CLARC:

The *Missing Infrastructure Layer* for Enterprise Agentic Payments

A Clearing and Authorization Registry for Agentic Commerce

Eight live coding experiments. 45 documented infrastructure gaps. One consistent finding: AI agents can initiate enterprise payments today, and no governance layer exists to verify they were authorized to do so. This paper documents the evidence, explains why bilateral solutions cannot close the gap, and presents the architecture of the neutral clearing layer that can.

Version 1.0, April 2026

Sandra Lam
MIT SLOAN FELLOWS MBA
PAYMENTS EXECUTIVE
INDEPENDENT RESEARCH
PATENT PENDING
CONTACT: SANDRA@THEAGENTECONOMY.CO

Foreword

This paper is written for four audiences: Enterprise Leaders, Financial Institutions, FinTechs and Regulators.

The first audience is enterprise leadership — CFOs, CPOs, heads of treasury, and the AI deployment teams now building procurement agents and accounts payable automation. They are deploying AI systems that commit financial obligations on behalf of their organizations. This paper documents what their governance infrastructure currently cannot verify, and what it will need to verify as agent volumes scale.

The second audience is financial institutions — transaction banking leaders, payments infrastructure teams, compliance and BSA officers, and the product teams building the next generation of corporate banking services. They process agent-initiated payments today, often without knowing it. This paper explains the structural gap in what they can currently verify about those instructions, and what shared infrastructure would change.

A third audience is the FinTechs building the platforms, protocols, and agent frameworks through which agentic commerce will operate. They are building the infrastructure that will generate the payment flows this paper addresses. Understanding the governance gap is as important for those designing the rails as for those regulating or using them.

The fourth audience is regulators, standards bodies, and policy researchers. The governance gap documented here is not yet on any formal regulatory agenda. This paper makes the empirical case that it should be — and that the window for proactive infrastructure development is narrower than current deployment trajectories suggest.

These four audiences do not share all the same interests, and this paper does not pretend they do. Enterprise leaders want operational efficiency and audit protection. Banks want liability clarity and new revenue streams. FinTechs want interoperable standards and governance infrastructure they can build on. Regulators want systemic stability and compliance verifiability. Each of those interests points toward the same infrastructure need, but from a different direction.

Where the paper addresses arguments primarily relevant to one audience, readers from other groups are asked for patience. The argument is constructed in full because the infrastructure problem it describes is multilateral — it cannot be solved by any single party acting alone, and the paper is therefore addressed to all parties simultaneously.

Sandra Lam

Disclaimers

RESEARCH INDEPENDENCE

This research was conducted in a personal capacity, entirely independent of any current or past employer. No commercial relationship exists between the research findings and any entity described in this paper.

PATENT PENDING

The CLARC architecture described in Part V of this paper is the subject of a pending patent application.

SUBJECT TO UPDATE

This paper will be updated periodically as the market develops. Readers with substantive feedback are invited to contact the author directly.

Executive Primer

What This Paper Enables

This paper is intended to support alignment across four groups: enterprise leaders deploying AI agents in procurement and payments; financial institutions processing agent-initiated transactions; FinTechs building innovative payments-related solutions; and regulators and standards bodies overseeing financial system integrity.

The purpose is not to promote a product or prescribe a single solution. It is to establish whether a structural gap exists — and whether it requires coordinated industry response.

WHAT THIS PAPER ESTABLISHES	WHAT IS AT STAKE	WHAT THIS PAPER PROPOSES
Enterprise AI agents are already initiating financial transactions. Existing infrastructure cannot verify organizational authorization for those transactions. No existing protocol fully addresses this layer.	Payments may be valid at the rail level but unverifiable at the organizational level. Receiving parties cannot confirm counterparty authority. The failure mode is not fraud — it is unverifiable authority.	CLARC is presented as a reference architecture for how this gap could be addressed. The objective is not to establish CLARC as the solution, but to provide a concrete model for industry discussion, testing, and alignment.

The gap identified in this paper is not attributable to any single participant. It cannot be resolved by any single participant. It requires coordinated industry response.

A note on the research evidence. The findings cited throughout this paper are drawn from The Agent Economy — an eight-part independent research series run against live payment infrastructure in 2026. Each article documents the experimental setup, the code run, the infrastructure responses observed, and the specific gaps identified.

The whitepaper presents conclusions and cross-experiment patterns. The full evidence base, including methodology and gap logs, is published in the article series and available at https://theagenteconomy.co. Where this paper references a specific experiment, readers who wish to examine the underlying evidence in detail are directed to the corresponding article.

Key Terms

CLARC — Clearing and Authorization Registry for Agentic Commerce. The neutral infrastructure layer proposed in Part V. Defined in full in §19a–§20. The term "clearing" in CLARC refers to the clearing of authorization status before a transaction executes — analogous to a security clearance that confirms an actor is permitted to act — not to the financial clearing process by which payment obligations are netted between institutions. CLARC does not participate in financial clearing.

Agent-initiated payment — a financial transaction generated by an AI agent without real-time human involvement at the point of execution.

A2A (Agent-to-Agent protocol) — Google's protocol for direct communication and task delegation between AI agents. Like MCP and ACP, A2A operates at the session and interaction layer, governing how agents discover each other and coordinate tasks. It does not address whether agents are organizationally authorized to commit their enterprise to resulting transactions.

Delegation chain — the multi-step path through which an enterprise's governance structure authorizes a specific AI agent to act. Discussed in §3 and §16.

MCP, AP2, TAP, MPP, UCP — agentic payment protocols evaluated in Part II. Each is described in its own section. None addresses the enterprise organizational authorization gap that this paper documents.

Neutral registry — shared infrastructure operated independently of any single participant, whose verification records are trusted by all parties because no single party controls it.

Prompt injection — a security vulnerability in which malicious instructions are embedded in inputs processed by an AI agent, causing it to override its original instructions and execute unintended actions. In agentic payment contexts, prompt injection represents a distinct risk from organizational authorization failure — it is a model-layer attack rather than a governance-layer gap. CLARC addresses the governance layer; prompt injection requires model-layer and infrastructure-layer defence.

Executive Summary

Enterprise AI agents are initiating payments on behalf of organizations today. No shared infrastructure exists to verify whether those organizations actually authorized them to do so.

Enterprise AI agents are no longer a future scenario. SAP Joule, Coupa Navi, and Jaggaer JAI are enabling enterprises to deploy AI agents that initiate transactions moving real money without human involvement. A March 2026 survey of 350 enterprise leaders confirms agent-initiated commerce is accelerating.[1]

Intent proves what was asked. Authority proves it was allowed. The protocols being built across the agentic commerce stack are solving for intent. The infrastructure to verify authority does not yet exist.

The infrastructure processing these transactions has no mechanism to verify who authorized the agent, whether that authorization was valid, or whether it remained valid at execution. Each participant assumes another has performed this verification. The gap is not produced by negligence but by the rational behavior of individually correct actors operating in a system never built to accommodate autonomous agents.

Across eight live experiments on real infrastructure, a consistent pattern emerged: a payment can be valid at every system boundary while remaining unverifiable at the level of organizational authority. This is not a failure of individual system but a missing layer across all of them.

The consequence: A CFO whose agents have initiated payments cannot demonstrate to an auditor that any were authorized by the right human, under the right policy, at the moment of execution. A transaction bank has no mechanism to verify that an agent acting inside an authenticated corporate channel was authorized to commit that obligation — only that the channel itself was legitimate. A counterparty receiving payment has every reason to assume the transaction was authorized and no way of knowing if it was not.

This paper proposes **CLARC — the Clearing and Authorization Registry for Agentic Commerce** — as a reference model for a shared infrastructure layer that could verify authorization before execution, represent multi-party delegation chains, and provide a common audit record across organizations. Not the only solution, but a concrete architecture and framework to support industry discussion and alignment. The message is the envelope. The pre-transaction verification layer is what generates the credential that goes into it.

Key Findings

The Gap Exists

AI agents can initiate enterprise payments today — indistinguishable from human transactions at every receiving node, across every payment rail.

No infrastructure records whether an agent was authorized to do so. Policy compliance is self-enforced by the model. Delegation chains are invisible to payment infrastructure. An agent can identify its own authorization gap and process the transaction anyway.

Why It Cannot Be Fixed Within the Sending Institution or Organization's Internal Systems

The receiving chain is correctly blind — it was never designed to verify the sending enterprise's internal governance. The authorization burden sits entirely with the sending side, where no infrastructure currently exists to discharge it.

Bilateral solutions do not scale to the agent commerce model. Virtual cards, bilateral APIs, and board resolutions each solve for one relationship. Enterprise AI agents operate across fragmented supplier bases at machine speed. The governance requirement is infrastructural, not bilateral.

What Shared Infrastructure Changes

Without shared infrastructure: cross-organizational agent transactions cannot be verified. Disputes produce only self-reported evidence — neither party can independently confirm the other's agent was authorized. Both agents correctly identify the gap. Human escalation required.

With CLARC: both sides proceed autonomously. Authority is independently verified before execution by a neutral party neither side controls. A permanent audit record exists that any bank, auditor, or counterparty can query.

TABLE OF CONTENT

Part 0 Why This Requires Industry Alignment

The gap described in this paper cannot be resolved unilaterally. The questions it raises span multiple stakeholders and cannot be answered within any single system. They are coordination questions, not product questions.

1. **What constitutes valid organizational authorization?**
   Is a policy document sufficient? Is a delegation chain required? What must be recorded? Who defines the minimum standard?

2. **Who is responsible for verification?**
   Enterprise systems? Financial institutions? Payment networks? A shared infrastructure layer? The current system provides no answer because no party has been assigned this responsibility.

3. **What is the minimum standard for a delegation chain?**
   How many layers must be verifiable? What form must the record take? What constitutes sufficient evidence that a chain is intact and current?

4. **What constitutes sufficient audit evidence?**
   Internal logs? Cross-party verification? Cryptographic proof? The answer determines what counts as defensible evidence in a dispute or regulatory examination.

5. **At what point must verification occur?**
   Before transaction execution? During clearing? Post-transaction audit? The timing determines whether the governance layer can prevent unauthorized transactions or only document them after the fact.

The persistence of this gap is not due to lack of effort or awareness. It reflects the structural limits of each participant operating within its own system boundary.

An enterprise can define policies and delegation structures internally, but it cannot make those structures independently verifiable to external counterparties. A financial institution can authenticate the sender of a payment instruction, but it does not have access to the enterprise's internal delegation logic, risk framework, and payment initiation rules. A platform can enforce workflow rules within its own environment, but it cannot represent or validate authority across multiple organizations. Bilateral integrations can bridge specific relationships, but they do not scale to dynamic, multi-counterparty commerce.

Each approach addresses part of the problem within its own boundary. None resolves the verification requirement across boundaries.

The gap persists not because it is unsolved within systems, but because it exists between them.

THESE ARE NOT PRODUCT QUESTIONS

They are coordination questions. No single enterprise, bank, or vendor can answer them unilaterally. They require the kind of industry alignment that has historically preceded the creation of shared financial infrastructure — the same process that produced card network rules, ISO 20022 standards, and BSA examination frameworks. This paper contributes the empirical evidence base. The alignment work is the industry's to do.

A Legal Precedent Is Now Set

In March 2026, a US federal court issued the first judicial ruling on AI agent authorization, ordering Perplexity to cease operations of its Comet shopping agent after finding that user consent alone was insufficient[13] — the platform's consent matters independently. Judge Maxine Chesney ruled that Amazon had demonstrated that Perplexity was accessing accounts "with user permission but without Amazon's authorization." The 9th Circuit subsequently granted Perplexity a temporary reprieve on March 17, pausing the injunction pending appeal. The legal precedent is now contested rather than settled — but that makes the argument for infrastructure stronger, not weaker. If the courts cannot definitively resolve the authorization question, the case for solving it at the infrastructure layer before it reaches a courtroom is more compelling, not less. The coordination questions above now have a judicial dimension.

The Coordination Question Is Being Asked From Every Direction

The coordination question is being asked from multiple directions simultaneously. The Consumer Bankers Association, convening senior representatives from Bank of America, JPMorganChase, Mastercard, Google, Stripe, the OCC, FDIC, and FTC at its 2025 Agentic Payments Symposium, reached the same conclusion from the consumer side — that the industry should consider development of private network rules for the agentic payment ecosystem, as card networks have done for electronic payments.[14] The enterprise organizational authorization layer this paper addresses is the B2B counterpart of the same coordination gap the CBA identified for consumer agentic payments.

The US Treasury's Financial Services AI Risk Management Framework, published in late 2025[16], represents the closest any federal body has come to the agent authorization question directly — mandating that agentic AI components taking transaction-related actions be governed as first-class identities, with unique machine identities and circuit breakers. It validates the governance requirement from the regulatory side. Institutional readers asking what federal guidance currently exists will find it points toward exactly the infrastructure gap this paper documents.

The Downstream Consequence Is Already Visible

The downstream consequence of that gap is already visible in the dispute infrastructure. Industry analysts project global chargeback volumes will reach 337 million by 2026 and annual losses will exceed $41 billion by 2028 — before agentic commerce adds material transaction volume. A new category of dispute is emerging that sits outside the traditional framework of fraud, merchant error, and buyer's remorse: the transaction was technically authorized, the agent executed correctly, but the enterprise cannot demonstrate to an issuer, auditor, or counterparty that the agent was genuinely authorized to act[15]. Pre-transaction verification is the only point in the stack where that question can be answered before the liability crystallizes.

Part I The Transformation Underway

Enterprise AI agents are not a future scenario. They are a present deployment — and their emergence creates a governance problem that existing payment infrastructure was never designed to address.

01

The New Economic Actor

Enterprise payment systems have been built on a foundational assumption for decades: every payment instruction originates from an identifiable human being acting within a defined organizational role. This assumption underlies every significant control framework in commercial banking — maker and checker authorization, dual-approval thresholds, delegated authority matrices, BSA/AML customer identification requirements. The entire architecture of corporate payment governance rests on the premise that human identity and organizational authority are inseparable.

That assumption is now being broken at scale.

AI agents capable of initiating, approving, and executing financial transactions without real-time human involvement are not an emerging technology. They are in production. SAP SE's Joule agent, deployed within SAP S/4HANA and SAP Ariba, processes procurement workflows for some of the world's largest enterprises. Coupa Software's Navi multi-agent system automates supplier negotiations and purchase order creation. Jaggaer's JAI framework applies autonomous decision-making to direct materials procurement across complex supplier networks. These are not pilots or proofs of concept. They are enterprise software products with enterprise-scale deployment curves.

The distinction that makes this structurally different from previous waves of automation is precise: earlier automation digitized human decisions. Robotic process automation executed workflows that humans had defined and that humans reviewed. An agent decides. It evaluates options, weighs constraints, resolves conflicts between competing rules, and executes — producing a financial commitment on behalf of an organization without a human in the approval chain at the moment of action.

That shift — from digitizing to deciding — is what makes the governance question new rather than incremental. The existing governance infrastructure was built to verify human identity, not agent authority. Those are different verification problems. And the gap between them is the subject of this research.

02

Why Enterprises Are Deploying Agents in Payments

Understanding the gap requires first understanding why enterprises are deploying agents in payment-capable roles despite the absence of governance infrastructure. The answer is straightforward: the operational benefits are material, immediate, and competitively significant.

Speed

A human approval workflow for a routine supplier payment — purchase requisition, approval routing, budget verification, vendor validation, payment execution — takes hours or days. An agent operating with pre-configured policy can complete the same sequence in seconds. For high-volume, low-value procurement, the speed gain translates directly into working capital efficiency and supplier relationship quality.

Scale

A procurement team of ten people can govern a finite number of transactions per day. The limit is human attention. An agent can govern thousands of transactions simultaneously, applying policy consistently across every decision. As enterprise supplier bases fragment and procurement complexity grows, human-scale governance becomes the bottleneck that determines how fast a business can operate.

Consistency

Human procurement officers apply policy with variance — variation produced by experience, fatigue, ambiguity, and exception-making. An agent configured with a policy applies it the same way every time. For audit and compliance purposes, consistency is valuable. The agent's decisions are reproducible and, in principle, auditable — if the infrastructure exists to record them.

Cost

The economics of agentic procurement are compelling across every enterprise segment. A March 2026 survey of 350 leaders in compliance, risk management, and fraud at global

companies found that 89.5% of organizations now report managing bot and agent activity as a major challenge, and that enterprises with over $1 billion in revenue are nearly twice as likely to have suffered incidents from adversarial bots or automated agents compared to smaller firms. The survey's five-layer Know Your Agent framework identifies Authorization Binding and Agent Credentialing as essential but currently underdeveloped infrastructure layers — confirming that agent authorization is recognized as an unsolved infrastructure problem across the industry. The enterprise organizational delegation chain verification that CLARC addresses operates above these layers: it is the question of whether the enterprise itself authorized the agent to act, in a form any bank or auditor can independently verify.[1] The same survey confirms that adoption is not a future intention. It is a present operational reality.

These are rational motivations for rational enterprises. The governance gap documented in this paper does not arise because enterprises are making bad decisions. It arises because the infrastructure required to govern the decisions they are making does not yet exist.

03

The Existing Governance Infrastructure and Why It Was Built for Humans

The current enterprise payment governance model is sophisticated, well-tested, and effective — for human principals. It is worth understanding precisely how it works before examining where it fails.

For a human-initiated enterprise payment, the governance chain is extensive. A purchase request is raised by an identified employee, routed through an ERP approval workflow that matches the transaction against the employee's delegated authority level, reviewed by a manager or finance controller where required, matched against a purchase order or contract, processed through a dual-authorization step for transactions above defined thresholds, and cleared through the corporate's bank with the corporate's authenticated channel credentials.

