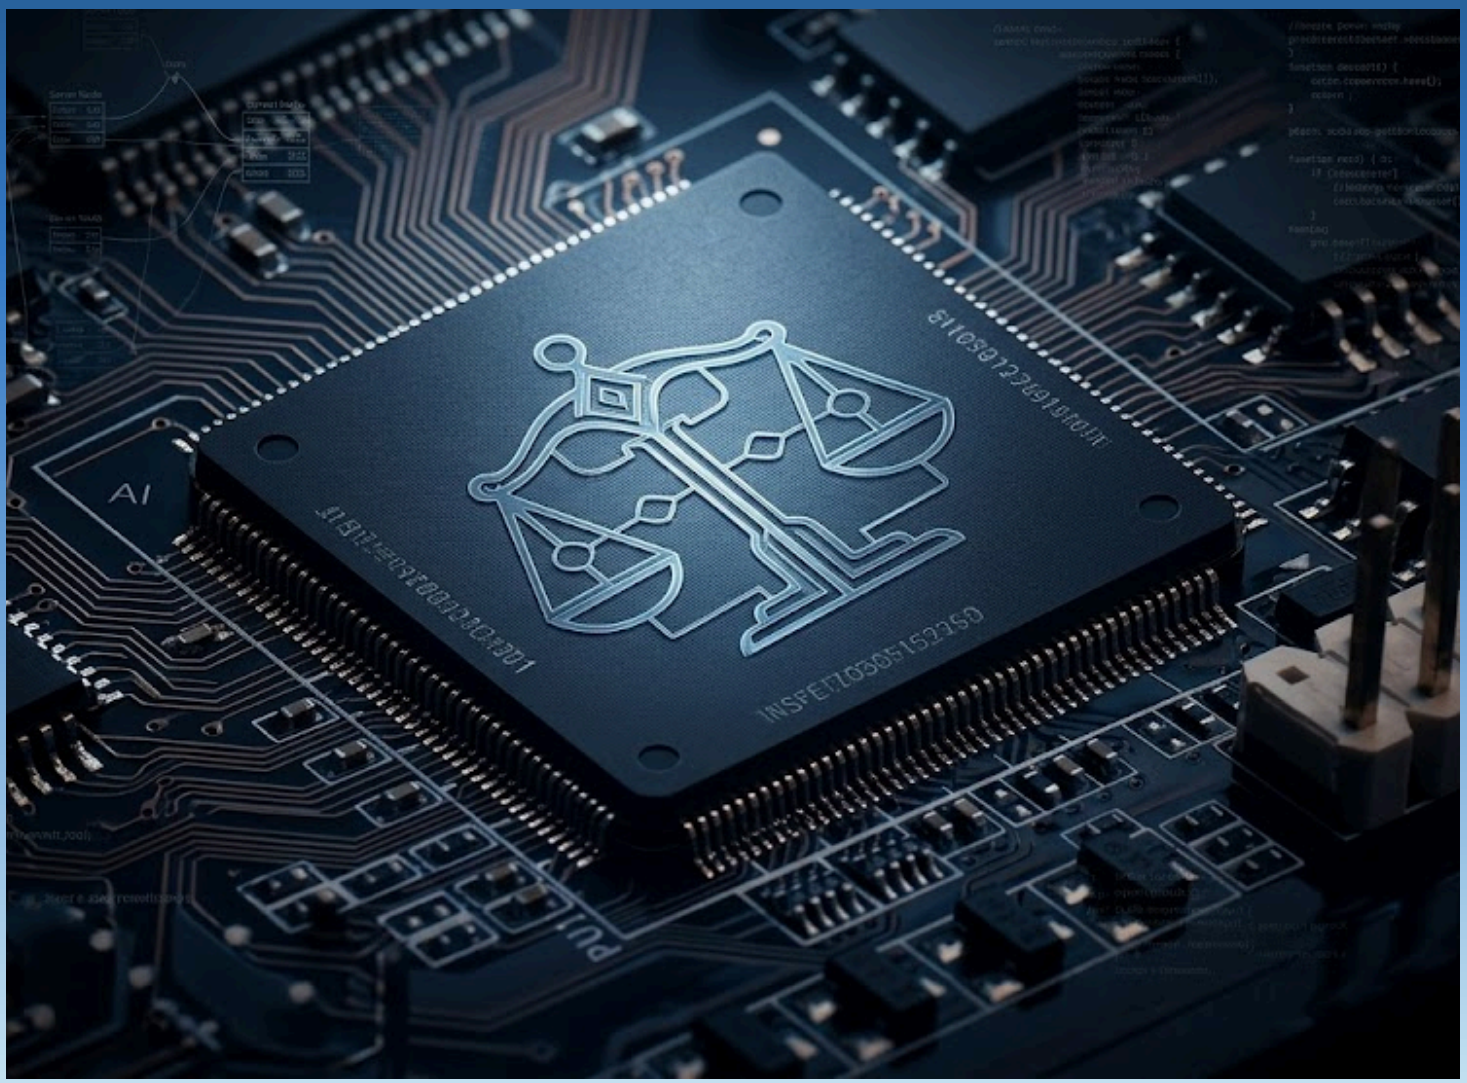

Commerce at Machine Speed, Part II

The Governance, Accountability, and Regulatory Framework Agentic Commerce Requires

Authors

Nicole Sandler, Chief Ecosystem Officer at Ubyx Inc.

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

Key Contributors

Commerce at Machine Speed, Part II

The Governance, Accountability, and Regulatory Framework Agentic Commerce Requires

Authors: **Nicole Sandler**, Chief Ecosystem Officer at Ubyx Inc.
Colin Payne, Head of Innovation at the Financial Conduct Authority (FCA) and Chair of the Global Financial Innovation Network (GFIN)

Key Contributors: A&O Sherman, AllUnity, Amazon, Ava Labs Inc., Emerging Payments Association Asia (EPAA), European Commission, Fireblocks, Hedera, Latham & Watkins LLP, Quantoz, Standard Chartered, Thunes

This paper examines the governance, liability, and regulatory frameworks that agentic commerce requires, and is the second in a three-paper series. It is not an endorsement of any specific payment instrument, issuer, or platform.

Section 1: From Requirements to Rules - Creating the Regulatory Framework for Agentic Commerce

When a human authorises a payment, the accountability chain is direct: a person approved the transaction, their identity is on record, and liability can be attributed to them if something goes wrong. Agentic commerce removes that person from the point of execution, replacing individual transaction approval with mandate design: the parameters within which an agent is authorised to act, defined before deployment and without human review at each transaction. Governance relocates upstream, and the questions that relocation creates are what this paper addresses: how accountability is attributed when the transacting party is a machine, how liability is allocated when an agent acts outside its mandate, and how regulators across jurisdictions are building the frameworks to answer those questions.

The first paper in this series established the infrastructure that agentic commerce requires: on-chain settlement with smart contract finality, a global mutualised acceptance network operating across issuers, chains, and jurisdictions, and a wallet architecture that does not require a human account holder. The governance frameworks this paper examines apply wherever a non-human actor executes transactions on behalf of a principal without real-time human oversight, and are designed to remain relevant as settlement infrastructure continues to evolve; tokenised deposits and central bank digital currencies (CBDCs) developing to meet the same requirements will qualify for the same framework.

The infrastructure question and the governance question are running in parallel, and each requires the other to be complete. The regulatory and trust frameworks this paper examines are a key condition for the confidence that commercial participants need to engage with agentic commerce at scale.

Section 2: Governing a System Built for Agents

A mandate is the governance instrument that agentic commerce requires: it must define what an agent is authorised to do, under what conditions, and with what constraints, in advance of deployment and without the possibility of human review at each transaction. The terms under which an agent is permitted to act, the data inputs that trigger settlement, the spending limits and counterparty constraints within which it operates all require rigorous design and auditing before deployment. Consent moves with governance, operating at the point of mandate design.

