale. National sandboxes and cross-border forums such as the Global Financial Innovation Network (GFIN) provide the coordination mechanism for that testing across markets.

> *"A standard only helps an agent if it is recognised in every market the agent reaches. Coordinating regulators across borders is how principles agreed in one jurisdiction come to be trusted in the next, so an agent does not meet a different rulebook at every crossing."*
>
> **— Colin Payne, Head of Innovation at the Financial Conduct Authority (FCA) and Chair of the Global Financial Innovation Network (GFIN)**

For issuers, the practical step is building redemption infrastructure that can read and verify an agent's mandate and principal chain at the point of settlement, designed to interoperate across a shared network rather than in isolation. For PSPs, it is developing the capability to read and verify an agent's mandate at transaction speed, a verification the rest of the chain can then rely on. For merchants and large marketplaces, establishing policies for credentialed agent transactions, including commercial terms and dispute resolution processes, opens agent-initiated volume as a managed commercial channel.

Banks, PSPs, and regulators need to agree on which participant in the agent transaction chain carries the AML monitoring obligation before agent transaction volumes make that question consequential. Standards bodies and card networks need to align on authentication and authorisation protocols before competing proprietary approaches become entrenched across markets.

Each participant's steps call for network infrastructure capable of operating credential verification, clearing, and obligation allocation across issuers, chains, and jurisdictions. The institutions that build and operate that layer, with reach across markets and no commercial dependency on any single issuer or chain, are positioned to give the four agreements the scale at which they function. Whether agentic commerce develops as a coherent commercial system or across competing proprietary approaches will depend on whether that infrastructure exists before transaction volumes reach commercial scale.

The KYA framework, alongside the settlement and governance infrastructure the earlier papers established, gives agentic commerce the foundation it needs to be deployed safely, responsibly, and with confidence across the commercial chain.

> *“A platform processing agent-initiated transactions at scale needs to know what the agent was authorised to do and who bears responsibility when it acts outside that authority. Agreed standards for establishing those facts are what make large-scale agent commerce operationally viable.”*
>
> **— Amira Karim, Head of Payments and Financial Services Public Policy - International Public Policy at Amazon Consumer**

Agentic commerce will be built by the institutions that act on these decisions first, and the standards they establish will set the terms for all who join the network.