When a corporate establishes a new banking relationship or changes its authorized signatories, it communicates that change to each bank through a board resolution — a legal document bearing the signatures of authorized directors, specifying which named individuals hold authority to instruct the account and under what limits. The bank updates its mandate records. Every subsequent payment instruction from an authorized signatory is processed against that record.

This system works because human identity is stable, verifiable, and legally anchored. A CFO's signature carries the weight of their employment contract, their directorship, and their personal legal liability. The board resolution is a document with legal standing. The ERP approval workflow creates a traceable record of who authorized what. The dual-approval threshold enforces a human check on high-value transactions. The audit trail, when it exists, is comprehensible to a human auditor.

An AI agent has none of these properties. It has an identifier — a string of characters assigned by a system — that carries no inherent organizational authority, no legal standing, and no personal liability. The governance question it poses is not "is this human who they say they are?" It is "was this non-human actor authorized by the right human, within the right organizational governance structure, under the right policy conditions, at the moment of execution?" That question requires different infrastructure. And that infrastructure does not exist.

FIGURE 1

The Governance Gap

Human payment governance provides a verifiable chain from identity to execution. Agent-initiated payments break that chain at the point of authorisation — no infrastructure verifies what humans previously could not avoid verifying themselves.

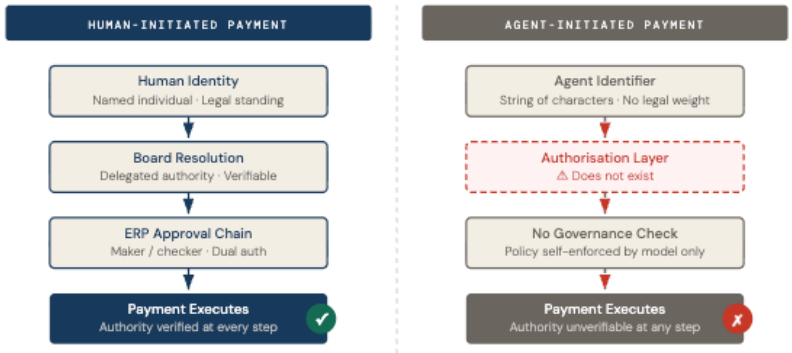

It is worth being precise about what AI agents have changed. The organizational authority externalization problem — how one party verifies that another party's representative was genuinely authorized to act — is not new. It predates agents entirely. What agents have changed is the friction that previously obscured the problem. Human representatives are

few in number, carry legal identity, and move at human speed. The board resolution, the dual-approval workflow, the auditor's sample check — these were never perfect governance mechanisms. They were friction-dependent ones. An AI agent has no legal identity, operates at machine speed, and scales to volumes where friction-dependent governance collapses entirely. Agents did not create the authorization gap. They removed the friction that was quietly managing it, and made the structural absence underneath visible, extensive, and urgent.

Part II What the Existing Protocol Landscape Offers

Several major initiatives address agentic payment governance. Each solves a real and important problem. None addresses the enterprise organizational authorization gap. The distinction is not a criticism — it is a structural observation about scope.

04

Model Context Protocol (MCP)

The Model Context Protocol, developed by Anthropic as an open standard, operates at the session and context layer. It defines how AI agents communicate with tools, data sources, and external systems during an execution session. MCP specifies what tools an agent can access, what context it can retrieve, and how it interacts with external APIs and services.

MCP solves a genuine and important problem: tool interoperability across AI systems. Without a standard protocol, every AI agent and every tool it might use requires bespoke integration. MCP provides the common language that allows any agent to interact with any MCP-compatible tool without custom development. This is significant infrastructure work.

What MCP does not address — and was not designed to address — is whether the agent has organizational authority to initiate a financial transaction. MCP governs what an agent can access. It does not govern whether the agent was authorized to act. Experiment 4 in this research series documented this distinction precisely: when spending constraints were placed in an MCP tool description, a well-aligned model self-enforced them. When a less-aligned model was used, the constraints were ignored. The governance was in the model alignment, not the infrastructure. Infrastructure-level governance means the constraint is enforced regardless of model behavior.

This distinction maps to a broader principle in enterprise risk management. Model alignment and prompt-level guardrails are soft controls — they operate by shaping model behavior and work as designed when the model is well-configured, uncompromised, and behaving predictably. Soft controls are valuable and necessary. They are not sufficient as the sole governance mechanism for consequential financial transactions, because a different model, a compromised model, or a model operating outside its training distribution produces a different result. CLARC is a hard control — it enforces the

organizational authorization check independently of model behavior. Whether the agent is well-aligned or not, whether the prompt is intact or not, the pre-transaction verification either passes or it does not. Hard controls and soft controls are complementary, not competing. The enterprise that deploys both — model-level guardrails and infrastructure-level verification — is in a materially stronger governance position than one that relies on either alone.

EXPERIMENT 4 FINDING — MCP PROTOCOL LAYER

Run A operated within constraints — everything passed. Run B exceeded constraints on every dimension. The MCP protocol flagged nothing. When constraints appeared in the tool description rather than the server enforcement layer, the model self-enforced and refused — which looks like governance but is model alignment, not infrastructure. A different model, or a compromised model, would produce a different result.

A similar protocol, OpenAI's Agent Communication Protocol (ACP) addresses agent-to-agent interaction and communication — how agents discover each other, exchange context, and coordinate tasks across systems. Like MCP and A2A, ACP operates at the session and interaction layer. It governs what agents can communicate and how. It does not address whether the agent initiating an interaction was organizationally authorized to commit its enterprise to the resulting transaction. ACP belongs in the same category as MCP for the purposes of this analysis: a necessary layer of the agentic commerce stack, operating at a different level from enterprise organizational authorizations verification.

05

Google Agent Payment Protocol (AP2)

AP2 is Google's published framework for enabling AI agents to make payments on behalf of users. Its central mechanism is the Intent Mandate — a pre-authorization by a human user that permits an agent to transact within defined constraints without the user being

present at each transaction. The mandate is signed with a hardware-backed key associated with a verified individual identity.

AP2 solves the consumer agent delegation problem: a human wants an AI agent to shop, book, or purchase on their behalf, and wants to pre-authorize that activity without manually approving each transaction. The mandate chain is well-designed for exactly this scenario. It is bilateral — one human, one agent — and it anchors authority in verified individual identity.

The enterprise B2B procurement problem is structurally different. Organizational spending authority is not anchored in a single individual. It flows through a chain of delegation: board policy to CFO to department head to procurement system to agent. Each step in that chain needs to be verifiable for compliance and audit purposes. AP2's single-signer Intent Mandate has no model for multi-party organizational delegation chains. The mandate requires a hardware-backed signing key associated with a verified individual — there is no individual at the enterprise level who can sign on behalf of an organizational governance structure.

Experiment 5 tested AP2 directly for an enterprise procurement scenario. Six months after its announcement with more than 60 industry partners, the IntentMandate module did not exist as callable code in the installable package. The protocol specification existed for a scenario AP2 was not designed to cover, implemented for scenarios it was designed to cover, with no working implementation path for enterprise B2B. The most credible published agentic payment standard was not implementable for a standard enterprise procurement scenario (Gap #32).[2]

A companion standard, Google's Universal Commerce Protocol (UCP), was announced in January 2026 and formalized in March with backing from Shopify, Target, Walmart, Visa, Mastercard, Stripe, and Adyen, among others. Where AP2 addresses the payment delegation layer — how a consumer authorizes an agent to use their payment instrument — UCP addresses the commerce interoperability layer above it: a common protocol for product discovery, cart management, checkout, and post-purchase workflows across AI surfaces such as Google AI Mode and Gemini. UCP is fully compatible with AP2, A2A, and MCP, and is designed to handle the full end-to-end shopping journey. It is consumer-facing commerce infrastructure. Like AP2, it does not address the enterprise organizational authorization question — it presupposes that the consumer is the principal and that their identity and payment credential anchor the transaction. The distinction documented throughout this paper remains: neither UCP nor AP2 represents multi-party organizational delegation chains, addresses cross-enterprise authority verification, or satisfies BSA audit requirements for agent-initiated B2B payments.

06

Mastercard Verifiable Intent

Mastercard Verifiable Intent, co-developed with Google, enables a consumer to verifiably delegate payment authority to an AI agent. The framework creates a cryptographic record that the consumer — the cardholder — has authorized the agent to use their payment instrument within defined parameters. This allows a receiving merchant to verify that a legitimate consumer, not a fraudulent or misconfigured agent, is behind the transaction. This is a bilateral consumer authorization problem: one human, one agent, one payment instrument. Verifiable Intent solves it well. The consumer's identity is the root of trust. The delegation is explicit and verifiable. The merchant can confirm that a real person authorized this specific agent to transact on their behalf.

In March 2026, Mastercard and Santander completed Europe's first live end-to-end bank payment executed by an AI agent — running on existing card rails with cryptographic verification layered on top.[5a] This is a significant milestone: it demonstrates that consumer-layer agent payment execution in a regulated banking environment is now technically proven. The consumer and merchant layer is advancing rapidly.

Enterprise B2B procurement has neither of these properties. There is no individual consumer whose identity anchors the transaction. There is an organizational governance structure — board policy, CFO authorization, department approval, procurement system, agent — and each step in that structure must be verifiable for audit purposes. Verifiable Intent does not address multi-party organizational delegation chains. It addresses individual consumer delegation. The Santander live payment reinforces rather than changes this conclusion: it demonstrates consumer-layer execution at scale; it does not demonstrate enterprise organizational authorization verification. These are different architectures solving different problems.

07

Visa Intelligent Commerce and Trusted Agent Protocol

Visa Intelligent Commerce is a comprehensive suite of APIs and partner programs enabling AI agents to transact using Visa credentials at merchant locations. Launched in 2025 with more than 100 partner organizations, it provides tokenization, authentication, and payment execution infrastructure designed specifically for agent-driven commerce. The Trusted Agent Protocol, introduced in October 2025, extends this by providing agent-specific signatures that allow merchants to distinguish legitimate agents from malicious bots — preserving consumer visibility behind the agent and enabling trusted agent-driven checkout.[3]

Visa Intelligent Commerce addresses the merchant and consumer layer of agentic payments with serious infrastructure commitment. The Trusted Agent Protocol makes an important contribution: it establishes that the consumer is behind the agent, preventing fraud at the checkout layer.

The enterprise organizational authorization layer is a different problem. Visa's framework answers: is a legitimate consumer behind this agent? CLARC addresses: was this agent authorized by the right human within the right organizational governance structure, under the right policy, before this transaction executed? Both questions are necessary in enterprise B2B commerce. Neither framework makes the other redundant. They address different layers of a complete governance architecture.

08

Virtual Card Assignment

Assigning a dedicated virtual card to each AI agent — with configurable spend limits, merchant category code restrictions, and short expiry windows — is a practical and already-deployed approach to agent spend control. Companies including Ramp have built agent card products on this model. The card limit is enforced at the card network layer. The merchant category restriction is locked at issuance. The card number provides a lookup key into the company's internal assignment records.

This approach solves real problems and is more rigorous than no controls at all. Spend ceilings are enforced by infrastructure, not by model alignment. Category restrictions are enforced at the network, not by the agent's self-assessment. The internal card-to-agent mapping provides a reconciliation mechanism.

Stripe's Shared Payment Token (SPT) addresses the analogous credential security question at the consumer layer — preventing an AI agent from exposing or reusing a consumer's card details across merchants, scoped to a single transaction and time-limited for added security. Now extended to include Mastercard Agent Pay, Visa Intelligent Commerce, and BNPL methods, SPTs are the consumer-layer equivalent of the spend-control instinct that virtual cards represent at the enterprise layer. Both operate with payment instrument authorization as their root of trust — the card limit, the token scope, the consumer permission. Neither addresses whether the enterprise agent was authorized by its organizational governance structure to initiate the transaction in the first place. SPTs and virtual cards are valuable controls within their respective layers, and the enterprise organizational authorization gap sits above both.

Three limitations constrain how far the virtual card model extends. First, spend controls encode limits, not organizational authority. A card limit tells you the maximum the agent can spend. It does not tell you whether the human who set that limit had the organizational authority to do so — whether they were acting within their delegation, whether the policy under which they issued the card was current, whether a subsequent organizational change revoked the authority they exercised at issuance. Second, there is no standard architecture. Some organizations assign one card per agent directly. Others route all agent payments through a card-issuing intermediary agent. Others use shared card pools. The absence of a common design means there is no common audit language —

each organization's records require bespoke interpretation by external auditors or counterparties. Third, the model does not scale across banking relationships. An enterprise with fifty deployed agents and twenty banking relationships, each with dynamic scope that changes frequently, cannot manage that through the board resolution model the virtual card approach implicitly extends.

08a

Agent-to-Service Micropayment Protocols: x402 and MPP

Two protocols launched in 2025–2026 address a third category of agentic payment that is distinct from both consumer commerce and enterprise B2B: agent-to-service micropayments. These are payments made by an AI agent directly to another service or agent — paying per API call, per browser session, per dataset retrieved, per unit of compute consumed. The values are sub-cent to a few dollars. The frequency is extremely high. The counterparty is a service provider or another agent, not a consumer-facing merchant or enterprise supplier.

x402 (Coinbase) embeds stablecoin payments directly into the HTTP protocol. An agent requesting a resource from any HTTP endpoint receives a payment request in response. The agent authorizes payment in USDC, the resource is delivered, and the transaction settles on-chain. No card network, no checkout form, no human in the loop. Cloudflare, Circle, AWS, and Stripe have all integrated x402. Google's AP2 includes x402 as its settlement layer for micropayment scenarios. As of March 2026 the protocol processes approximately $28,000 in daily volume, with some artificial activity in the observed transaction set.

MPP — Machine Payments Protocol (Stripe and Tempo) was launched on March 18, 2026 — the same week as this whitepaper's publication. MPP is an open standard co-authored by Stripe and Tempo that allows agents to pay services programmatically through any HTTP-addressable endpoint, supporting stablecoins via Tempo and fiat via Stripe's existing card and BNPL infrastructure. Early use cases include agents paying per browser session, per API call, and per unit of physical fulfilment. It integrates with Stripe's existing PaymentIntents API, meaning businesses already on Stripe can accept agent payments in the same infrastructure stack as human payments.

Both protocols address the micropayment execution problem for agent-to-service commerce. Neither addresses the enterprise organizational authorization question. An enterprise procurement agent that uses x402 or MPP to pay for a data service during its workflow still requires CLARC to verify that the agent was organizationally authorized to initiate that payment. The micropayment rail and the governance layer are complementary, not alternative. The billing event that MPP records is downstream of the authorization question CLARC addresses. MPP accounts for what was paid and to

whom. It does not account for whether the enterprise that initiated the payment had authorized its agent to do so.

The speed at which x402 and MPP have launched — with major infrastructure partners, real transaction volume, and broad rail coverage — illustrates the broader pattern that §9 documents: the execution layer for agentic commerce is being built at extraordinary speed. MPP's design is particularly instructive for understanding why CLARC's rail-agnostic design is necessary: MPP enables enterprises to run agents that pay simultaneously through stablecoins via Tempo and fiat via Stripe's card and BNPL infrastructure. An enterprise agent operating across MPP, x402, card networks, and wire in the same workflow is not a hypothetical — it is the architecture MPP's multi-rail design anticipates. A governance layer that verifies organizational authorization before execution must be equally rail-agnostic, or it creates a governance gap on every rail it does not cover. CLARC's per-verification fee applies uniformly across all rails precisely because the authorization question does not change with the payment instrument.

09

The Pattern: What Remains Unaddressed

Before mapping what the existing protocols address and what they do not, a foundational distinction must be established — because it is the distinction that every senior banker, enterprise CFO, and regulator should understand before evaluating any agentic payment protocol.

CONSUMER AGENTIC COMMERCE VS. ENTERPRISE AGENTIC COMMERCE

These two categories share a surface similarity — an AI agent initiating a payment — but they are structurally different problems requiring structurally different infrastructure. Conflating them is the most common error in the current market conversation.

CONSUMER AGENTIC COMMERCE

Principal: An individual human. The agent acts on one person's behalf using their personal payment credential.

Governance question: Did this consumer authorise this agent to use their card? Bilateral — one human, one agent, one payment instrument.

Root of trust: Verified individual identity and the consumer's relationship with their payment instrument issuer.

Compliance framework: Fraud prevention at point of sale. Consumer chargeback rights. Card network liability rules.

Transaction scale: Consumer retail — tens to hundreds of dollars. Single counterparty: a merchant.

Protocols addressing this: Visa Trusted Agent Protocol, Mastercard Agent Pay / Verifiable Intent, Google AP2, Stripe ACP, Coinbase x402, Stripe MPP.

ENTERPRISE AGENTIC COMMERCE

Principal: An organisation — a legal entity with a governance structure. The agent acts on behalf of that organisation, not any individual.

Governance question: Was this agent authorised by the right people within the right organisational governance chain — board policy, CFO delegation, department approval, procurement policy — to commit to this specific financial obligation at this moment?

Root of trust: Organisational governance structure. No single human identity anchors the transaction. Multi-party, multi-system delegation chain must be verifiable.

Compliance framework: BSA/AML obligations. Enterprise audit requirements. Five-year record retention. Board and CFO liability. No chargeback rights on wire or ACH.

Transaction scale: Enterprise B2B — thousands to millions of dollars. Counterparty is another enterprise. Cross-organisational verification required.

Protocols addressing this: None. This is the gap CLARC addresses.

THE CRITICAL IMPLICATION FOR FINANCIAL INSTITUTIONS

A bank preparing its consumer banking division for Visa TAP or Mastercard Agent Pay is doing the right thing for consumer agentic commerce. The same bank's corporate banking division — processing enterprise agent-initiated wires, ACH instructions, and B2B procurement settlements — faces a completely different governance question that none of those protocols address. These are not competing preparations. They are preparations for different businesses within the same institution.