> *"Mandate design is a legal and contractual exercise as much as a technical one. The agent's scope of authority, conditions under which it can act, and liability that attaches when acting outside those conditions all raise novel questions about regulatory compliance and accountability that will percolate through the agentic commerce value chain."*
>
> **– Shruti Ajitsaria, Partner and Head of Fuse at A&O Sherman**

Value stability is a design requirement of any clearing infrastructure serving agentic commerce. Par value redemption, the settlement of instruments at their declared face value, is what makes any tokenised money instrument viable for commercial use, and clearing infrastructure enforcing that standard across issuers is what makes it consistent in practice.

> *"Maintaining par value confidence requires more than reserves. The redemption infrastructure, the audit standards, and the speed at which a holder can redeem at face value also matter in determining whether an instrument is trusted for commercial settlement."*
>
> **– Alexander Höeptner, CEO at AllUnity**

The same concentration of governance at the design stage applies to fraud risk: an agent operating at scale can execute a compromised instruction across thousands of transactions before any human reviews it. Rigorous governance at the design stage provides the answer: the parameters within which agents operate are defined at contract design, with spending limits, counterparty whitelists, velocity caps (transaction frequency limits), and settlement conditions set and audited before deployment. Every transaction is recorded on an immutable, timestamped ledger, producing an audit trail that correspondent banking chains cannot replicate. The fraud risk in stablecoin clearing concentrates at the design stage, and the governance response is pre-deployment audit standards and mandate controls.

Privacy obligations run across the same chain. The behavioural data and financial information that agents generate are subject to data protection frameworks that vary by jurisdiction, few of which were designed for automated, multi-party transaction flows. A well-designed authorisation contract can confirm sufficient funds without disclosing the account balance, verify identity without exposing personal details, and complete a transaction without revealing its terms to parties not required to know them. The result is a clearing infrastructure that carries governance logic and selective disclosure (the capacity to reveal only what a counterparty needs to verify) as properties of the transaction itself, a capability that conventional payment rails, passing data through intermediary chains, cannot replicate in the same way.

Section 3: The First Mover Advantage

Banks, fintechs, payment providers, and wallet providers are the institutional layer for what comes next, and the institutions that engage with tokenised money clearing infrastructure early are positioned within the settlement layer as agentic commerce grows. The pattern is familiar from every major payment network that preceded this one: early participants set the standards, the architecture, and the terms of access, and those who followed integrated on the terms that had already been established.

The clearing layer for agentic commerce will develop with or without early engagement from regulated institutions; the question is whether they shape how the compliance, liability, and interoperability standards are designed.

Banks and payment providers face a specific challenge that existing governance frameworks have not yet resolved: how to handle anti-money laundering (AML) and know your customer (KYC) obligations when the initiating party in a transaction chain is an autonomous agent. Existing compliance frameworks assume a human customer whose identity can be verified and whose transaction behaviour can be monitored against known patterns. Agent-initiated flows present different characteristics: higher frequency, lower individual values, and conditionality that varies significantly by use case. The institutions building compliance infrastructure for agent-initiated flows now are generating the operational evidence from which industry standards will be made.

> *“Agent-initiated transactions challenge foundational assumptions around identity and authentication. Verifying a counterparty with no legal standing requires a different approach and, ultimately, new frameworks for regulation, liability and trust built specifically for machine-to-machine interactions.”*
>
> **– René Michau, Global Head of Digital Assets at Standard Chartered**

> *"The unique challenge of agentic flows goes beyond the complexity of the counterparty chain - it's also about the fragmentation of the response. We believe the way to solve for agentic friction is through an interoperable infrastructure layer that translates local AML and settlement standards into a single, machine-readable protocol. We are engineering the network which can ensure that autonomous triggers, no matter how high the transaction velocity is, always land within the safety of regulated financial rails."*
>
> **– Elie Bertha, Chief Product Officer at Thunes**

Wallet providers hold identity credentials, authorisation mandates, and payment instruments for the agents that initiate transactions. The wallet layer presents what matters at the point of transaction: the mandate an agent carries, what it is authorised to execute, and the identity signal it presents to merchant and payment provider at settlement. How mandates are scoped, how identity credentials are structured, and how authorisation is expressed are design questions without consistent industry answers. The wallet providers building for agent commerce now are establishing those standards in practice, ahead of any formal regulatory specification.