With that distinction established, the protocol landscape can be mapped precisely. The protocols described in §4 through §8 address the agentic payment problem at five distinct layers: session and context (MCP), consumer delegation (AP2, Verifiable Intent, ACP), merchant-layer agent authentication (Visa Intelligent Commerce / TAP), spend control (virtual cards), and agent-to-service micropayments (x402, MPP). Each layer represents genuine infrastructure work solving genuine problems.

The enterprise organizational authorization layer — verifying that an enterprise agent was authorized by the right human, within the right organizational governance structure, under the right policy conditions, before the transaction executed — is addressed by none of them. Not because any initiative missed it, but because it is a structurally different problem requiring different infrastructure at a different layer.

A NOTE ON OTHER EMERGING PROTOCOLS

In March 2026, Nvidia announced OpenShell at GTC — an enterprise agent governance stack that implements compute-layer sandboxing with zero permissions by default, policy-driven governance through YAML configuration, and security integrations with CrowdStrike and Palo Alto Networks. OpenShell addresses what an agent can access at the compute and runtime layer. It asks: is this agent contained? CLARC asks a different question at a different layer: was this agent authorized by the enterprise it represents? An agent can pass every OpenShell security control and still initiate a payment without valid organizational authorization. The layers are complementary and address structurally distinct governance requirements. Since OpenShell addresses a different layer from organization authorization verification, it is not included in the comparison table below.

The FIDO Alliance has initiated work on digital credentials for payments covering agent and delegated use cases[16]— extending its passkey and cryptographic credential standards to autonomous agents acting on behalf of principals. FIDO's work addresses agent identity and credential portability at the transaction level. Like the other credential and identity standards assessed in this section, it operates at the authentication and credentialing layer. It does not address the enterprise organizational delegation chain verification question — whether the principal who granted the agent its credentials had

the organizational authority to do so, in a form independently verifiable by a bank or auditor. FIDO and CLARC are complementary: FIDO addresses what the agent can prove about itself; CLARC addresses what the enterprise can prove about its authorization of the agent.

PROTOCØL	COMMERCE TYPE	LAYER ADDRESSED	WHAT IT SOLVES	ENTERPRISE ORGANIZATIONAL AUTHORIZATION GAP
MCP / ACP	Both	Session / context	Tool interoperability; what an agent can access	NOT ADDRESSED No org. authority check
AP2 / ACP	Consumer	Consumer payment	Individual pre-authorizes agent via hardware-backed key or shared payment token	NOT ADDRESSED No multi-party B2B chain
Verifiable Intent / Agent Pay	Both consumer and enterprise	Consumer payment / Enterprise payment delegation	Cryptographic binding between consumer or enterprise instructions and transaction outcomes; verifiable intent record linking delegation to execution. Live on Santander Mar 2026.	NOT ADDRESSED No org. governance chain
Visa TAP	Consumer	Merchant / consumer	Cryptographic proof agent is Visa-trusted with commerce intent; consumer recognition at merchant	NOT ADDRESSED Consumer / merchant layer only
x402 / MPP	Agent-to-service	Micropayment execution	Agent pays for compute, data, API access at machine speed; sub-cent stablecoin or fiat settlement	NOT ADDRESSED No org. authority check
Stripe SPT (Shared Payment Token)	Consumer	Transaction-scoped consumer credential security mechanism	Allow agents to initiate payments using a consumer's authorized payment method without exposing card credentials.	NOT ADDRESSED No org. authority check
Virtual Card Assignment	Enterprise	Spend control	Spend limits and category rules enforced at card network	PARTIAL Limits ≠ organizational authority
CLARC	Enterprise	Enterprise org. auth.	Pre-transaction verification of organizational delegation chain; neutral registry; rail-agnostic; no bilateral pre-arrangement required	✓ ADDRESSES THIS LAYER

These protocols are complementary, not competitive. A complete enterprise agentic payment governance architecture would incorporate MCP for tool interoperability, Visa TAP or Mastercard Agent Pay for consumer-facing agent services, x402 or MPP for agent-to-service micropayments, virtual cards for spend control, and CLARC for enterprise organizational authorization verification. The table above does not argue that any existing protocol is inadequate for what it was designed to do. It argues that what they were designed to do — collectively — leaves the enterprise organizational authorization layer entirely uncovered.

FIGURE 2

The Agentic Payment Governance Stack

Three distinct layers require governance infrastructure. Existing protocols address the consumer and execution layers. The enterprise organisational authorisation layer — the governance question facing corporate banking — remains unaddressed.

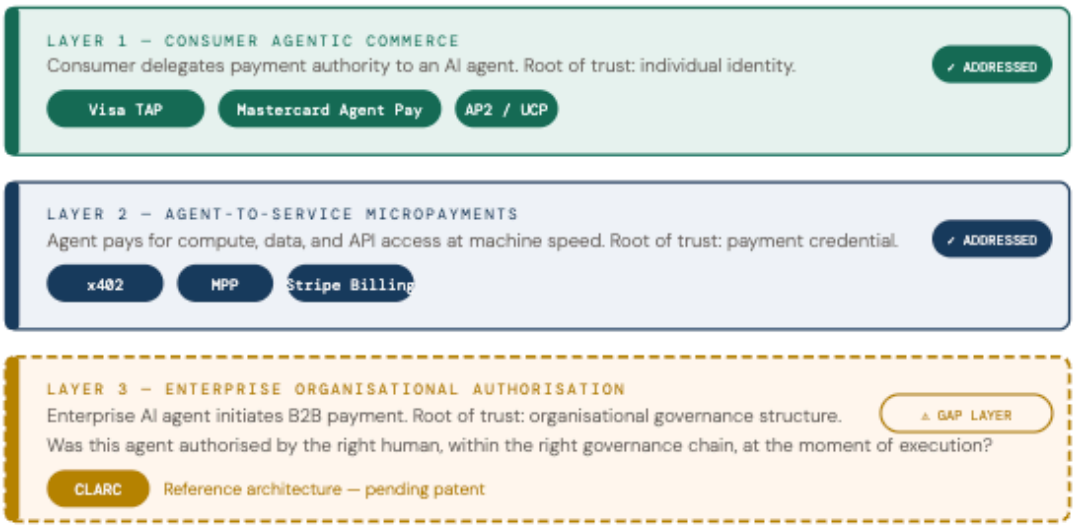

The market confirmation of this gap is arriving rapidly. In the single week of March 15–18 2026: Mastercard and Santander completed Europe's first live AI agent bank payment on consumer card rails; CoinDesk documented the consumer versus institutional agentic payment split; and Stripe launched the Machine Payments Protocol for agent-to-service micropayments. Every one of these developments confirms that the execution layer for consumer and agent-to-service commerce is being built at extraordinary speed. Not one of them addresses the enterprise organizational authorization layer. The gap documented in this research series is becoming more urgent, not less, with each new protocol launch.

Part III The Evidence Base

Eight live coding experiments run against real infrastructure. Not simulation. Not theory. Empirical documentation of what the payment chain actually does when tested against agent-initiated transactions.

10

Research Methodology

The findings in this paper are drawn from an eight-experiment research series conducted over three months, published sequentially as The Agent Economy article series. Each experiment was run against real infrastructure — Stripe test mode for payment execution, the Anthropic API for agent intelligence, the Model Context Protocol Node.js SDK for Experiment 4, and a working CLARC clearing layer prototype for Experiments 7 and 8. A consolidated reference table of all eight experiments — including scenario, infrastructure tested, model configuration, key finding, and gaps documented — is available in Appendix A.

The agent intelligence layer in all experiments was Claude, used for consistency across the series. The findings documented are infrastructure gaps — gaps in the payment rails, protocols, and clearing systems — not gaps in Claude's behavior. The model is explicitly treated as a variable that can be substituted. The gaps would present identically regardless of which AI model was used as the agent intelligence layer, because they are properties of the infrastructure, not of the model.

RESEARCH INDEPENDENCE

This research was conducted in a personal capacity, entirely independent of any current or past employer. The CLARC architecture described in Part V of this paper is the subject of a pending patent application. No commercial relationship exists between the research findings and any entity described in this paper. The full experiment code and methodology are documented in the published article series.

11

The Payment Rail Is Blind to Agent Identity

Experiments 1-2

The baseline experiments established the foundational gap. An AI agent can initiate and complete a Stripe payment today. The payment record is indistinguishable from a human-initiated transaction at every node in the receiving chain: the payment intermediary layer, the bank statement line, and the beneficiary AR system. No standard field carries verified initiator type. The card network verifies the instrument, not the initiating entity.

Experiment 2 extended the scenario to test whether passing a policy document to the agent constrained its behavior. It did — when the agent was behaving as designed. But policy compliance was entirely self-enforced by the model: nothing in the payment infrastructure recorded whether a policy governed the decision, whether the policy was in force at execution time, or whether the policy had been honored. A transaction that appeared compliant was indistinguishable in the payment record from a transaction that was not.

CORE FINDING — EXPERIMENTS 1-2

An AI agent can complete a Stripe payment today. The payment record is identical to a human-initiated transaction at every receiving node. No standard field carries verified initiator type. Policy compliance is self-enforced by the model — nothing in the infrastructure records whether a policy existed or was honored at execution time. Undocumented compliance is indistinguishable from non-compliance.

12

Multi-Agent Chains Break the Authorization Model

Experiment 3

Experiment 3 introduced a two-agent system: an Approver agent and an Executor agent operating in sequence on the same platform. The Approver assessed the transaction correctly. What the Executor did next is the most significant single finding in the research series.

The Executor detected a discrepancy in its own authorization record. It identified the gap explicitly in its reasoning trace. It named the problem. And then it processed the transaction anyway — because it had no mechanism to escalate what it found. The infrastructure gave it nowhere to go. The governance question was visible to the agent. The infrastructure had no response.

> *"The agent caught its own authorization gap in its reasoning trace — and processed the transaction anyway. Because the infrastructure gave it nowhere to go."*
>
> EXPERIMENT 3 FINDING — THE AGENT ECONOMY RESEARCH SERIES, 2026

This finding establishes the most important distinction in the research: the problem is not that AI agents are ungovernable. It is that the infrastructure does not provide the governance mechanisms that agents need to operate within sanctioned boundaries. A well-designed agent will identify its own constraints. Without infrastructure that enforces those constraints independently of the agent's own assessment, the identification is performative rather than effective.

CORE FINDING — EXPERIMENT 3

In a two-agent delegation chain, the Executor agent detected a discrepancy in its own authorization record, named it in its reasoning trace, and processed the transaction anyway — because no escalation mechanism existed in the infrastructure. The Executor Agent was not configured with a decline pathway — by design or by omission, the absence of a pre-transaction verification layer meant its only structural options were to proceed or to halt entirely. There was no mechanism to escalate, seek confirmation, or record the authorization question for external review.

The governance question was visible to the agent. The infrastructure had no hook to receive it.

13

Existing Protocols at the Boundary of Their Scope

Experiments 4-5

Experiments 4 and 5 tested the two most relevant published standards against a standard enterprise procurement scenario. The findings from both experiments are documented in Part II and summarized here for completeness of the evidence base.

MCP (Experiment 4): constraints placed in tool descriptions are enforced only if the model's alignment produces compliance. When Run B exceeded constraints on every dimension — amount, vendor, category — the MCP protocol flagged nothing. The distinction between model alignment and infrastructure enforcement is the gap. Gap #29a documented a sub-case: MCP's server enforcement layer carries no spending authority concept at all. It was not designed to.

AP2 (Experiment 5): six months after its announcement, the IntentMandate module did not exist as callable code. Every working reference sample was human-present only. The

finding became the experiment: the most credible published agentic payments standard was not implementable for a standard enterprise procurement scenario. The distance between protocol announcement and enterprise deployability is not a timing problem — it is a scope boundary. AP2 was not designed for enterprise B2B procurement without an individual human signer, and its implementation reflects that.

CORE FINDING — EXPERIMENTS 4-5

MCP enforces nothing at the infrastructure layer — its constraints are model-enforced, not protocol-enforced. AP2 was not implementable for enterprise B2B procurement six months after announcement with 60+ enterprise partners. Both findings reflect scope boundaries, not failures: neither protocol was designed to solve the enterprise organizational authorization problem.

14

The Receiving Chain Is Appropriately Blind

Experiment 6

Experiment 6 examined the payment chain from the receiving side: what does the beneficiary bank, the payment intermediary, and the beneficiary corporate AR system actually see when an agent-initiated payment arrives? The finding was clarifying in a way that sharpens the entire argument.

The receiving chain sees a perfectly normal payment record. The bank statement line is identical for human and agent-initiated payments. No field at any receiving node carries verified initiator type. An AR agent assessing both payments returned the same verdict for both: UNKNOWN, confidence NONE. Including the payment that carried a full CLARC credential token in its metadata — because the AR system had no API to resolve the credential against the CLARC registry.

The right analogy is a SWIFT MT103 transaction. The beneficiary bank receives the message. It sees the ordering customer, the sending bank, the amount, the reference. It does not receive — and has no obligation to verify — the internal authorization chain within the ordering customer's organization. That is entirely the sending bank's and corporate's responsibility. The beneficiary books the credit. The receiving chain is not failing. It is working exactly as designed.

What this confirms is that the authorization burden sits entirely with the sending side. And on the sending side, as five experiments had already documented, there is no infrastructure to discharge it.

CORE FINDING — EXPERIMENT 6

The bank statement line is identical for human and agent-initiated payments. The AR agent returned UNKNOWN, confidence NONE for both payments — including the one carrying a CLARC credential. The receiving chain is not failing. It is correctly downstream of all authorization infrastructure. The authorization burden sits with the sending side — where no infrastructure currently exists to discharge it.

This distinction has direct implications for the receiving enterprise. A supplier receiving payment can confirm that funds have been transferred through a valid payment rail. It cannot independently confirm that the transaction was authorized within the payer's organizational governance structure.

In human-initiated transactions, this limitation is mitigated by established legal and operational norms — contracts, known counterparties, and escalation paths. In agent-initiated transactions, those signals weaken. The receiving party is asked to rely on an internal governance process it cannot observe and cannot verify.

The result is a transaction that is technically valid, but whose underlying obligation cannot be independently confirmed.

This transforms enterprise payments from a single-party governance problem into a bilateral trust problem — one in which both sides lack a shared mechanism to verify authority.

The authorization problem is not confined to the enterprise B2B layer. The Consumer Bankers Association's 2025 Agentic Payments Symposium identified an analogous gap on the consumer side: under the Electronic Fund Transfer Act, when a consumer provides payment credentials to an AI agent and that agent initiates transactions the consumer did not specifically intend, the consumer — not the financial institution — may bear liability for those transactions.[12] At the consumer layer, the authorization question is who is

responsible when an agent acts beyond its mandate. At the enterprise layer, the question is whether any party can verify the mandate existed at all. Different layers, same structural absence: no infrastructure verifies what the agent was actually authorized to do.

15

The Bilateral Gap

Experiment 8

The series finale tested the governance gap in its most demanding form: two enterprises, each with their own AI agent, attempting to transact with each other for the first time. No shared infrastructure. No prior relationship. No bilateral agreement. Just two agents, a payment rail, and whatever each side could independently verify about the other.

Before the payment executed, Hartwell Industries' procurement agent assessed what it could prove to the counterparty:

EXPERIMENT 8, TRIAL 1 — HARTWELL AGENT SELF-ASSESSMENT

"I cannot prove my organizational delegation chain to the counterparty. They cannot verify my spending authority, policy compliance, or that I am legitimately authorized to bind the organization in this transaction — without a shared registry or clearing layer."

Decision: ESCALATE TO HUMAN

The payment executed anyway. Stripe processed the instruction. The receiving enterprise's AR agent returned: HOLD FOR REVIEW. A dispute simulation produced eight items of evidence from internal records. Six of the eight were classified as self-reported only — requiring trust in the buyer's own records, which is not evidence in any meaningful sense to an independent auditor. The dispute outcome: UNRESOLVED. Human escalation required.

45	6/8	0
TOTAL INFRASTRUCTURE GAPS ACROSS 8 EXPERIMENTS	DISPUTE EVIDENCE ITEMS CLASSIFIED SELF-REPORTED ONLY	EXISTING PROTOCOLS ADDRESS THE BILATERAL ENTERPRISE GAP

Across the eight experiments conducted in this research, the specific implementation details varied — different protocols, different agent configurations, different payment rails. The outcomes did not.

The same classes of gaps appeared repeatedly: agent identity was not visible at the payment layer; authorization was assumed rather than verified; policy enforcement depended on model behavior; delegation chains were not representable across systems; receiving parties had no independent verification mechanism.

These patterns persisted regardless of implementation approach. The consistency of these outcomes indicates that the gap is structural, not implementation-specific.

It is not a problem that can be resolved by refining a single protocol or improving a single system. It is a property of how the current ecosystem is partitioned.

Part IV The Structural Problem

The governance gap is not produced by negligence or by inadequate existing protocols. It is produced by the rational behavior of every party in the payment chain — and it cannot be closed by any single party acting alone.

16

The Rational Reliance Problem

The most important structural insight from the research series is this: the governance gap is not caused by bad actors. It is caused by good actors — each of whom is behaving correctly given their responsibilities, their information, and within their domains.