Chains are the settlement infrastructure on which stablecoin transactions run, and institutional settlement places specific demands on that layer: compliance architecture that regulated institutions can work within, reliability standards that hold at transaction volume, and interoperability with the clearing infrastructure above. The chains that meet the compliance and reliability requirements that institutional participants demand will carry the commercial volume as it grows.

Merchants face the most immediate and concrete engagement decision. Agentic systems will interact with merchant infrastructure in one of two ways: through browser-based navigation, where agents behave as a human user would, or through direct application programming interface (API) connections, where structured links connect agents directly to merchant and payment systems. In an environment where agents optimise across available merchants for price, availability, and transaction success rate, the connection architecture becomes a direct commercial variable. Large-scale marketplaces already operate sophisticated trust infrastructure including escrow mechanisms, automated refunds, and structured dispute resolution; on-chain compliance infrastructure provides a parallel protection layer at the settlement level, and the governance framework for agentic commerce must define how these interact when an agent-initiated transaction is disputed.

Without a global acceptance layer there is a risk that agentic commerce concentrates volume toward merchants with the infrastructure to serve automated counterparties, excluding smaller participants in ways that card networks were designed to prevent.

> *"Retail checkout experiences were built to authenticate a customer completing a purchase. When an AI agent acts on a customer's behalf, the question evolves: not only who is this customer, but what has this agent been authorised to do, and can it be trusted to act on accurate, up-to-date information?"*
>
> **– Amira Karim, Head of Payments and Financial Services Public Policy - International Public Policy, Amazon Consumer**

Section 4: Trust Architecture - Introducing Know Your Agent (KYA)

Agentic commerce operates across five connected layers: the merchant layer, through which agents interface with product catalogues, pricing systems, inventory, and checkout flows; the artificial intelligence (AI) agent initiating transactions; the authorisation layer governing what it is permitted to do and on whose behalf; the programmability layer embedding conditional logic in the transaction; and the settlement layer completing it. The authorisation layer is the design problem the trust architecture for agentic commerce must address.

Figure 1: The five-layer architecture of agentic commerce

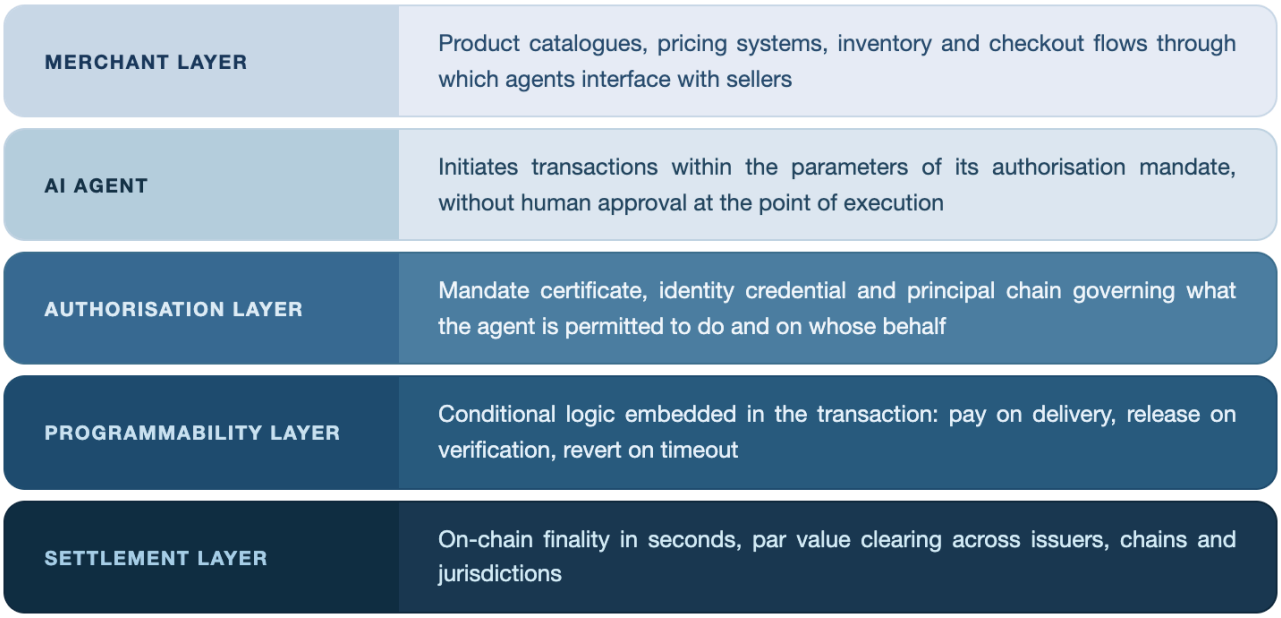

Every current payment framework places a human at the point of authorisation: a person who can be identified, whose consent is on record, and against whom liability can be assessed when something goes wrong. Agentic commerce moves that person upstream. Consent and authorisation shift to mandate design: the parameters within which the agent operates are defined and validated before deployment. At the point of transaction execution, no human is present to verify, override, or be held immediately accountable. No jurisdiction currently recognises an AI agent as a legal person: when an agent executes a transaction, the legal act must be attributed to a human or corporate principal, but the delegation chain can be long, technically complex, and opaque to any party outside it. When a transaction is executed outside its mandate or a smart contract triggers on a contested condition, accountability disperses across agent developers, platforms, payment providers, and merchants in ways that existing dispute frameworks were not built to resolve.

Figure 2: The accountability chain in agentic commerce

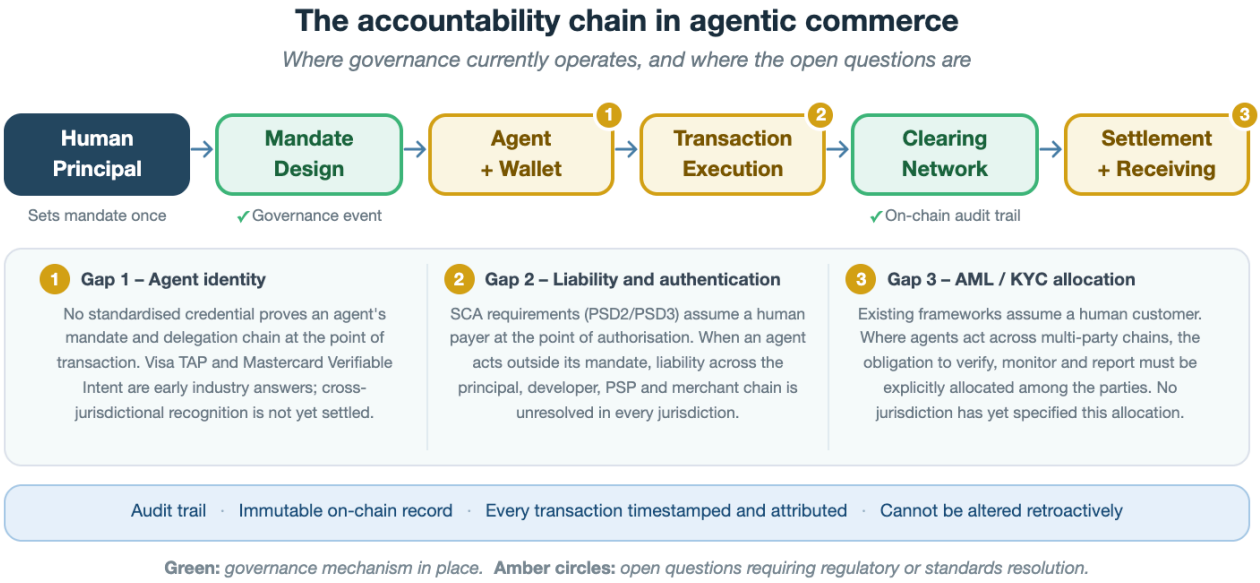

The existing common law of principal-agent relationships provides the legal foundation from which that accountability gap should be read. Three concepts are already embedded in current payment frameworks: actual authority, covering the express permissions granted at mandate design and the reasonable actions implied by them; implied authority, governing the discretion an agent exercises in carrying out those permissions; and apparent authority, which binds a consumer to transactions executed by an agent they have authorised until that authorisation is affirmatively revoked. In the United States, Regulation Z, which implements the Truth in Lending Act governing consumer credit transactions, already incorporates apparent authority in its definition of unauthorised use. Agentic commerce on stablecoin rails inherits that doctrine and requires the technical infrastructure to express, transmit, and verify authority in machine-readable form across a settlement layer that those frameworks did not anticipate.