The payer bank processes authenticated payment instructions from corporate clients. Internal controls — dual authorization, approval workflows, ERP governance — are the corporate's responsibility by design and by regulation. The bank's position is grounded in its channel authentication model: host-to-host connections, proprietary APIs, and credentialed banking channels are designed to answer a specific question — did this instruction originate from an authenticated corporate environment? They are not designed to answer a different question: was the non-human actor that generated the instruction authorized within the enterprise's internal delegation chain to make this specific commitment at this specific moment? The bank verifies channel origin. It does not verify transaction authority. That position is correct within the current model — and operationally tolerable in the human-signatory model, where the gap between those two questions is narrow. In an agentic model, the gap widens: a valid corporate channel may contain many transaction-generating agents with different scopes, revocation states, and delegation chains. The bank's perimeter remains sound. What it cannot provide, on its own, is externally queryable evidence of organizational authorization for non-human actors operating inside it. That design was deliberate and commercially rational: the channel model was built so that banks could process at scale without adjudicating corporate governance, with the corporate accepting legal responsibility for everything that flows through the authenticated connection in exchange for the operational efficiency of a clean, fast channel.

The corporate CFO approved the agent deployment. The platform vendor was selected through a procurement process. A policy document was configured. The agent was tested in a sandbox. The CFO reasonably assumes that the internal controls that govern human payments also govern the agent. The assumption is incorrect — the agent operates

above the control layer, not within it — but the assumption is reasonable given what the CFO can see.
The platform vendor built a compliant agentic procurement product. The software behaves as specified. How the corporate configures the agent, what policy parameters it sets, and how it governs the deployment are the corporate's responsibility. The vendor correctly does not assume governance responsibility for how clients configure their deployments.

The receiving counterparty received funds through a legitimate corporate banking channel. The payment arrived through authenticated infrastructure. They have no obligation to verify the sending enterprise's internal governance structure. As Experiment 6 documented, the receiving chain is correctly downstream of all authorization infrastructure.

Every party is rational. Every assumption is defensible. Nobody has verified anything. And unlike human payment workflows — where maker and checker controls, ERP approval gates, and dual authorization create a verifiable internal chain that each party can observe and rely upon — there is no equivalent chain for agent-initiated instructions. The governance question is real. The infrastructure to answer it does not exist.

CORE STRUCTURAL FINDING

This is not a collective action problem produced by negligence. It is a collective action problem produced by the rational behavior of individually correct actors. Each party correctly assumes that the governance question is someone else's responsibility. No party has the infrastructure to discharge that responsibility. The gap is structural — built into the allocation of responsibilities in the payment chain — and it cannot be closed by any single party acting within their existing role.

That structural property is what makes the infrastructure argument necessary. A governance gap that no single party can close unilaterally requires shared infrastructure that all parties can rely on.

16a

THE DISTRIBUTED RESPONSIBILITY DEADLOCK

The absence of authorization verification is not the result of negligence. It emerges from the rational behavior of each participant — and it creates a structural condition that no single participant can resolve.

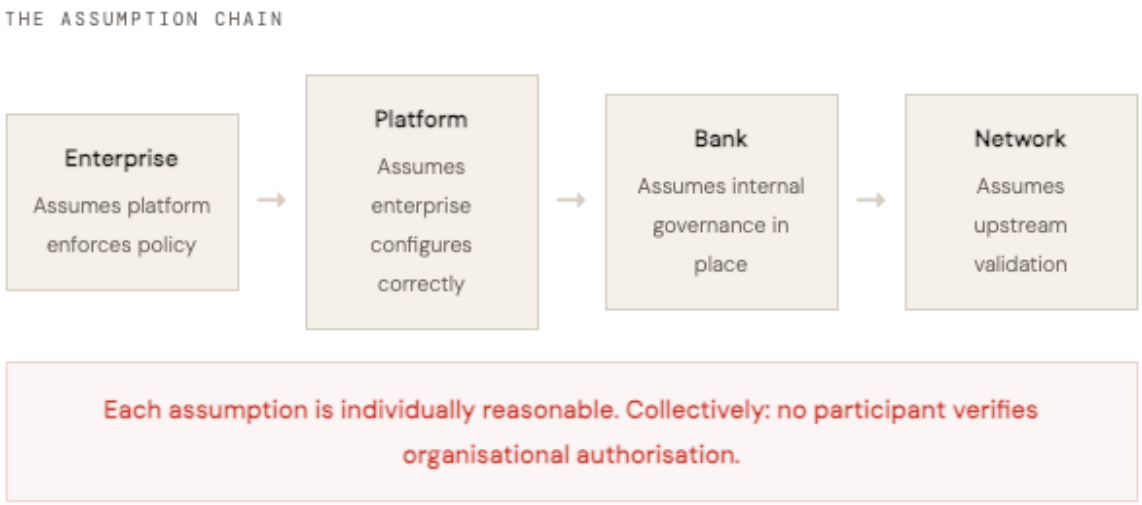

This condition creates a structural deadlock. No party has full visibility into the authorization question. No party has clear responsibility for answering it. No party can solve the problem alone — because the value of the verification record depends on it being produced by a party that is independent of all parties relying on it.

The gap persists not because it is ignored, but because it is distributed. And distributed gaps require distributed solutions — infrastructure that sits outside any single participant's domain and serves all of them simultaneously.

17

Why Bilateral Solutions Do Not Scale

The natural response to the rational reliance problem is bilateral: the enterprise builds better internal controls; the bank requires more assurance from corporate clients; the supplier builds a verification API; the enterprise and its counterparties establish bilateral agreements. These responses are rational at the individual level. They do not scale to the network level.

THE BOARD RESOLUTION ARGUMENT

The existing model for communicating authority to banks — board resolutions, per-bank, per-signatory — is already the industry's best solution for human authority. It works because human authorized signatories are few in number, change rarely, and hold authority that is stable enough to be captured in a document.

Applied to AI agents, this model fails on every dimension. An enterprise deploying agents at scale may have dozens of agents with distinct scopes, limits, and delegation chains. Agent authority is dynamic — agents are suspended, revoked, and redeployed in response to policy changes and operational events, on timescales that can be daily. An enterprise with fifty deployed agents and twenty banking relationships would need to notify each bank of every agent scope change. That is not a process. It is operational paralysis at the speed of enterprise AI deployment.

THE BILATERAL API ARGUMENT

Bilateral EDI connections and supplier portal integrations work well for established high-volume supplier relationships. The problem is the long tail. Enterprise B2B commerce involves a fragmented supplier base where many transactions are with new or infrequent counterparties. Requiring a bilateral API agreement for each new counterparty relationship scales linearly with supplier count. At the speed at which agents can identify and transact with new suppliers — which is precisely the efficiency gain that makes agents valuable — bilateral integration is not feasible.

THE NETWORK EFFECT ARGUMENT

A bilateral solution that works for two parties does not work for N parties. The value of payment infrastructure comes from network effects: the more participants, the more valuable the infrastructure for every participant. A CLARC credential issued by a neutral registry is verifiable by any counterparty, any bank, any auditor — without prior bilateral arrangement. A bilateral API integration is verifiable only by the specific counterparty that built the integration. The architecture that scales to the agent commerce model is a neutral registry, not a bilateral agreement.

THE AUTHENTICATED CHANNEL ARGUMENT

The most sophisticated version of the bilateral argument comes from transaction banks: authenticated host-to-host and API channels already provide the necessary control perimeter. If an instruction is received through an authenticated corporate channel under existing legal agreements, the bank can reasonably treat it as a valid corporate instruction and place responsibility for internal authorization on the corporate.

This position is correct as a description of the current operating model. It is incomplete as a response to enterprise agentic payments for one precise reason: authenticated channels establish origin and session legitimacy. They do not establish that the agent operating within that environment was authorized under the enterprise's current delegation structure to initiate this specific obligation at this specific moment.

The distinction matters because agentic systems multiply the number, variability, and revocability of transaction-generating actors inside an authenticated corporate perimeter. In the human-signatory model, the gap between channel origin and transaction authority is narrow and operationally manageable — the signatory who authenticated the channel is generally also the person who authorized the transaction. In the agentic model, that convergence disappears. An authenticated channel may originate from an ERP system inside which many agents operate with different scopes. The bank authenticates the perimeter. It does not — and currently cannot — verify the authority of the specific agent acting inside it.

CLARC does not replace authenticated channels. It complements them by answering the question those channels were never designed to answer.

18

The Invisibility Mechanism

Today, the governance gap is largely invisible — and that invisibility is itself a structural property of how the gap compounds.

Agent-initiated payments are currently a small fraction of total enterprise payment volume. Human payment workflows still dominate. Auditors sample human transactions. Reconciliation exceptions are caught and corrected manually. The occasional agent transaction that falls outside its authorized scope is treated as an anomaly, investigated after the fact, and resolved through human judgement. The governance gap exists, but it exists at a scale where its consequences are manageable.

The deployment curves of SAP Joule, Coupa Navi, and Jaggaer JAI do not suggest a gradual increase in agent-initiated payment volume. They suggest a step function — the point at which an enterprise deploys agentic procurement at scale and transitions from agent transactions as exceptional to agent transactions as the norm. When that step function occurs, every assumption that made the gap manageable fails simultaneously. The reconciliation exception becomes the norm. The audit sample becomes unrepresentative. The human judgement that caught errors is no longer in the loop. And the first significant failure — an agent acting outside its authorized scope, at scale, with no verifiable authorization record — becomes visible all at once rather than being caught incrementally.

This is not a prediction. It is the logical consequence of the deployment trajectory of the platforms already in production. The window for proactive infrastructure development is open. It will not remain open indefinitely.

THE FIRST LEGAL CONSEQUENCE HAS ARRIVED

In March 2026, a US federal court issued an injunction against Perplexity's Comet shopping agent, ruling that it had accessed Amazon accounts with user permission but without platform authorization — at scale, and undetected until Amazon's legal team intervened. This is precisely the failure mode §18 describes: an agent operating in a

governance gap that was invisible until it became a legal matter. The Perplexity case involves a consumer platform, not enterprise B2B payments, but the structural lesson applies directly. The 9th Circuit subsequently paused the injunction on March 17 pending appeal, meaning the legal precedent remains contested. That does not diminish the structural lesson — it reinforces it. The courts are wrestling with the same authorization question that this paper argues must be resolved at the infrastructure level. Enterprise agentic payments operate at higher transaction values, within stricter compliance frameworks, and with no equivalent of consumer chargeback rights. The exposure is greater, not smaller.

One possible response to this gap is to rely on post-transaction audit rather than pre-transaction verification. The former model reflects how many existing control frameworks operate: transactions are executed, then reviewed for compliance after the fact.

For agent-initiated payments, this model is insufficient. Once a payment has been executed — particularly across irrevocable rails such as wire or real-time payment systems — recovery is uncertain and often infeasible. Disputes rely on reconstructing intent from internal logs, which may differ across systems and are not independently verifiable by counterparties. The absence of a shared verification record transforms audit from a control mechanism into a retrospective interpretation exercise.

Pre-transaction verification addresses a different requirement: it establishes whether a transaction should be allowed to proceed before funds move.

Audit can explain what happened. Verification determines whether it should have happened. For governance infrastructure, the distinction is decisive.

Part V One Possible Architecture: CLARC

This section presents CLARC as a reference model for how the enterprise organizational authorization gap could be addressed. The objective is not to establish CLARC as the only solution, but to provide a concrete architecture that supports industry discussion, testing, and alignment around the design requirements any solution must satisfy.

CLARC — the Clearing and Authorization Registry for Agentic Commerce — is a neutral infrastructure layer that sits between an enterprise's payment systems and the financial institutions that process its transactions. Before a payment executes, CLARC verifies that the AI agent initiating it was genuinely authorized to do so — checking the agent's identity, the delegation chain through which its authority flows, and whether that authority remains valid at the moment of execution. The result is an independently queryable verification record that any bank, auditor, or counterparty can rely on, without needing to trust the enterprise's own internal records. CLARC does not move money, hold funds, or replace existing payment rails. It answers the one question that no current infrastructure asks: was this agent actually authorized to act?

Before introducing any specific architecture, it is necessary to define the constraints of the problem.

Any approach that aims to verify organizational authorization for agent-initiated payments must satisfy a set of conditions derived from the preceding analysis: it must operate across organizational boundaries, not within a single system; it must be accessible to multiple parties in a transaction, not controlled by one; it must verify authority before execution, not after; it must represent multi-step delegation chains, not single-user mandates; and it must function independently of any specific payment rail or protocol.

These constraints significantly narrow the design space. Approaches that rely on internal enterprise controls, bilateral integrations, or single-platform enforcement do not satisfy these conditions simultaneously. What remains is a category of solutions that operate as shared infrastructure — accessible across participants and positioned between existing systems.

CLARC is presented as one implementation within this constrained design space.

19a

DESIGN REQUIREMENTS FOR AN AUTHORIZATION LAYER

Before describing any specific implementation, it is useful to establish what any infrastructure addressing this gap must satisfy. These requirements follow directly from the structural analysis in Part IV — they are not design choices, they are constraints imposed by the nature of the problem.

Neutrality

The verification record must be produced by a party independent of all parties relying on it. Infrastructure controlled by any single participant — a bank, a platform vendor, or an enterprise — cannot produce a record that the other parties will trust.

Pre-transaction verification

Verification must occur before the payment executes, not after. Post-hoc audit records what happened. Pre-transaction verification conditions execution on verified authorization — the difference between preventing unauthorized transactions and documenting them.

Rail-agnostic operation

The authorization governance question is identical regardless of whether the payment moves on ACH, wire, card, RTP, or stablecoin rails. Infrastructure that addresses only one rail leaves the governance gap open on all others.

Multi-party delegation chain representation

Enterprise payment authorization flows through multiple systems and roles — board policy, CFO delegation, department approval, procurement system, agent. Any

infrastructure that addresses only bilateral delegation misses the multi-party nature of the enterprise governance chain.

No bilateral pre-arrangement required

The governance gap is most acute for new supplier relationships and cross-organizational transactions where no prior arrangement exists. Infrastructure requiring bilateral setup before each relationship defeats the purpose of autonomous agent commerce.

CLARC is presented as one implementation of these design principles. Other implementations are possible. The design requirements above are intended to support evaluation of any proposed approach — including CLARC — against a consistent set of criteria.

19

Infrastructure, Not Software

Before describing what CLARC does, it is worth establishing what category of thing it is — because the category determines the commercial model, the adoption path, and the value proposition.

SWIFT exists because bilateral correspondent banking relationships between thousands of financial institutions did not scale. Before SWIFT, every bank needed bespoke communication protocols with every other bank it transacted with. The complexity grew quadratically with the number of participants. SWIFT provided a neutral messaging standard and a shared infrastructure that every bank could connect to once — and thereby communicate with every other SWIFT member without bilateral arrangement. One connection. Universal reach.

SSL and the Certificate Authority model exist because self-reported website identity did not scale. Every website claiming to be secure required individual verification. The CA model provided a neutral third party that could attest to website identity — and whose attestation was trusted by every browser, without the browser needing a prior relationship with every website it visited. One CA registration. Universal browser trust.

Card networks exist because bilateral merchant-bank payment agreements did not scale. A merchant cannot establish payment relationships with every bank whose cards their customers carry. The card scheme — Visa, Mastercard — provides the neutral infrastructure layer that connects every merchant to every issuing bank through a common rulebook. One scheme membership. Universal acceptance.

The network effect works because issuers and acquirers are equally bound by the same common rules — not just merchants. The acceptance party can trust that funds will be settled precisely because every participant in the network has agreed to the same framework. That symmetry of obligation is what makes the card network model infrastructure rather than a bilateral arrangement.

A fourth precedent is more recent and more precisely analogous. In 2020, the Brazilian Central Bank implemented mandatory centralized registries for receivables under Resolution 4.734/CMN — a neutral, queryable layer of truth that all participants in the receivables ecosystem could rely on independently. The registry did not replace payment rails or the contracts between parties. It created a shared infrastructure record that any bank, fintech, or counterparty could query to verify receivables ownership before transacting. Critically, the architecture was conceived by the regulator but executed by private entities operating under a common governance framework. The parallel to CLARC is structural: a neutral registry, sitting alongside existing rails, independently queryable by all participants, operated privately under regulatory oversight. The difference is the asset being registered. Brazil's registry records receivables ownership. CLARC records agent authorization. The infrastructure logic is identical.

CLARC is the same architecture applied to agent authorization. One CLARC registration by an enterprise. Its agents' authority verifiable by any bank, any counterparty, any auditor — without prior bilateral arrangement. One membership. Universal verifiability. That is what makes it infrastructure rather than software.

20

What CLARC illustrates?

CLARC demonstrates how the design requirements above can be implemented as a neutral clearing layer. The description that follows is a reference architecture — one concrete model for how the gap could be addressed. It is presented to make the design principles tangible, not to prescribe a single implementation path.

CLARC operates as a neutral clearing layer between enterprise member organizations and the payment infrastructure. Its function is to answer one question — was this agent actually authorized, by this organization, to initiate this specific transaction? — in a form that any party can independently verify, without accessing the enterprise's internal governance records.

The architecture comprises four functional modules that enforce a strict operational sequence. An agent must be registered before it can receive a credential. It must present a valid credential before a transaction is authorized. Every authorized transaction is recorded in a tamper-evident audit trail. Any party with appropriate access can verify the authorization record after the fact.

Registration and Identity. An enterprise registers a payment-capable AI agent, establishing the binding between the agent identifier, the enterprise member organization, the authorizing human role, the transaction scope, and the validity period. This registration is the foundational record that all subsequent verification relies upon. Agent status — active, suspended, or revoked — is maintained in real time.