> *"The mandate is where governance lives in agentic commerce. Getting the credential structure and scope right at the design stage prevents unauthorised activity. It creates the foundation for counterparties to verify the agent's authority and attribute accountability when something falls outside it."*
>
> **– Varun Paul, Senior Director for Financial Markets at Fireblocks**

For policymakers and regulators, resolving the accountability gap requires decisions that industry cannot deliver on its own. The legal personhood of AI agents, the cross-border allocation of AML and KYC obligations, and the consent frameworks that apply when an agent executes across multiple jurisdictions all require regulatory decisions. The jurisdictions and industry working groups engaging with those matters now, through coordinated pilots and international initiatives such as the Global Financial Innovation Network (GFIN), are where those frameworks will form.

> *"The liability question in agentic commerce is about how the chain from consumer instruction to agent action to transaction execution is documented, audited, and attributed. When one party is a machine operating under a delegated mandate, existing dispute resolution frameworks require re-examination at every stage."*
>
> **– Stuart Davis, Partner and Global Co-Chair of Fintech at Latham & Watkins LLP**

Transparency in agentic commerce requires that the principal can reconstruct the agent's decision chain after the fact: which merchant was selected, on what criteria, what terms were accepted, and whether the transaction fell within the scope of the original mandate.

One emerging approach draws on the architecture of KYC compliance. Where KYC establishes the identity and standing of a human or corporate counterparty, a Know Your Agent (KYA) framework applies the same logic to the AI agent itself: what it is, what it is authorised to do, on whose behalf it acts, and what record exists of its actions. A KYA framework would require four elements for each agent executing financial transactions: an identity credential establishing the agent as a registered entity within a defined technical and governance framework; a mandate certificate specifying the scope of authority delegated to it; a principal chain tracing the delegation from the originating human or corporate principal through to the acting agent; and an audit log capturing each transaction in a form attributable to the agent, the mandate, and the principal. The identity and mandate standards that KYA requires are the subject of the third paper in this series.

Governance provenance infrastructure is the complementary layer: making that authorisation portable and verifiable as the transaction moves across chains, networks, and jurisdictions, so that the proof of permission travels with the asset and remains verifiable at every point in the transaction chain.

> *"The architectural requirement that machine-driven commerce adds to the settlement layer is provable judgment: demonstrating that a transaction was authorised to execute under defined policy, in a form portable enough to travel with the asset and auditable enough to attribute liability when something falls outside it."*
>
> **— Tom Zschach, Former Chief Innovation Officer at Swift**

The layers described above are structurally interdependent rather than sequential. Without identity, liability cannot be allocated. Without interoperable tokenisation infrastructure, identity cannot be authenticated. Without accurate data access at the merchant layer, mandate compliance cannot be assessed. Governance frameworks that treat these as separable workstreams will find themselves building accountability structures on incomplete foundations.

Section 5: Regulatory Landscape - Frameworks in Formation

The accountability questions Section 4 raises are live regulatory questions. Regulators in multiple jurisdictions are working through them, and the frameworks that emerge will set the practical operating conditions for agentic commerce. Those frameworks are developing on different timelines, within different regulatory traditions, and against different assumptions about what agentic commerce is and how it works.

The legal foundations for agent accountability in commercial transactions are older than the current regulatory debate suggests. The US Electronic Signatures in Global and National Commerce Act (ESIGN) and the Uniform Electronic Transactions Act (UETA), both enacted in 1999-2000, already contemplated autonomous electronic agents capable of completing transactions without human review, establishing a three-part structure of agency relationship, reliable attribution of acts to a principal, and defined security procedures that anticipate the KYA framework in its essentials.

In the United Kingdom, the Financial Conduct Authority's (FCA's) Mills Review is examining how agentic AI intersects with existing financial regulation, with specific attention to delegated authority and authentication.

The FCA's AI Lab is expected to produce good and poor practice guidance during 2026, giving financial institutions practical direction ahead of any formal framework. The FCA's AI Compass provides a self-assessment tool for institutions mapping their AI deployments against existing regulatory requirements in the interim. At the government level, HM Treasury's Payments Forward Plan, published in February 2026, identifies agentic AI as a specific focus for the modernisation of payment services regulation, and the consultation it commits to on stablecoin payments and agent-initiated transactions is a government-level signal that the payment services architecture itself requires updating.

> *"Where obligations are shared across every participant in the agent transaction chain, accountability has to be made explicit at each point. Regulatory frameworks that do not specify that allocation will find institutions defaulting to the narrowest interpretation of their own obligations."*
>
> **– Camilla Bullock, CEO at the Emerging Payments Association Asia**

In Singapore, the Monetary Authority of Singapore's (MAS's) draft AI Risk Management Guidelines address autonomous transaction execution directly. The MAS Single-Currency Stablecoin framework, finalised in August 2023, requires one-to-one reserves and redemption at par within five business days for Singapore dollar and G10-pegged stablecoins. Singapore's shared responsibility model allocates compliance obligations explicitly among developers, deployers, financial institutions, and end users and offers one approach to how liability in agentic commerce might be structured across the transaction chain.

> *"Agentic commerce systems will operate across regulatory borders as a matter of design, and the accountability and governance frameworks that apply to them need to reflect that from the outset."*
>
> **– Peter Kerstens, Adviser on Technological Innovation at DG FISMA, European Commission**

In the European Union (EU), the Markets in Crypto-Assets Regulation (MiCA) has been in force for stablecoins since June 2024, establishing issuance, reserve, and operational requirements for euro-denominated stablecoins. The EU AI Act establishes a risk-based framework for AI systems operating in financial services contexts, and eIDAS 2.0, the EU's digital identity framework, provides the infrastructure through which agent authorisation mandates could be anchored.

> *"MiCA creates a governance baseline across regulated stablecoin issuers that is audited rather than asserted. As agentic commerce drives automated settlement volumes through these instruments, that regulatory foundation is what institutional confidence in multi-issuer settlement depends on."*
>
> **– Arnoud Star Busmann, CEO at Quantoz**

In the United States (US), the Guiding and Establishing National Innovation for US Stablecoins (GENIUS) Act, signed into law in July 2025, establishes a federal framework for stablecoin issuance requiring full reserve backing, monthly attestation, and licensing requirements for issuers. The White House AI Policy Framework, published in March 2026, is largely silent on agentic AI in financial services, reflecting the administration's position that industry standards should develop ahead of federal regulation.

The Four Structural Gaps

Each of these frameworks addresses part of the problem, and none yet addresses it in full. Four structural gaps run across all of them.

The first is agent identity and delegated authority. No jurisdiction has legally defined agentic AI in the context of payment services, and no authentication framework currently specifies how an AI agent proves its delegated authority at the point of transaction. Strong Customer Authentication (SCA) requirements under the Payment Services Directive assume a human payer actively authorising at the point of transaction, and how agent-initiated payments satisfy that requirement remains an open question. Industry protocols are beginning to address different aspects of this gap at different layers of the transaction stack. Visa's Trusted Agent Protocol (TAP) operates at the trust layer, establishing a network of known, whitelisted, and identified agents that counterparties can rely on before mandate verification takes place. Visa Intelligent Commerce (VIC) and Mastercard Agent Pay address the mandate authentication layer, synchronising agent mandates through each network's authorisation infrastructure to verify that the mandate a merchant receives matches the one the consumer originally authorised. Mastercard's Verifiable Intent standard, co-developed with Google as an open standard, provides the cryptographic proof of that authorisation within the framework. Both VIC and Mastercard Agent Pay work within their respective rails. Google's Agent Payments Protocol (AP2) and OpenAI's Agentic Commerce Protocol (ACP) define how an agent communicates what it intends to pay for and on whose authority, and are payment method agnostic, operating across instruments and rails rather than within a single network. The regulatory question these protocols collectively raise is whether solutions operating within individual networks can close the authentication gap, or whether the gap requires a standard that functions consistently above the rail, across issuers and jurisdictions simultaneously.