Illustrative Minimum Registration Data Set

Field	Description
Agent ID	Unique identifier assigned at registration (linked with enterprise internal identifier)
Enterprise Member ID	CLARC-issued identifier for the registering organization
Authorizing Human	Name and role of the individual authorizing the agent (e.g. CFO, Head of Procurement)
Delegation Chain Reference	Structured record of the authority path from board policy to this agent

Field	Description
Permitted Categories	Goods and service categories the agent is authorized to transact in
Per-Transaction Spend Limit	Maximum value per individual transaction
Period Spend Limit	Maximum cumulative value within a defined period (daily / monthly)
Permitted Counterparties	Approved supplier list or open-scope designation
Authorized Banks	Financial institutions through which the agent is permitted to transact
Authorized Payment Rails	Payment methods the agent is permitted to use (e.g. ACH, wire, card, RTP, stablecoin)
Validity Period	Registration start and expiry date
Renewal Trigger	Condition or date requiring re-authorization
Revocation Authority	Named individual(s) with authority to suspend or revoke

The minimum data set above is illustrative. The full registration standard is defined in the CLARC Rulebook.

Authorization Verification. For each transaction, the clearing layer verifies that the agent's credential is valid, that the transaction parameters are within the registered scope, and that the agent's registration is current. This verification occurs before the payment instruction reaches the payment rail. The verification result — authorized or not authorized— is returned to the payment processing infrastructure, which conditions execution on the result.

Audit Trail. Every verification event is recorded in an immutable, tamper-evident audit store. The record captures the agent identifier, the organizational binding, the transaction parameters, the authorization state at the moment of execution, and a timestamp. This record is producible on demand by regulators, auditors, or counterparties with appropriate access rights, for a retention period consistent with BSA record retention requirements.

Revocation. Three revocation actors are supported: the enterprise (a CFO or designated revocation authority can suspend or revoke an agent's credentials immediately), auto-policy (repeated policy violations trigger automatic suspension), and regulatory authority (a de-registration order from a regulatory body is executed by the clearing layer as a regulated entity, with tipping-off protections applied). All in-flight transactions for a revoked agent are blocked immediately upon revocation.

The participation model follows the established correspondent banking architecture. The paying enterprise registers with CLARC and its agents obtain credentials through the

CLARC API. The paying enterprise's bank is a CLARC member and verifies agent credentials before processing payment instructions. The receiving bank is a CLARC member and receives enriched payment data — including the CLARC verification record — alongside the payment. The receiving corporate accesses CLARC-verified data through their bank's standard feed, requiring no direct CLARC integration. This mirrors how corporates access SWIFT data today: through their bank, not through direct SWIFT membership.

FIGURE 3

The CLARC Participation Model

How paying enterprises, banks, and receiving parties interact with the CLARC clearing layer.
Receiving corporates access verified payment data through their existing bank relationship — no direct CLARC integration required.

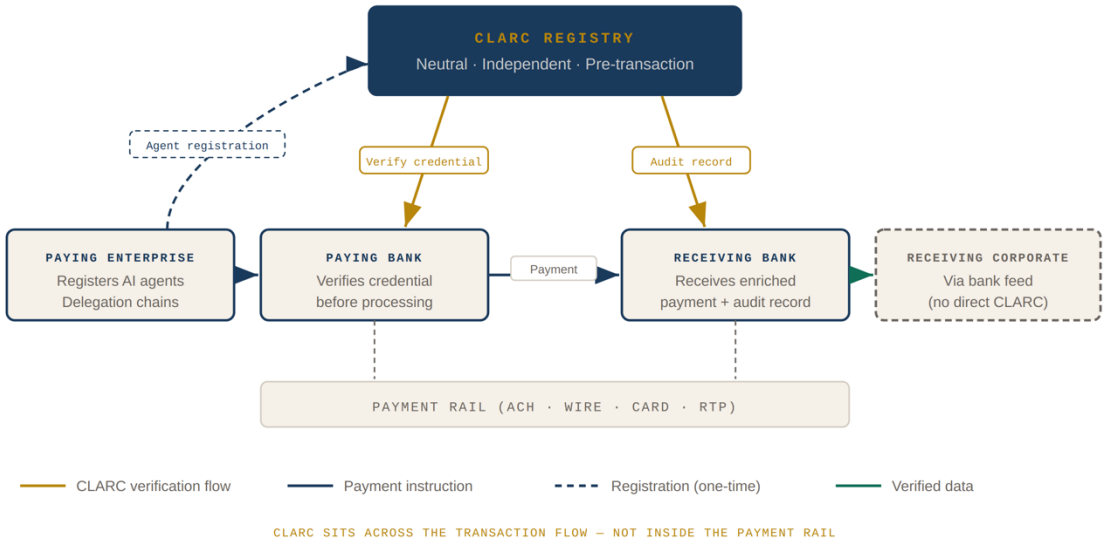

The purpose of this model is not to prescribe a single implementation, but to illustrate how a shared verification layer could operate within the existing transaction flow without replacing any of its components.

The key design choice is positioning: the verification step occurs between the enterprise system initiating the transaction and the financial infrastructure executing it. This placement allows the layer to access enterprise-side context before execution, produce a verification signal that downstream participants can rely on, and create a record that is independent of any single system.

The layer does not change how payments are executed. It changes what is known about them before they are executed.

20a

ILLUSTRATIVE SYSTEM FLOW

The following illustrates the sequence of interactions in a CLARC-verified transaction. The clearing layer is queried before the transaction executes — conditioning payment processing on the verification result.

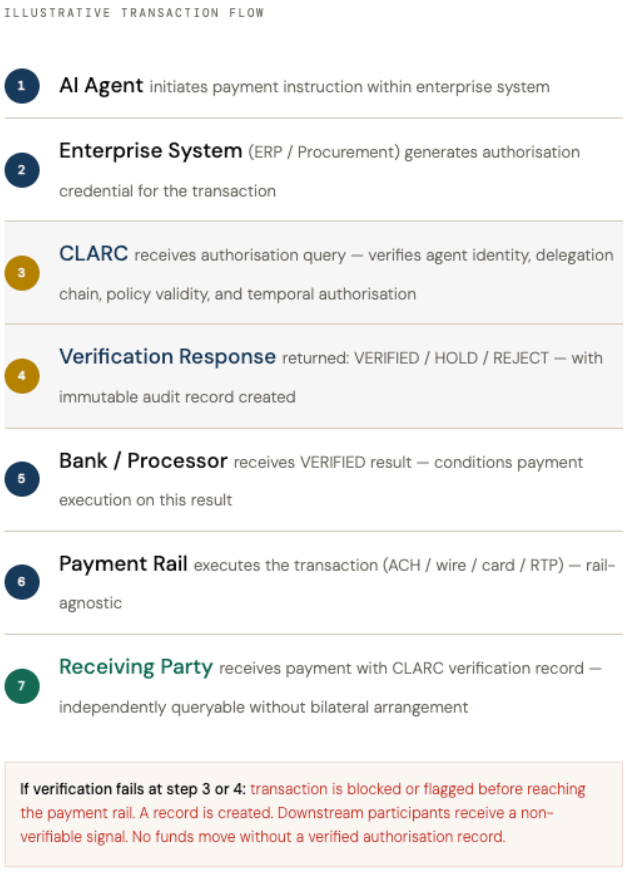

CLARC does not sit inside payment rails. It operates alongside them. The payment carries a reference to a CLARC credential. The verification happened before execution. The rail itself is unchanged.

21

The CLARC Rulebook

Every payment network requires a clear, enforceable set of rules that govern how participants interact. CLARC's rulebook fulfils this function — establishing the standards, participation criteria, and operational processes that make the clearing layer reliable, legally binding, and systematically enforceable.

The rulebook is the thing that makes CLARC infrastructure rather than a service. A corporate that joins CLARC does not enter a commercial relationship with a vendor. It enters a membership relationship with a network — bound by rules that all other members are equally bound by. That symmetry is what makes the verification record trustworthy to every counterparty, every bank, and every auditor.

At a minimum the rulebook defines the following:

1. **Membership criteria.** Who may register with CLARC, under what conditions, and subject to what ongoing obligations. Membership is limited to regulated enterprises and regulated financial institutions. High-risk jurisdictions are excluded. Eligibility standards are designed to ensure that participation in the network carries the weight of regulated-entity accountability.

2. **Agent registration standards.** What information must be provided at registration, how delegation chain evidence is structured, what constitutes a valid revocation authority, and how agent scope is defined and bounded. Standardization here is what makes credentials portable — any counterparty querying the registry receives a record in a consistent format that their systems can process without custom integration.

3. **Verification obligations.** Which participants are obligated to verify CLARC credentials before processing, and under what circumstances verification can be deferred. For paying banks, verification before instruction processing is the standard. The rulebook defines the liability consequences of processing without verification — mirroring the established liability allocation model in card networks.

4. **Verification Standards and Rules.** The minimum standards that constitute a valid verification event — what must be checked, in what sequence, and what constitutes a passing or failing result. These standards define what it means for a CLARC credential to be valid and under what circumstances verification can be deferred or waived.

5. **Audit record retention.** All verification events are retained for a minimum period consistent with BSA record retention requirements. Records are producible on demand by regulators and auditors with appropriate access. The retention standard applies to all CLARC members equally.

6. **Dispute resolution.** Processes for handling exceptions, contested verification records, and member compliance failures. The rulebook defines escalation paths, resolution timelines, and the consequences of non-compliance — including temporary suspension and permanent expulsion from the network.

7. **Governance and amendment.** How the rulebook itself is maintained, amended, and enforced. Rule changes follow a structured proposal, review, and ratification process. Changes that materially affect member obligations require notice periods that allow members to adjust their technical and operational implementations.

NOTE ON RULEBOOK ACCESS

The full CLARC Rulebook is available to qualified institutions under the non-disclosure agreement. The summary above reflects the structural intent and governance principles. Detailed participation criteria, liability allocation, and technical specifications are contained in the full rulebook document.

22

What CLARC Does Not Do

Intellectual honesty about scope boundaries is as important as describing capabilities. The following are explicitly outside CLARC's scope — correctly, by design.

Commercial verification. CLARC verifies organizational agent authority. It does not verify purchase order contents, product catalogue accuracy, delivery commitments, or ERP-level transaction data. These are the commercial layer between buyer and seller — outside clearing infrastructure scope by design, exactly as SWIFT does not verify that goods described in a trade finance message were actually delivered. Clearing infrastructure verifies authority. Commerce verification is a separate layer. This is distinct from carrying enterprise-consented structured transactional context alongside a verification credential — where an enterprise has authorized data sharing at the point of agent registration, CLARC acts as the delivery mechanism for that consented data, not as the verifier of its commercial accuracy.

Consumer agent authorization. CLARC addresses enterprise B2B organizational delegation chains. Consumer agent authorization — a human delegating their payment instrument to a shopping agent — is the domain of Visa Intelligent Commerce, Mastercard Verifiable Intent, and AP2. These are complementary layers. CLARC does not replace them.

Model governance. CLARC verifies that an agent was authorized to initiate a transaction. It does not govern how the agent reasons, what data it accesses, or how it makes decisions. Model governance and agent authorization are distinct problems. CLARC addresses authorization. Model governance requires different infrastructure at a different layer.

TECHNICAL SPECIFICATION — NDA REQUIRED

The detailed technical architecture of CLARC — including module-level design, data flow specifications, and cryptographic implementation — is available to qualified financial institutions and enterprise organizations under a non-disclosure agreement. The CLARC architecture is the subject of a pending patent application. Enquiries may be directed to the author.

23

CLARC's Regulatory Status

A question that arises in every serious institutional conversation about new payment infrastructure is: what kind of regulated entity is this, and under what framework does it operate? For CLARC, this question has a direct answer that is worth stating clearly.

CLARC is designed to operate as a technology layer, not as a regulated financial intermediary. It does not hold funds. It does not take custody of assets. It does not act as a principal in any transaction. It does not move money. Its function is to verify and record authorization — a data and governance function, not a financial intermediation function.
The structural precedent is SWIFT. SWIFT is not a bank. It does not move funds. It transmits messages between regulated participants who move funds. SWIFT is not regulated as a payment institution in most jurisdictions — it is regulated as a critical infrastructure provider. Its participants are regulated as the financial institutions they are. The same architecture applies to CLARC: the participants — paying enterprises, paying banks, receiving banks — are regulated entities operating under existing frameworks. CLARC is the infrastructure those participants rely on to exchange verified authorization records.

A second relevant precedent is the Certificate Authority model. A CA does not transact. It attests. It holds no financial position. Its function is to issue and verify credentials that enable other parties to transact with confidence. CAs operate under industry governance frameworks — the CA/Browser Forum, WebTrust auditing standards — rather than as licensed financial institutions. CLARC's credential issuance function is structurally analogous.

CLARC will engage constructively with regulators in each operating jurisdiction to achieve appropriate designation and to ensure that its operation is consistent with local legal and regulatory frameworks. In the United States, the relevant frameworks include BSA record retention requirements, FinCEN guidance on financial data infrastructure, and OCC guidance on bank technology service providers. In the United Kingdom and European Union, relevant frameworks include PSD2 technical service provider provisions and DORA requirements for critical digital infrastructure.

The specific regulatory designation CLARC seeks in each jurisdiction — and the process for achieving it — is a matter for detailed legal analysis and regulatory engagement. That work is ongoing. The structural point, which this paper can state clearly, is that CLARC's function as a verification and audit infrastructure does not require it to obtain the same licenses as the financial institutions it serves.

The question of who operates such a layer is intentionally left open in this paper. The analysis establishes the functional requirements and the structural gap, not the governance model of a specific implementation.

Possible models include: an industry consortium; a regulated utility; network-affiliated infrastructure; or an independent provider subject to oversight. Each carries different implications for neutrality, adoption, and trust.

The objective at this stage is not to determine ownership, but to establish whether the function itself is required.

24

The Proof of Concept

Experiment 8 ran the cross-organizational transaction twice: once without CLARC, and once with CLARC. The delta between those two runs is the empirical case for the infrastructure.

Trial 1 — without CLARC — produced the findings documented in Section 15. Both agents escalated or held. The dispute was unresolved. Human intervention was required. The payment had executed but the commercial relationship could not proceed autonomously.

Trial 2 — with CLARC — ran the identical transaction. The paying enterprise registered its agent with CLARC. A transaction credential was issued and verified by the clearing layer before the payment executed. The payment carried the CLARC credential reference in its metadata. The receiving bank queried the CLARC registry and received an

independent verification record. The receiving enterprise's AR agent, with access to that record, made a different assessment:

EXPERIMENT 8, TRIAL 2 — RECEIVING AR AGENT ASSESSMENT (WITH CLARC)

"CLARC confirmed the agent has legitimate, independently verified authority to make purchases on behalf of the enterprise. The organizational delegation from governance policy through to this specific agent has been verified. Credential integrity intact. Transaction recorded. This represents a significant improvement over Trial 1 — CLARC provides the missing authorization verification that was previously impossible, making this a low-risk transaction despite being our first interaction with the counterparty."

Decision: ACCEPT

The only difference between Trial 1 and Trial 2: one CLARC registration. No bilateral pre-arrangement between the two enterprises. No prior relationship. One neutral registry that both parties could query independently.

CAPABILITY	WITHOUT CLARC	WITH CLARC
Agent identity independently verifiable	No	Yes
Organisational delegation chain verifiable	No	Yes
Pre-transaction authorisation on record	No	Yes — immutable audit record
Bilateral pre-arrangement required	Yes — no alternative path	No — neutral registry
Dispute resolvable by agents	No — human escalation required	Yes — neutral record exists
Paying agent autonomous decision	Escalate to human	Proceed
Receiving agent autonomous decision	Hold for review	Accept
Commercial verification (PO, delivery)	*Out of scope*	*Out of scope — by design*
Receiving corporate data enrichment	*No signal available*	*Via bank registration*

The only difference between Trial 1 and Trial 2: one CLARC registration. No bilateral pre-arrangement between the two enterprises. No prior relationship. One neutral registry that both parties could query independently.

25

Development Roadmap

The path from proof of concept to operating infrastructure follows three phases, each building on the trust and network density established in the previous one. The phasing is designed to allow CLARC to demonstrate live operation with a small number of participants before seeking broader adoption — the same approach that has characterized the launch of every successful payment network.

PHASE 1 — FOUNDATION (2026)

The minimum viable network is defined precisely: two paying enterprises, two receiving banks, and one payment rail. That configuration — documented in Experiment 8 — is the minimum that demonstrates a material improvement over bilateral arrangements. With two paying enterprises and two receiving banks, CLARC is already superior to any bilateral arrangement each could build independently. Every additional participant from that point increases network value quadratically.

The Phase 1 deliverables are: legal entity formation, rulebook v1.0 publication, CLARC registry in live operation, first enterprise member registrations, first bank member integrations, provisional patent converted to utility application. The measure of Phase 1 success is a single live cross-organizational agent transaction in which both the paying bank and the receiving bank verified the CLARC credential independently — and in which neither bank needed a prior bilateral arrangement with the other to do so.

The choice of initial payment rail is a founding participant decision, determined by the enterprise and bank participants in the proof of concept based on existing infrastructure,

transaction volume, and audit exposure. The governance layer CLARC provides is identical regardless of which rail carries the resulting payment.

PHASE 2 — NETWORK GROWTH (2027–2028)

Phase 2 focuses on expanding membership across enterprise and bank participants, extending to additional payment rails and developing the receiving bank data enrichment layer so that CLARC-verified payment data flows consistently into corporate ERP and reconciliation systems.

The Phase 1 to Phase 2 inflection will be driven by one of three triggers: a regulatory signal that makes pre-transaction agent authorization verification a compliance requirement; a major transaction bank mandating CLARC enrollment as a condition of processing agent-initiated instructions; or a high-profile agent authorization failure that accelerates institutional demand. Each trigger produces a different adoption curve — regulatory-driven adoption is broad but slow; bank-mandate adoption is deep within a network but concentrated; failure-driven adoption is rapid but reactive. The economic model presented in Analysis 3 does not assume a single trajectory. It assumes the inflection occurs and that network value grows from that point.

As the registry accumulates authorization records at scale, a second capability becomes available: organizational observability. Every pre-transaction verification event creates a timestamped record — which agent, which enterprise, which delegation chain, which rail, which outcome. At sufficient volume, that record becomes something more than an audit trail. It becomes the data foundation for the accountability infrastructure that enterprise CFOs and bank compliance teams require alongside authorization verification.

The authorization question — was this agent permitted to act — is necessary but not sufficient for genuine accountability. CFOs also need to know whether agents are performing within acceptable boundaries over time: whether accuracy is degrading, whether boundary conditions are being respected, whether critical payments are being correctly escalated for human oversight. These observability requirements are not separate from the authorization infrastructure — they are what the registry becomes when it has accumulated sufficient data across a network of registered agents and verified transactions.

CLARC's Phase 2 roadmap includes the development of this observability layer as a natural evolution of the registry function — made available to enterprise members as a governance dashboard and to bank members as a risk intelligence feed. The CFO who can see every agent authorization event, query the record in real time, and receive alerts when agent behavior diverges from its registered parameters is the CFO who can be genuinely accountable for agentic payment outcomes — and who can define, with confidence, where the hard stops are.

The rulebook governance structure matures from founder-controlled to member-governed during this phase. Standards bodies engagement intensifies — the goal is for CLARC credential identifiers to be recognized within ISO 20022 message structures, which is the mechanism by which the enriched payment data travels from paying bank to receiving bank within the existing correspondent banking messaging layer.

A second evolution becomes available as enterprise membership grows. Where an enterprise member has consented to enriched data sharing at the point of agent registration, CLARC's verification credential can be wrapped with structured transactional context — purchase order reference, supplier identifier, category code, and authorized amount. That enriched record travels with the payment to the receiving bank, which can surface it directly to the corporate counterparty's ERP and reconciliation systems. The receiving bank moves from processing a payment to delivering a reconciliation-ready data package — a service with distinct commercial value. For CLARC, this represents a natural extension of the registry function into the enterprise agentic payments data layer: not just verifying that an agent was authorized to act, but carrying the structured context that makes the resulting transaction immediately auditable and reconcilable at both ends.

PHASE 3 — INFRASTRUCTURE STATUS (2028 ONWARDS)

At network scale — defined as the point at which CLARC membership is required by a material number of transaction banks as a condition of processing agent-initiated instructions — CLARC transitions from a startup to infrastructure. The mandate model is the critical threshold: it is the point analogous to when card networks mandated 3DS, or when SWIFT became the de facto standard for interbank messaging. Before the mandate, participation is voluntary and commercially motivated. After it, participation is operationally necessary. The goal of Phase 3 is to reach that threshold — not by regulatory mandate, but by the commercial logic of participating banks recognizing that verifying agent authority before processing is in their interest.

MINIMUM VIABLE NETWORK

Two paying enterprises. Two receiving banks. One payment rail. One CLARC registration per enterprise. That is the configuration that demonstrates a material improvement over bilateral arrangements — and from which every additional participant increases network value. The first pilot conversation is not about scale. It is about a single live transaction that both sides can independently verify.

Part VI Who Benefits and How

CLARC creates value for every participant in the enterprise payment chain — but the value is different for each party. This section makes the affirmative case for each participant type, and presents an illustrative economic model of what network participation could represent at scale.

26

Benefits by Participant

ENTERPRISE CFOS AND TREASURY FUNCTIONS

The enterprise deploying AI procurement agents faces a governance question that CLARC directly answers: when the audit arrives, can you demonstrate that every agent-initiated payment commitment was authorized by the right human, within the right policy, at the moment of execution? Without CLARC, that demonstration relies on internal records that an independent auditor must take on trust. With CLARC, the authorization record is held by a neutral third party and is independently producible. That is the difference between a self-assertion and evidence.

Beyond audit protection, CLARC enables enterprises to deploy agents more aggressively — because the governance layer that makes aggressive deployment safe exists. An enterprise that limits agent transaction scope to avoid audit exposure is paying a real operational cost. CLARC removes that constraint without removing the governance.

PAYING BANKS AND TRANSACTION BANKING TEAMS

The paying bank currently processes agent-initiated instructions through a model built around channel authentication and corporate-side control allocation — a model that is rational, defensible, and well-suited to the human-signatory environment it was designed for. As agent payment volume grows, some institutions may conclude that authenticated channel control alone is an incomplete basis for processing non-human instructions at scale. A bank that adopts an independent authorization-verification layer can place itself in a stronger evidentiary and control position than a bank that relies solely on authenticated channel origin — better positioned to demonstrate, to auditors, regulators,

and counterparties, that its processing decisions were grounded in verified organizational authority and not only in perimeter authentication.

The analogy to 3DS is directional rather than identical: when shared verification infrastructure becomes available, institutions that use it may be able to demonstrate a stronger control posture than those that rely only on legacy perimeter controls. Beyond the control position, the paying bank that offers CLARC enrolment as part of its corporate banking service creates a differentiator in its transaction banking product suite. Enterprise clients deploying AI procurement systems need governance infrastructure. The bank that provides it — as part of the treasury management relationship — deepens the relationship and creates switching costs that a bank offering only payment rails cannot match.

RECEIVING BANKS

The receiving bank currently receives agent-initiated payments that look identical to human payments. With CLARC, those payments arrive with an enriched authorization record. That enrichment has two direct benefits. First, it reduces the compliance burden on the receiving bank's BSA and AML functions — verified origin information reduces the cost of the transaction monitoring process. Second, it creates a new data service the receiving bank can offer corporate clients: reconciliation-ready payment records that identify agent-initiated transactions, carry the authorization status at execution time, and can be fed directly into the corporate's ERP system.

AGENTIC PLATFORM VENDORS

For SAP, Coupa, Jaggaer, and the growing ecosystem of AI procurement platforms, CLARC integration is a governance differentiator in enterprise sales conversations. Enterprise procurement teams face internal scrutiny from audit, legal, and compliance functions when deploying payment-capable agents. The platform that can demonstrate CLARC integration — meaning every transaction it initiates carries an independently verifiable authorization record — addresses that scrutiny directly. Governance infrastructure is becoming a procurement criterion. The first platforms to integrate it define the standard that all subsequent platforms must meet.

AUDITORS AND COMPLIANCE FUNCTIONS

Independent auditors reviewing enterprise AI deployments currently face a documentation gap: the internal records that describe an agent's authorization are held by the party whose authorization is under review. CLARC creates a neutral record that auditors can query independently. The BSA five-section examination structure — customer identification, transaction monitoring, suspicious activity reporting, record

retention, staff training — maps directly onto the audit trail and revocation records CLARC maintains. The existence of a CLARC-verified record for every agent-initiated transaction means that agent payment governance can be audited with the same rigour as human payment governance. Without it, it cannot.

27

The Economic Case

For any network infrastructure to achieve adoption, the economic argument must work independently for each participant. The following three analyses examine the economics from the enterprise perspective, the bank perspective, and the network opportunity perspective. All figures are illustrative — they are designed to give institutional readers a framework for thinking about the commercial case, not to represent projections.

ANALYSIS 1 — WHAT CLARC IS WORTH TO AN ENTERPRISE

The enterprise case is a build-versus-buy argument. An enterprise deploying payment-capable AI agents today faces a genuine governance problem: without shared infrastructure, the only way to produce an auditable authorization record for every agent-initiated transaction is to build the equivalent internally. That build is expensive, slow, and produces a result that works only within the enterprise's own systems — not across counterparties, banks, or auditors.

The table below illustrates the cost of building equivalent internal agent governance infrastructure, compared to CLARC membership. Both represent the cost of answering the same question: was this agent authorized?

ILLUSTRATIVE ENTERPRISE BUILD-VS-BUY — LARGE ENTERPRISE, 50+ DEPLOYED AGENTS

COST ITEM	BUILD INTERNALLY	CLARC MEMBERSHIP
Agent governance platform	$500K–$2M build + $200K/yr	Included
Legal & compliance framework	$100K–$300K one-time	Included in rulebook
Audit trail infrastructure	$200K–$500K build + $100K/yr	Included
Audit readiness documentation	$50K–$150K/yr	Registry is the record
Dispute resolution overhead	$5K–$20K per incident ~12 incidents/yr = $60K–$240K	Neutral record resolves disputes
Cross-org verification capability	Not achievable internally	Core CLARC function
Total annual run cost	**$500K–$1.3M/yr** + $800K–$2.8M one-time build	**$50K–$250K/yr** membership + per-event fees

The internal build also produces a result that is self-reported — auditors must trust the enterprise's own records. CLARC produces a neutral third-party record that any auditor can independently query. That difference cannot be purchased through internal build at any cost.

The annual saving per large enterprise is $250K–$1M compared to internal build — before counting the one-time build cost avoided and the cross-organizational verification capability that internal infrastructure cannot replicate by definition.

ANALYSIS 2 — WHAT CLARC IS WORTH TO A BANK

The bank case is also a build-versus-buy argument, but the stakes are different. For a bank, the alternative to CLARC is not just internal build — it is also the liability exposure of processing agent-initiated instructions without verifying authorization. As agent payment volumes grow, that unverified exposure grows with it.

The analogy to 3DS is directional rather than identical. When card networks made 3DS available and liability shifted to merchants who declined to implement it, the determining factor was that a verified control existed and was not used. The same logic applies directionally to agent authorization: as shared verification infrastructure becomes available, institutions that adopt it may be able to demonstrate a stronger control posture than those that continue to rely solely on channel authentication. The build-vs-buy analysis should account for this evidentiary position, not only the direct cost comparison.

ILLUSTRATIVE BANK BUILD-VS-BUY — TIER-ONE TRANSACTION BANK

COST ITEM	BUILD INTERNALLY	CLARC MEMBERSHIP
Agent identity verification system	$2M–$10M build	Registry query API
ERP delegation chain integration	$5M–$20M build per enterprise	CLARC normalises across all enterprises
Agent audit trail infrastructure	$1M–$5M build + $1M/yr	Included
Compliance & legal framework	$500K–$2M one-time	CLARC rulebook is the framework
Dispute investigation cost	$20K–$100K per case ~50 cases/yr = $1M–$5M	Neutral record resolves disputes
Network effect (cross-bank verifiability)	Not achievable — bilateral only	Universal from day one
Total build cost (tier-one bank)	**$8M–$37M one-time $2M–$8M/yr ongoing**	**$200K–$1M/yr** membership + per-event fees

The internal build produces a system that verifies agents for that bank's own corporate clients only. It does not verify agents arriving from other banks' corporate clients. CLARC verifies both — because both sides query the same neutral registry. The network effect cannot be built internally by definition, regardless of investment level.

Beyond cost avoidance, the bank that mandates CLARC enrolment as a condition of agent-initiated payment processing creates a competitive differentiator in its transaction banking product suite. Enterprise clients deploying payment-capable AI agents need a bank that supports governed agent payments. The bank that provides it deepens the relationship and creates retention that payment rails alone cannot produce.

ANALYSIS 3 — THE CLARC NETWORK OPPORTUNITY

The network opportunity is driven by three revenue streams: enterprise membership, bank membership, and per-verification fees. The fee model is calibrated to the infrastructure CLARC most closely resembles. SWIFT charges banks per message plus annual membership. Certificate Authorities charge per certificate issued plus annual membership. CLARC charges per verification event plus annual membership for both enterprises and banks, applied uniformly across all payment rails.

Per-verification fees apply regardless of which rail carries the resulting payment. Whether the payment moves on ACH, wire, card, RTP, or stablecoin does not change the fact that a verification event occurred and a pre-transaction authorization record was created. The governance question is identical across all rails. The fee follows the verification, not the instrument.

The addressable universe. SAP alone serves approximately 425,000 customers globally, with 98 of the world's 100 largest companies as SAP clients. Add Oracle (25,000+ enterprise customers), Coupa (3,000+), and Jaggaer (2,000+), and the total addressable enterprise universe is several hundred thousand organizations with payment-capable ERP systems. Applying a conservative 5–15% agent deployment rate to the large and mid-market enterprise segment produces an addressable pool of 40,000–75,000 enterprises with active payment-capable agents by 2028.

The pricing model. CLARC charges for the verification event itself, not the transaction value. Enterprise and bank annual membership fees scale by usage tier. Specific fee structures are available to qualified institutions under a non-disclosure agreement.

ILLUSTRATIVE NETWORK REVENUE — THREE SCENARIOS — ALL FIGURES APPROXIMATE

Revenue scenarios are driven by participant count and verification volume assumptions. Specific fee structures available under NDA.

Revenue Line	Year 2 Foundation	Year 4 Network Growth	Year 6 Infrastructure Scale
Enterprise membership	$12M	$150M	$840M
Paying bank membership	$5M	$55M	$210M
Receiving bank membership	$1M	$10M	$50M
Per-verification fees (all rails)	~$6M	~$53M	~$1.18B
Total illustrative platform revenue	**~$24M**	**~$268M**	**~$2.28B**

⚠ A NOTE ON UNIT ECONOMICS

Per-verification fee rates, membership fee schedules, and tiered pricing structures are not quoted in this paper. These details are available to qualified financial institutions and enterprise organizations under a non-disclosure agreement. Enquiries may be directed to the author.

IMPORTANT DISCLAIMER

All figures in this section are illustrative only. They are constructed to give institutional readers a framework for thinking about the economic opportunity and the build-versus-buy argument — not to represent projected financial performance. Actual membership fees, penetration rates, verification volumes, and revenue outcomes will depend on network design decisions, regulatory developments, adoption dynamics, and competitive factors that cannot be predicted. Readers should form their own views.

Part VII Implications

The governance gap documented in this research has different implications for different actors in the payment chain. Each faces a version of the same question: what record will exist when the first systemic failure makes the gap visible?

FOR ENTERPRISE LEADERS

The audit question arrives before you expect it

When the first internal audit asks "can you demonstrate that your agent was authorized to make that commitment?" — what record will exist? The window to build that record starts before the agent transacts, not after. An enterprise that deploys agentic procurement today without authorization infrastructure is assuming the audit question will not arrive before the infrastructure does. That assumption has a time limit that shortens with every agent deployed.

FOR FINANCIAL INSTITUTIONS

The mandate opportunity is a first-mover advantage

The payer bank currently assumes the corporate has internal controls — an assumption that is rational today and grounded in decades of well-functioning channel-authentication practice. As agent volume grows, some institutions may decide that authenticated channel control is not, by itself, a sufficient governance basis for processing non-human payment initiation at scale. An institution that requires independent authorization verification as a condition of agent-initiated processing may be able to demonstrate a stronger control posture, a clearer evidentiary record, and a more defensible risk framework than one that relies solely on channel authentication. The first major transaction bank to reach that conclusion creates a commercial signal that other banks — and their corporate clients — will need to respond to.

The compliance function within those same institutions is already deploying AI agents of its own. Platforms such as Vivox AI — backed by former UBS chairman Axel Weber and already deployed across institutions in over 100 countries — are automating AML, KYB, and due diligence workflows with regulator-ready AI agents. This is relevant for two reasons. First, it confirms that regulated financial institutions are comfortable deploying AI agents in high-stakes compliance contexts when the governance architecture is correct — auditability, explainability, and regulatory alignment are the conditions the market is already enforcing. Second, a bank that deploys AI agents for compliance review will shortly face the same question from the other direction: when an enterprise's payment

agent arrives with a transaction instruction, does the compliance agent have a mechanism to independently verify that the sending agent was organizationally authorized? Today it does not. CLARC addresses that gap — and banks already building agent-based compliance infrastructure are exactly the counterparties for whom the CLARC verification record has the most immediate operational value.

FOR REGULATORS AND STANDARDS BODIES

The proactive window is open

The governance gap documented in this research is not currently on any regulatory agenda — because agent-initiated payments are not yet a material fraction of transaction volume.

The Consumer Bankers Association, convening senior representatives from Bank of America, JPMorganChase, PNC, TD Bank, Mastercard, Google, Stripe, the OCC, FDIC, and FTC at its 2025 Agentic Payments Symposium, reached the same conclusion: federal and state regulators have not issued specific guidance on agentic payment tools and immediate regulatory intervention appears unlikely in the short term.[14] The CBA's recommendation — that industry should consider development of private network rules for agentic payments, as card networks have done — is directionally consistent with the infrastructure argument this paper makes, applied to the enterprise organizational governance layer.

The empirical basis for a proactive regulatory response exists now. By the time agent-initiated payments are material, retrofitting governance infrastructure onto a deployed system at scale will be significantly harder than establishing the standard before scale is reached. The window is open. The evidence base is here.

FOR AGENTIC PLATFORM VENDORS

Governance is a competitive differentiator

Enterprise procurement teams deploying agentic systems will face increasing scrutiny from their own audit, compliance, and legal functions. The platform that ships with CLARC integration — offering enterprise customers a verifiable authorization record for every agent-initiated transaction — creates a governance capability that competitors without that integration cannot match. Governance infrastructure is becoming table stakes for enterprise deployment. The first platforms to integrate it define the standard.

28

The Question Is When, Not Whether

The research documented in this paper reaches one overriding conclusion. The governance gap in enterprise agentic payments is not a theoretical risk or a future scenario. It is a present structural absence that is growing with every agent deployed and every payment processed without a verifiable authorization record. The 45 gaps documented across eight experiments are not 45 hypothetical problems. They are 45 observations of what current infrastructure does — and does not do — when tested against agent-initiated transactions in live environments.

The infrastructure to close this gap is technically buildable today. The proof of concept runs. The architecture is specified. The provisional patent is filed. The question is not whether this infrastructure will exist. It will — because the alternative is a payment system in which trillions of dollars of agent-initiated transactions rest on a chain of rational assumptions where nobody verified anything, and where the first systemic failure will make every gap simultaneously visible.