The distinction that current frameworks have not yet closed is between cryptographic validity and policy validity. Cryptographic verification can confirm that a message was signed by an authorised key. It cannot confirm that the underlying action was permitted under defined policy, by a party with the standing to permit it, at the time it was executed. Closing that gap is the precise requirement that agent identity standards must meet.

> *"Agentic payments require a verifiable record linking user intent to transaction outcome. Cryptographic proof of delegated authority is what makes that record auditable at scale. For regulated institutions operating under transaction monitoring obligations, that auditability is how any meaningful volume moves through agentic channels."*
>
> **– Nilmini Rubin, Chief Policy Officer at Hedera**

The second is liability allocation. When an agent executes a transaction outside its mandate, liability must be distributed across a consumer, AI developer, payment provider, and merchant. Existing chargeback and dispute resolution frameworks were designed for human error and transaction disputes, and how they apply to agent error or misaligned optimisation is an active question for regulators in each jurisdiction.

For card-based agentic commerce in the United States, the existing Regulation Z framework provides a workable starting point: a consumer is bound by transactions an agent executes within the authority they have created the appearance of granting, and the issuer bears liability for transactions executed without any such authority until revocation is confirmed across the network. That framework does not extend to on-chain stablecoin settlement, where transactions reach finality in seconds, chargeback rights have no structural equivalent, and the Regulation Z and Regulation E coverage underpinning card dispute resolution applies at best ambiguously. Liability allocation for stablecoin-settled agentic commerce therefore requires deliberate construction, drawing on card and ACH doctrine where it applies and filling the gaps where it does not.

The third is AML and KYC in multi-party agent chains. Where agents act across multi-party chains, the obligation to verify identity, monitor transactions, and report suspicious activity must be allocated explicitly among the user, AI developer, payment provider, and merchant. Current frameworks do not specify that allocation for non-human principals, and the absence of a standard creates compliance uncertainty for every institution building on stablecoin clearing infrastructure. Singapore's shared responsibility model is one approach; it has not yet been adopted or adapted elsewhere.

Cutting across all three gaps is a fourth challenge: protocol interoperability. An agent operating across jurisdictions faces different technical standards for authentication, authorisation signalling, and settlement. The credential chain proving delegated authority in one jurisdiction may not be recognised in another, and settlement finality achieved on one architecture may not be recognised as final by a receiving institution on a different one.

The same fragmentation extends to payment credential tokenisation. The mechanism protecting consumer payment details when an agent transacts on their behalf is developing along proprietary, non-interoperable lines: Visa's AI-ready tokenised credentials and Mastercard's Agentic Tokens each operate within their own network rails, with no cross-network standard in place. PCI DSS v4.0.1, the most relevant existing compliance framework for entities handling card data, was not designed for agentic transaction flows and does not address non-human identity governance or agent-specific privilege management under delegated payment authority.

Resolving this requires coordinated cross-border engagement on technical standards.

> *"Commercial settlement at institutional scale requires more than throughput. What determines whether regulated volume moves is the combination of compliance architecture, verifiable authority, and legally reliable finality."*
>
> **– Lee A. Schneider, General Counsel at Ava Labs, Inc.**

Section 6: Infrastructure, Governance, and the Standard Ahead

The governance gaps this paper has examined are live questions, already subject to regulatory engagement, standards work, and institutional decision-making: the legal status of agents in commercial transactions, the allocation of AML obligations across multi-party chains, the consent frameworks that apply when an agent transacts across multiple jurisdictions, and the liability model for agent error are each in active consideration across the jurisdictions and working groups this paper has reviewed. The clearing infrastructure and the governance framework are interdependent, and the institutions building both together are establishing the operating conditions under which institutional-scale agentic commerce becomes possible.

The standards that emerge from these deliberations will determine what the compliance architecture for agentic commerce looks like, how liability is distributed across the transaction chain, and which participants can operate within it with confidence at scale. The identity and credential standards that complete that architecture are the subject of the third paper in this series.