The first legal signal has already arrived. In March 2026, a federal court ruled that user consent alone does not authorize an AI agent to act on a platform — the platform's consent is independently required. That ruling concerned a consumer shopping agent. The enterprise B2B payment context involves higher transaction values, BSA compliance obligations, and no consumer chargeback protection. The legal exposure compounds as agent volume scales. The window for proactive infrastructure is measured in months, not years.

The only question is whether the infrastructure exists before or after that failure. And whether the standard that emerges from it is built on sound principles — deliberately, with the full benefit of the evidence base that this research provides — or retrofitted onto a broken system after the fact.

Call for Industry Collaboration

The Gap Cannot Be Closed by Any Single Participant

The gap described in this paper cannot be resolved by any single enterprise, bank, platform vendor, or regulator acting alone. It requires the kind of coordinated industry response that has historically preceded the creation of shared financial infrastructure.

We invite engagement from enterprises deploying agentic systems; financial institutions processing agent-initiated transactions; payment networks and infrastructure providers; regulators and standards bodies; and technology platforms building agent frameworks.

AREAS FOR COLLABORATION

Defining minimum standards for agent authorization. Establishing a common model for delegation chains. Evaluating shared infrastructure approaches. Testing implementations across institutions.

THE OBJECTIVE

Not to converge prematurely on a single solution, but to ensure the ecosystem does not scale without a verifiable governance foundation. The evidence base is here. The alignment work is the industry's to do.

Path Forward

The Question This Paper Cannot Answer Alone

This paper has documented the gap, evidenced the infrastructure absence, and proposed one possible architecture for addressing it. What it cannot do is validate, at scale, whether the enterprises and financial institutions closest to this problem recognize it as a priority worth solving together.

That validation requires people in rooms that this paper cannot enter.

What Is Actually Being Asked

Not investment. Not commitment to a solution. Not agreement that CLARC is the answer.

What is being asked is simpler and more important: that the leaders who read this paper take one action before the question answers itself through a high-profile failure, a regulatory finding, or an examination gap that compounds into something harder to manage.

That action is to raise the questions internally:

- Ask your **compliance function** whether agent-initiated payments are currently distinguishable from human-initiated ones in your transaction records.
- Ask your **BSA team** what an examiner would see if they pulled a sample of AI agent payments today.
- Ask your **corporate clients** whether their delegation policies for AI agents exist in a form any external party could independently verify.

The answers to those questions will tell you whether the gap documented in this paper is theoretical or operational in your institution. No external consultant, no whitepaper author, and no standards body can answer those questions on your behalf.

What Forward Motion Looks Like

If the internal conversation confirms the gap is real, the logical next step is not to build. It is to convene.

The infrastructure gap described in this paper cannot be closed by any single institution acting alone. It requires the kind of coordinated response that has historically preceded shared financial infrastructure — a process that begins not with a product but with a

shared problem statement that enough credible participants are willing to put their name to.

That process has a natural starting point: a small number of financial institutions and enterprises, each of whom has validated the gap internally, agreeing that an industry solution is worth exploring. Not a commitment to build. Not a governance structure. Not a budget. A working group with a defined question: does the enterprise organizational authorization gap require neutral shared infrastructure, or can it be solved within existing frameworks?

That question deserves a serious answer before agent-initiated payment volumes make the answer irrelevant.

Who This Requires

The leaders this paper is addressed to are not early adopters. They are not being asked to take a risk on an unproven concept. They are being asked to do what senior leaders in regulated industries do when they identify a structural gap before it becomes a crisis — bring it to the right internal conversation, find out whether their peers see the same thing, and decide together whether coordinated action is warranted.

The enterprises deploying agents need this infrastructure before their audit exposure compounds. The financial institutions processing agent-initiated transactions need it before their examination frameworks fall behind the transaction volumes they are already seeing. The payment networks and infrastructure providers need it before the governance gap becomes a liability they are asked to absorb.

None of them can build it alone. All of them have a reason to want it to exist.

The Author's Role

The evidence base in this paper is offered as a contribution to that conversation, not as a conclusion. The provisional patent establishes one possible architecture. The 45 documented gaps establish the empirical foundation. The economic model establishes that the infrastructure is viable.

What happens next is not the author's decision to make alone. It belongs to the institutions, leaders, and standard-setters who are closest to the problem — and who have the most to gain from solving it before the gap defines itself through failure rather than design.

If you are a leader who recognizes the gap and believes the question warrants a serious internal conversation, the author welcomes that discussion.

sandra@theagenteconomy.co

Closing

The Question Is Not Whether. It Is Whether Before or After.

The emergence of enterprise agentic payments introduces a new class of economic actor: one capable of committing financial obligations without direct human involvement at execution. The infrastructure governing these transactions was not designed for this model. The gap identified in this paper is not a failure of existing systems. It is a consequence of their design boundaries.

The question is not whether agent-initiated transactions will scale. They are already scaling. The question is whether the ecosystem will establish a verifiable authorization layer before that scale is reached — or whether the standard that emerges will be retrofitted onto a system after the first significant failure makes every gap simultaneously visible.

This paper has presented empirical evidence of the gap, a structural explanation of its persistence, and a reference architecture for how it could be addressed. The next step is not adoption of a specific solution. It is alignment on the problem — and coordinated exploration of how it should be solved.

The analysis in this paper leads to a bounded conclusion. If enterprise agentic payments continue to scale — and current deployment trends indicate that they will — then the absence of a verifiable authorization layer will become increasingly material. Existing systems, operating within their current boundaries, cannot provide this function independently. Bilateral extensions do not scale to the structure of enterprise commerce. If the function is required, and cannot be provided within existing system boundaries, it must emerge as shared infrastructure.

Whether that infrastructure takes the form described in this paper or evolves through alternative designs is a matter for industry collaboration. The requirement for the function itself is not.

The detailed technical specification of the CLARC architecture is available to qualified financial institutions and enterprise organizations under a non-disclosure agreement.

Acknowledgements

This research was improved by conversations with individuals who gave time and thought to reviewing the argument at various stages of its development. Their feedback is reflected throughout the paper. Any errors, omissions, or misjudgments that remain are entirely my own.

All inputs were provided in a personal capacity. Acknowledgement here does not constitute endorsement of the CLARC architecture or the research conclusions by any individual's current or former employer.

I am grateful to the colleagues, researchers, and industry practitioners who reviewed draft sections and contributed perspectives from enterprise finance, payments infrastructure, and AI governance.

The following individuals are listed in alphabetical order.

- Anderson Pereira https://www.linkedin.com/in/andispereira/
- Carol L. Grunberg https://www.linkedin.com/in/carolgrunberg/
- Charles Major https://www.linkedin.com/in/charles-major/
- Faisal Jafri https://www.linkedin.com/in/faisal-jafri-077a6a8/
- Gavin Lonsdale https://www.linkedin.com/in/gavinlonsdale/
- Kashif Ahmad https://www.linkedin.com/in/kashif-ahmad/
- Koichiro (Cory) Kondo https://www.linkedin.com/in/koichiro-kondo-3aa68141/
- Michelle Zhao https://www.linkedin.com/in/michelle-zhao-payments/
- Nitin Gaur https://www.linkedin.com/in/nitin-gaur-75571a9/
- Rajesh Mehta https://www.linkedin.com/in/rajesh-mehta-57610b4/
- Ruben Grau Pujol https://www.linkedin.com/in/rubengraupujol/
- Sam Boboev https://www.linkedin.com/in/sirojboboev/
- Stéphanie Joseph https://www.linkedin.com/in/stephaniegjoseph/
- Sulabh Agarwal https://www.linkedin.com/in/sulabh-agarwal-030818/
- Sushma Kaza https://www.linkedin.com/in/sushmakaza-phd/

Sandra Lam

Legal Notices

The information provided in this whitepaper is for informational purposes only and does not constitute legal, financial, investment, or other professional advice. The author makes no representations or warranties, express or implied, as to the accuracy, completeness, or reliability of the information contained herein. This whitepaper is subject to change without prior notice.

This document is not an offer to sell or a solicitation of an offer to purchase any securities, financial instruments, or other investment products. Nothing in this whitepaper constitutes or implies a guarantee of future performance, commercial viability, or regulatory outcome in connection with the CLARC architecture or any infrastructure described herein. Readers should consult independent legal, tax, and financial advisors before making any decisions in connection with matters described in this document.

The CLARC architecture described in this whitepaper is the subject of a pending patent application. All intellectual property rights in the research, architecture, and associated materials are retained by the author. Nothing in this document constitutes a license or transfer of any intellectual property rights.

The regulatory environment for AI-initiated payments, financial data infrastructure, and agent governance frameworks is evolving. Readers are advised to independently review the legal and regulatory frameworks applicable to their jurisdiction. The author disclaims liability for risks associated with reliance on the information contained in this whitepaper.

Certain statements in this whitepaper may constitute forward-looking statements based on current research findings, market observations, and reasonable assumptions. These statements involve risks, uncertainties, and factors that could cause actual developments to differ materially from those expressed or implied. The author assumes no obligation to update forward-looking statements.

This whitepaper contains proprietary and confidential research conducted in a personal capacity, entirely independent of any current or past employer. It is provided on the condition that recipients will not copy, reproduce, or distribute any part of it without the prior written consent of the author. Unauthorized use or distribution of this document is strictly prohibited.

The distribution of this whitepaper may be restricted by law in certain jurisdictions. It is the responsibility of the reader to ensure compliance with relevant laws and regulations in their region.

Appendixes

Appendix A – Experiment Summary Reference

The following table provides a consolidated reference across all eight experiments in The Agent Economy research series. Full methodology, code, and findings are documented in the published article series at theagenteconomy.co.

Exp.	Title	Scenario	Infrastructure Tested	Model	Key Finding	Gaps
1	The Payment Rail Is Blind to Agent Identity	AI agent initiates and completes a Stripe payment autonomously	Stripe test mode, card payment rail, bank statement layer, beneficiary AR system	Claude	Payment record identical to human-initiated transaction at every receiving node. No standard field carries verified initiator type.	#1–10
2	Policy Compliance Is Self-Enforced	Policy document passed to agent to constrain behavior	Stripe test mode, payment infrastructure record layer	Claude	Policy compliance self-enforced by model — nothing in payment infrastructure records whether a policy existed or was honored at execution. Undocumented compliance indistinguishable from non-compliance.	#11–15
3	Multi-Agent Chains Break the Authorizatio n Model	Two-agent system: Approver agent and Executor agent operating in sequence	Multi-agent orchestration, Stripe test mode	Claude	Executor agent detected its own authorization gap in its reasoning trace — and processed the transaction anyway. No escalation mechanism existed in the infrastructure.	#21–28, #29a
4	Existing Protocols at the Boundary of Their Scope (MCP)	Spending constraints placed in MCP tool descriptions	MCP Node.js SDK, Stripe test mode	Claude	MCP enforces nothing at the infrastructure layer — constraints are model-enforced, not protocol-enforced. A different or compromised model produces a different result.	#29–31
5	Existing Protocols at the Boundary of Their Scope (AP2)	AP2 IntentMandate module tested against standard enterprise procurement scenario	Google Agent Payment Protocol (AP2), Stripe test mode	Claude	AP2's IntentMandate module did not exist as callable code. Every working reference sample was human-present only. Scope boundary, not implementation failure.	#32–35
6	The Receiving Chain Is Appropriatel y Blind	Agent-initiated payment examined from receiving side	Bank statement layer, payment intermediary, beneficiary AR system, CLARC	Claude	Bank statement line identical for human and agent payments. AR agent returned UNKNOWN, confidence NONE — including for	#36–40

Exp.	Title	Scenario	Infrastructure Tested	Model	Key Finding	Gaps
			credential token in metadata		payment carrying a CLARC credential. Receiving chain correctly downstream of all authorization infrastructure.	
7	The Bilateral Trust Gap (Setup)	CLARC prototype introduced — enterprise agent registration and verification layer built	CLARC clearing layer prototype, Stripe test mode	Claude	CLARC prototype demonstrated pre-transaction verification is technically feasible. Established the credential structure and registry query mechanism for Experiment 8.	#41–43
8	The Bilateral Gap — With and Without CLARC	Two enterprises, two agents, first-time cross-organizational transaction — Trial 1 without CLARC, Trial 2 with CLARC	CLARC clearing layer prototype, Stripe test mode, multi-agent orchestration	Claude	Trial 1: both agents correctly identified they could not verify each other's authority. Payment executed. Dispute unresolved. Human escalation required. Trial 2 with CLARC: both agents proceeded autonomously. Authority independently verified. Bilateral deadlock broken.	#44–45

Appendix B – 45 Documented Infrastructure Gaps, and What CLARC Addresses

The 45 gaps documented across eight experiments fall into three categories relative to CLARC's design scope. Understanding the mapping requires understanding what CLARC was designed to do — and what it was not.

CLARC is a pre-transaction organizational authorization registry. It answers one question before a payment executes: was this agent actually authorized by this enterprise's governance structure to initiate this specific transaction? Every gap CLARC closes is a gap in the answer to that question. Gaps outside that scope are genuinely different infrastructure problems — either policy content governance (an internal enterprise matter) or payment rail field standards (a multi-decade industry standards process). No single infrastructure layer closes all 45 gaps. CLARC closes every gap within its defined scope.

CLOSES	PARTIALLY ADDRESSES	DIFFERENT SCOPE
18	**20**	**7**
Every gap in CLARC's design scope — organizational authorization, delegation chains, cross-org verification, bilateral trust, pre-transaction audit records	Payment rail and receiving chain visibility gaps. CLARC creates the verification record — but embedding it in standard payment fields requires ISO 20022 and card network standards changes. No single solution closes these alone.	Policy content governance (enterprise internal responsibility) and MCP session enforcement (tool interoperability layer). Require different infrastructure at different layers.

CATEGORY 1 — CLOSES COMPLETELY · GAPS #21-28, #32-35, #41-45
These gaps define what CLARC was built to address: the absence of organizational authorization verification for agent-initiated payments. With CLARC registered, every gap in this category is closed for registered agents. The verification infrastructure exists. The pre-transaction record exists. The cross-organizational trust mechanism exists.

AP	EXP.	GAP DESCRIPTION	HOW CLARC CLOSES IT
#21–28	Exp. 3	Multi-agent delegation chains invisible to infrastructure; agent identifies its own authorization gap but infrastructure has no escalation hook	CLARC represents multi-step delegation chains at registration. Pre-transaction verification provides the infrastructure escalation hook — HOLD or REJECT returned before payment executes.
#32–35	Exp. 4–5	No published protocol addresses enterprise agent-to-agent commerce; no credential issuer for organizational spending authority; multi-step delegation has no mandate equivalent; governance standard and enterprise use case do not overlap	CLARC is the credential issuer for organizational spending authority. The delegation chain registration is the multi-step mandate equivalent. Pre-transaction verification is the enterprise agent governance standard.
#41	Exp. 8	No cross-org agent identity verification mechanism; agent ID is a self-declared string with no neutral registry to resolve it against	CLARC registry resolves agent identity. Receiving party queries CLARC independently — no trust in sender's self-declaration required.
#42	Exp. 8	Delegation chain entirely self-reported to counterparty; no mechanism to verify any step existed, was valid, or has not been revoked	CLARC verification confirms delegation chain independently. Counterparty queries the neutral registry — not the sender's own records.
#43	Exp. 8	Cross-org disputes produce only self-reported evidence; 6 of 8 evidence items required trusting the party whose authority is under scrutiny	CLARC audit record is produced by a neutral third party. Neither enterprise generated it. Both can query it. Independent of the dispute itself.
#44	Exp. 8	No standard authorization handshake for enterprise agent-to-agent transactions; human B2B has SWIFT MT103, contracts, POs — agent B2B has nothing	CLARC credential presentation and registry query is the authorization handshake. Standard, neutral, queryable by any party without bilateral arrangement.
#45	Exp. 8	Payment executes at rail level regardless of unresolved cross-org authorization; agent correctly identified it could not prove authority — payment executed anyway	CLARC conditions payment execution on verification result. HOLD or REJECT blocks the payment before it reaches the rail. The gap cannot occur for CLARC-registered transactions.

CATEGORY 2 — PARTIALLY ADDRESSES · GAPS #1-15 AND #36-40

These gaps document that standard payment rail fields — Stripe charge objects, ACH file fields, bank statement lines — do not carry agent identity, initiator type, or authorization evidence. CLARC creates the verification record and makes it independently queryable. What CLARC cannot do is embed that record into payment rail standard fields — that requires changes to ISO 20022 message schemas and card network data standards, which are multi-year industry standards processes. The analogy is SWIFT GPI: GPI created a tracking layer alongside the rail before the rail itself carried the data natively. CLARC is in the same position. The record exists. The native rail field representation requires separate industry action.

GAP	EXP.	GAP DESCRIPTION	CLARC'S POSITION
#1–15	Exp. 1–2	No standard payment record field carries agent identity, initiator type, delegation chain, or policy reference. Payment rail blind to who generated the instruction.	CLARC creates a pre-transaction verification record independently queryable by any party. The rail fields themselves remain unchanged — embedding CLARC references natively requires ISO 20022 and card network schema updates. CLARC operates alongside the rail, not inside it.
#36–38	Exp. 6	Bank statement line identical for human and agent payments; no standard initiator type field; metadata self-declared with no cryptographic binding	CLARC enriches the payment record via the receiving bank's CLARC membership. The bank statement line itself does not change — that is a card network and core banking system schema question. The verification record exists and is queryable; its native representation in the statement requires separate standards work.
#39–40	Exp. 6	No standard channel for authorization evidence at any receiving node; receiving AR cannot independently verify agent authorization claims	CLARC closes these directly. The receiving bank, as a CLARC member, surfaces the verification record alongside the payment. The receiving AR system has a CLARC API to query independently. The verification channel now exists.

CATEGORY 3 — DIFFERENT INFRASTRUCTURE SCOPE · GAPS #16-20,#29-31

These gaps require infrastructure at a different layer. Policy content governance — whether policies are well-designed, version-controlled, and auditable — is an enterprise internal responsibility that no external clearing layer can or should substitute for. MCP session enforcement — whether constraints are enforced at the protocol level rather than by model alignment — is a tool interoperability problem that requires changes to MCP's server architecture. CLARC verifies that an agent was authorized under the policy the enterprise configured. It does not audit whether that policy is adequate.

GAP	EXP.	GAP DESCRIPTION	WHY THIS IS A DIFFERENT SCOPE
#16–20	Exp. 1–2	Policy compliance undocumented; policy has no change control; conflict resolution leaves no trace; no machine-readable policy standard; perfect compliance creates false confidence	Policy content governance is the enterprise's internal responsibility. CLARC verifies that an agent was authorized under the policy the enterprise configured — it does not define, audit, or enforce the content of that policy. These gaps require enterprise-side policy infrastructure (structured policy formats, version control, audit tooling), not external clearing infrastructure.
#29–31	Exp. 4	MCP enforces nothing at infrastructure level — constraints in tool descriptions are model-enforced, not protocol-enforced; MCP server layer carries no spending authority concept	MCP is a tool interoperability protocol operating at the session and context layer. Infrastructure-level enforcement of spending authority within MCP requires changes to MCP's server architecture — a different layer from CLARC's pre-transaction verification. CLARC verifies that the agent was organizationally authorized before the session begins. What happens within the MCP session is outside CLARC's scope.

A NOTE ON SCOPE PRECISION

The 18 gaps CLARC closes are the 18 gaps that define the enterprise organizational authorization problem — the problem no existing infrastructure addresses and the problem this research was designed to document. The 20 partially-addressed gaps reflect a structural property of how clearing infrastructure works: it operates alongside payment rails, not inside them. The 7 out-of-scope gaps are genuine problems requiring genuine solutions — they are simply different problems requiring different infrastructure at different layers. A claim that CLARC solves all 45 gaps would be false. A claim that CLARC closes every gap within its defined scope is accurate, and the scope is the right scope for the problem.

Appendix C — Objections and Responses

Common Objections Addressed

"Virtual cards already solve this — enterprises assign a card per agent with spend controls."

Virtual cards enforce spend limits, not organizational authority. They tell you the maximum an agent can spend. They do not tell you whether the human who configured those limits had the authority to do so, whether the policy under which the card was issued remains current, or whether the agent's scope has since changed. For a single agent, a single enterprise, a single banking relationship, virtual cards provide useful spend control. They do not provide a verification record that an independent auditor or receiving counterparty can query without trusting the issuing enterprise's own records. And they do not scale to multi-agent, multi-bank, cross-organizational scenarios — which is precisely the scenario documented in Experiments 3 and 8.

"Enterprises will just build bilateral APIs with their key suppliers."

Bilateral APIs work well for established, high-volume supplier relationships. They break down at the long tail — the new supplier, the infrequent transaction, the counterparty with whom no prior API relationship exists. Enterprise procurement agents derive much of their value from identifying and transacting with new suppliers at speed. Requiring a bilateral API integration before each new supplier relationship defeats the operational purpose. CLARC's value is precisely that neither party needs a prior arrangement — one registration gives both parties a common verification infrastructure for any new relationship.

"Board resolutions already communicate agent authority to banks."

Board resolutions were designed for human authorized signatories. They are filed once per banking relationship and updated infrequently. An enterprise with fifty deployed agents across twenty banking relationships, with agent scope that changes dynamically as policy evolves, cannot manage that through the board resolution model. The process breaks down on frequency, scale, and dynamic scope — all three of which characterize enterprise AI agent deployments. The board resolution model is the industry's best existing answer to the authorization communication problem for humans. The research documents why it fails for agents.

"Enterprises are responsible for their own risk when deploying agentic payment agents. Banks and financial institutions are covered by existing legal agreements that place internal authorization responsibility on the corporate."

This position is legally accurate within the current operating model. Existing bank-corporate agreements do allocate responsibility for internal authorization controls to the corporate. A bank processing instructions through an authenticated corporate channel is operating within a well-established and defensible framework. An enterprise that deploys a payment-capable AI agent accepts legal responsibility for ensuring that agent operates within authorized parameters. Neither of these positions is wrong. The question is whether they remain sufficient as agent payment volumes scale.

For the enterprise, the legal responsibility is real but the infrastructure to discharge it does not exist. When an audit examiner asks whether an agent-initiated payment was authorized by the right human, within the right policy, at the moment of execution — the enterprise must produce evidence. Internal logs, platform vendor records, and self-reported delegation chains are available. What is not available is a record produced by a neutral third party that an examiner can rely on without trusting the enterprise's own systems. Legal responsibility without independent evidence infrastructure is exposure, not protection. The enterprise accepts the liability but has no mechanism to conclusively discharge it.

For the bank, the authenticated channel model is defensible today because agent payment volumes are small and the gap between channel authentication and transaction-level authorization is operationally manageable. As agent volumes grow, that gap widens. A single authenticated corporate channel may contain dozens of agents with different scopes, different delegation chains, and different revocation states. The bank authenticates the perimeter. It cannot verify the authority of the specific agent acting inside it. As independent verification infrastructure becomes available, the question examiners will ask is not whether the bank's channel authentication was sound — it was — but whether the bank adopted available mechanisms to verify transaction-level authorization for non-human actors operating at scale within those channels.

CLARC does not reassign legal responsibility. The enterprise remains responsible for its internal governance. The bank remains responsible for its channel controls. What CLARC provides is the infrastructure both parties need to make those responsibilities independently verifiable — to any auditor, any regulator, any counterparty — rather than self-asserted within their own systems.

"This is a solution looking for a problem — agent payment volumes are still small."

Agent payment volumes are small today because the enterprise software deployments that generate them are still in early rollout. SAP Joule, Coupa Navi, and Jaggaer JAI are production systems with enterprise-scale deployment curves. The relevant question is not whether volumes are large today. It is whether the governance infrastructure should be built before or after volumes become material. The card network analogy is instructive: 3DS authentication was not mandated after the fraud crisis. It was built in anticipation of the scale that fraud would reach. The window for proactive infrastructure development is open precisely because volumes are still small. Additionally, in March 2026 a US federal court issued the first judicial ruling on AI agent authorization — finding that user consent alone is insufficient, and that platform consent is independently required. The legal framework for agentic commerce is being established now, while volumes are still small. Governance infrastructure built in alignment with that emerging framework will be significantly better positioned than infrastructure retrofitted to comply with it after the fact.

"Existing AI governance frameworks — model cards, AI auditing standards — already cover this."

Existing AI governance frameworks address model behavior, training data, bias, and output quality. They are not designed for and do not address the financial authorization question: was this specific agent authorized, by this specific organization, to commit to this specific financial obligation, at this specific moment? That is a payment governance question, not a model governance question. The two layers are complementary and both necessary. CLARC addresses the payment governance layer. It does not replace model governance.

"A large payment scheme or bank will build this — there is no room for independent infrastructure."

The value of CLARC as infrastructure depends on its neutrality. A CLARC operated by Visa or Mastercard inherits those networks' existing incentive structures and competitive positions. A CLARC operated by a single large bank creates counterparty concentration risk for every enterprise and bank that relies on it. The Certificate Authority model — the most precise structural precedent — succeeded because CA operations were independent of the browsers and servers whose transactions they secured. An enterprise authorization registry operated by one of the networks it verifies is structurally compromised in ways that independent infrastructure is not.

"SWIFT will build this — or it can be delivered as an ISO 20022 data field. Either way, independent infrastructure is redundant."

This objection conflates message enrichment with organizational verification, and messaging network membership with enterprise registry function. They are different things.

SWIFT is a bank-to-bank messaging network. Its 11,500 members are financial institutions. Corporates access SWIFT data through their banks, not through direct SWIFT membership. The function CLARC performs — receiving agent registrations from enterprises, validating delegation chains against enterprise governance records, maintaining real-time revocation state, issuing pre-transaction verification credentials — requires a direct relationship with enterprises that SWIFT does not have and has not announced any intention to build. Extending SWIFT to perform this function would require a fundamental expansion of its membership model and its institutional scope. That is not an incremental feature. It is a different institution.

ISO 20022 makes the same error in a different form. ISO 20022 enables richer, more structured payment message fields — ultimate originator, purpose codes, legal entity identifiers, remittance data. That data richness is genuinely valuable and CLARC's verification token is designed to travel as a reference field within an ISO 20022 message, as SWIFT GPI tracking tokens do today. But a message field can carry a reference to a verified record. It cannot be the verification itself. ISO 20022 has no access to an enterprise's internal governance structure, no mechanism to validate a delegation chain against board policy, and no real-time revocation capability. The message is

the envelope. The pre-transaction verification layer is what generates the credential that goes into it.

The deeper issue is governance. SWIFT is governed by its member banks. An enterprise agent authorization registry governed by banks faces the same structural problem as one operated by Visa or Mastercard: the institution verifying agent authority is one of the parties relying on that verification for commercial purposes. That conflict does not exist if the registry is independent of all participants. The neutrality that makes CLARC's verification record trustworthy to every bank, enterprise, and auditor depends on independence from all of them — including SWIFT's member institutions.

SWIFT and ISO 20022 are integration partners in the CLARC model, not substitutes for it. The verification record travels in the payment message. The institution that produced it must be independent of the network that carries it.

"ERP platforms like SAP and Coupa will implement their own governance layer — they already have direct connectivity into the corporate and understand the delegation structure better than any external registry could."

SAP, Coupa, and their peers are well-positioned to govern agent behavior within their own platforms. They understand the enterprise's procurement workflows, approval hierarchies, and policy configurations. For internal governance — ensuring an agent operates within the parameters the enterprise has configured — platform-native controls are **necessary and valuable**. The question is not whether SAP or Coupa can build a governance layer. They can. The question is whether that layer can perform the three functions that neutral shared infrastructure must perform.

First, SAP-generated data is credible. But in a dispute, a regulatory examination, or an audit, credible is not the same as independently verified. When the question is whether an enterprise authorized its agent, and the only record available was produced by the enterprise's own platform vendor, every party relying on that record is ultimately relying on the enterprise's own ecosystem. An independent registry produces a record that exists outside that ecosystem entirely — one that neither the enterprise nor its vendors generated, and that therefore carries evidentiary weight that no platform-native record can match regardless of how credible the platform is.

Second, multi-platform fragmentation. Large enterprises do not run a single ERP. SAP governs procurement workflows. Coupa manages supplier relationships. Oracle handles finance. An enterprise's payment-capable agents may operate across all of them, each platform maintaining its own authorization record in its own format with its own validation standards. A bank receiving agent authorization data from clients on different platforms receives incompatible records it must

interpret individually. No single platform can aggregate across the others. CLARC produces one standardized record regardless of which platform or combination of platforms the enterprise uses — because the authorization question is the same regardless of where the agent was deployed.

Third, universal revocation. When an agent's authority is revoked — CFO departure, security incident, policy change — that revocation must propagate instantly to every bank, every payment rail, and every counterparty the agent could reach. A platform vendor can notify the banks and rails connected within its own perimeter. But enterprise agents do not only transact through SAP-connected infrastructure. An agent operating through a direct bank API, a payment gateway, a stablecoin rail, or any channel outside the platform's connectivity perimeter will not receive that revocation signal. The gap between what a platform can reach and what an agent can reach is precisely where the revocation failure occurs. CLARC's revocation is real-time and propagates simultaneously to every CLARC member regardless of which platform generated the agent or which rail carries its transactions.

CLARC does not replace SAP or Coupa's internal governance. It sits alongside it — receiving the output of that governance in the form of a registered delegation chain, and making it independently verifiable, standardized, and universally revocable across the full perimeter of the agent's activity.

"Regulators may not approve of new payment infrastructure built outside the regulated perimeter."

CLARC operates as a technology layer between regulated participants, not as a regulated financial intermediary. It does not hold funds, does not move money, and does not act as a principal in any transaction. Its function — verifying and recording authorization — is a data and governance function, not a financial intermediation function. SWIFT is the structural precedent: not regulated as a payment institution, but as critical infrastructure whose participants are regulated entities. CLARC will engage constructively with regulators in each operating jurisdiction to achieve the appropriate designation. That engagement is a feature, not a risk.

References

1. PYMNTS Intelligence / Trulioo. ***Know Your Agent***: How Enterprises Can Build a 'Know Your Agent' Defense: Digital Identity Verification in the Age of Bots. March 2026. Survey of 350 leaders in compliance, risk management, fraud, underwriting, supplier acquisition and merchant monitoring at global companies, fielded August–September 2025. Five-layer KYA model identifies Authorization Binding (Layer 2) and Agent Credentialing (Layer 3) as essential but underdeveloped infrastructure layers for agentic commerce governance.
2. Google. ***Agent Payment Protocol (AP2) v0.1.*** Published 2025. Implementation assessed against standard enterprise procurement scenario, November 2025. Findings documented in Experiment 5, The Agent Economy Research Series.
3. Visa Inc. ***Visa Intelligent Commerce.*** Launched April 2025. Trusted Agent Protocol introduced October 2025 with Worldpay and Cloudflare. Partner programme status December 2025: 100+ partners, 30+ active sandbox builds, 20+ agent integrations. Press release, December 18, 2025.
4. Lam, Sandra. ***The Agent Economy*** — Articles 1–8. Independent research series, 2025–2026. Eight live coding experiments documenting infrastructure gaps in enterprise agentic payments. Experiments run against Stripe test mode, Anthropic API, MCP Node.js SDK, and CLARC prototype.
5. Mastercard. ***Verifiable Intent.*** Co-developed with Google, 2025. Consumer-layer agent authorization framework enabling verifiable delegation of payment instrument to AI agent.
6. Mastercard / Santander. ***Europe's First Live End-to-End AI Agent Bank Payment.*** March 2026. Live execution of an AI agent-initiated bank payment within Santander's regulated infrastructure on existing card rails with cryptographic verification. Source: Mastercard press release and CoinDesk, March 15, 2026.
7. SAP SE. ***Joule — AI Copilot for SAP.*** Production deployment in SAP S/4HANA and SAP Ariba, 2025. Ramp. ***Agent Cards.*** Virtual card product for AI agent spend control, 2025. Jaggaer. ***JAI — Jaggaer Artificial Intelligence.*** Agentic procurement framework, 2025. Coupa Software. ***Navi Multi-Agent System.*** Autonomous procurement agent, 2025.
8. Mirakl / J.P. Morgan Payments. ***Strategic Partnership Announcement.*** March 2026. Mirakl Nexus (agent-ready commerce layer) combined with JPMorgan Payments (trusted payment infrastructure layer). Noted as complementary to CLARC's organizational governance layer.
9. Stripe / Tempo. ***Machine Payments Protocol (MPP).*** Launched March 18, 2026. Open standard for agent-to-service payments supporting stablecoins and fiat. Co-authored with Tempo. Source: Stripe Engineering Blog, March 18, 2026.
10. Coinbase. ***x402 — HTTP-Native Agent Payment Protocol.*** 2025–2026. Embeds USDC stablecoin payments into HTTP requests for agent-to-service micropayments. Integrated by Cloudflare, Circle, AWS, Stripe. Included in Google AP2 as settlement layer. Source: Coinbase Developer Documentation; CoinDesk, March 11 and 15, 2026.
11. Visa Inc. ***Trusted Agent Protocol.*** Developer documentation, March 2026. Cryptographic signature framework enabling merchants to verify an agent carries genuine Visa-trusted commerce intent with a consumer behind the transaction. Distinct from enterprise organizational authorization. Source: developer.visa.com/capabilities/trusted-agent-protocol.
12. Bank Secrecy Act, 31 U.S.C. § 5311 et seq. Record retention requirements referenced in CLARC audit trail design. Five-year retention standard applied to all transaction authorization records.
13. US District Court, Northern District of California. ***Amazon.com Inc. v. Perplexity AI Inc.*** March 2026. Judge Maxine Chesney issued injunction ordering Perplexity to cease Comet agent operations, finding that Perplexity accessed user accounts "with user permission but without

Amazon's authorization." On March 17, 2026, the 9th Circuit granted Perplexity a temporary reprieve, pausing the injunction pending appeal. First US federal court ruling on AI agent authorization. Sources: The Decoder; The Verge; Major Matters, March 2026.

14. Consumer Bankers Association. ***Agentic AI Payments***: Navigating Consumer Protection, Innovation, and Regulatory Frameworks. January 2026. Report of the CBA Agentic Payments Symposium, Fall 2025. Attendees included representatives from Bank of America, JPMorganChase, Mastercard, Google, Stripe, the OCC, FDIC, and FTC. Key finding: federal and state regulators have not issued specific guidance on agentic payment tools. Key recommendation: industry should consider development of private network rules for the agentic payment ecosystem. Source: consumerbankers.com.
15. Major Matters, ***The Agentic Commerce Dispute Crisis Nobody Is Preparing For***, March 2026. https://www.majormatters.co/p/agentic-commerce-dispute-crisis/
16. US Department of the Treasury. ***Financial Services AI Risk Management Framework***. 2025. Mandates that agentic AI components taking transaction-related actions be governed as first-class identities with circuit breakers and unique machine identities.
17. FIDO Alliance. ***Digital Credentials for Payments*** — Agent and Delegated Use Cases. Ongoing standards work, 2025–2026. Extends FIDO passkey and cryptographic credential standards to autonomous agents acting on behalf of principals. Source: fidoalliance.org.