The Payment Infrastructure Layer: How Network Rules Already Govern Agentic Commerce

May 05, 2026

Categories:

Agentic Commerce Series
Publications

Authors:

Nilesh (Neal) Patel
John S. Wagster
J. Dylan Grafe
Katelyn E. Moody

The most important rules governing agentic commerce may already be written. They are not in any statute, court decision, or regulatory guidance document specifically addressing AI agents. They are in Visa and Mastercard network operating regulations, EMV 3-D Secure authentication specifications, stored credential frameworks, and the emerging protocol layer that payment infrastructure providers are building right now to manage agent-initiated transactions before any formal legal framework requires them to.

The previous article in this series established that a legally valid transaction can still be commercially blocked and that the access-control disputes in agentic commerce will center on who is permitted to participate at all. This article addresses the transaction-level questions that follow: ***For agents that do get through the gate, how does the existing payment infrastructure govern what happens next? Who bears the loss when it goes wrong? And why might the answers already be more settled and more consequential than most practitioners realize?***

The Classification Problem: Agentic Transactions Do Not Fit Either Bucket Cleanly

The foundational architecture of card payment networks is built on a binary distinction. A Cardholder-Initiated Transaction (CIT) is one where the consumer is actively participating at the moment of execution, authenticating, and authorizing the specific transaction. A Merchant-Initiated Transaction (MIT) is one the merchant executes later under prior authority, using stored credentials linked to an original authenticated transaction. Under Mastercard's Transaction Processing Rules, MITs require a Trace ID generated during the original CIT authorization, a cryptographic link proving that strong customer authentication took place during the setup phase. Visa's framework operates on a parallel architecture using an

Original Transaction Identifier (OTID) to establish the same lineage.

An AI agent-initiated transaction fits neither category cleanly, creating a tripartite agency problem where the agent acts neither as the consumer nor as the merchant, yet exercises delegated authority that produces binding financial consequences for both. When a consumer instructs an agent to monitor a marketplace and purchase a product when it reaches a target price, the consumer's active participation ends at the moment of delegation. Days or weeks later, the agent executes the transaction autonomously. No network rule currently creates a dedicated classification for this scenario. Treating it as an MIT requires an initial CIT from which a Trace ID or OTID can be drawn. But when the merchant is selected by the agent rather than by the consumer, no prior relationship between consumer and merchant may exist, making the initial CIT structurally impossible. Treating it as a CIT expects active human participation at execution, which triggers authentication protocols that no agent can satisfy.

As of April 2026, Visa's public rulebook has moved from indirect treatment to express language. The current Visa Core Rules now contain provisions for agentic transactions requiring identity verification in accordance with Visa Intelligent Commerce specifications and use of the provisioned token. Mastercard's public Transaction Processing Rules remain silent on agentic terminology. The network's movement on the topic appears in product programs such as Agent Pay, Agentic Tokens, and its Acceptance Framework rather than in the operating rules themselves. The practical consequence of that asymmetry is significant: For Visa-processed agentic transactions, there is now an express rule of decision. For Mastercard-processed transactions, the answer remains an inference from stored credential, authentication, and dispute frameworks that were not designed with agents in mind.

The ambiguity is not an absence of applicable rules. The rules exist. What is uncertain is how quickly networks, issuers, and adjudicators will map agentic fact patterns onto preexisting hooks, and who bears the loss during the period before that mapping is settled.

FBT Gibbons

Authentication Without a Human: Where 3DS Fails and What the Industry Built to Compensate

EMV 3-D Secure is the global protocol for authenticating consumers in card-not-present transactions and satisfying Strong Customer Authentication requirements. Its frictionless flow, introduced in 3DS2 and refined in version 2.3.1, is designed to authenticate consumers without requiring a manual challenge by using real-time risk assessment based on device telemetry, behavioral signals, browser configuration, IP geolocation, and up to 150 unique data elements analyzed by the issuer's Access Control Server.

When an AI agent executes a transaction through an API or headless browser, the frictionless flow fails structurally. The Access Control Server detects anomalous device telemetry: data center IP addresses, absence of behavioral biometrics, no screen interaction, no typing cadence. It defaults to the challenge flow to verify human presence. The agent cannot receive an SMS one-time passcode. It has no biometric hardware. The transaction times out or is abandoned.

The payment industry's current solutions to this problem are structural workarounds rather than purpose-built fixes. The 3DS Requestor-Initiated indicator allows a merchant or payment platform to initiate authentication for subsequent payments without the cardholder in session, provided the flow is properly characterized and linked to an original authenticated mandate. The Secure Corporate Payments exemption within 3DS 2.3.1 covers transactions by cards held by third-party agents or virtual cards. This category is intended for enterprise procurement rather than consumer AI shopping, but it's currently the closest available structural fit. Neither solution was designed for consumer-grade agentic commerce, and authorization rates vary significantly across acquiring networks depending on how strictly individual issuers implement their Access Control Servers.

The more durable solution is emerging at the protocol level. Visa's Trusted Agent Protocol, launched in late 2025, operates as a three-signature cryptographic handshake between the agent and the merchant: an Agent Recognition Signature proving the agent is

network-approved, a Consumer Identity object linking the agent to an existing consumer profile, and a Payment Container Signature serving as a cryptographic hash of the payment credentials. If the agent attempts to spend beyond the consumer's authorized parameters, the protocol layer rejects the transaction pre-authorization and creates an audit trail that functions as non-repudiable evidence for the merchant. Stripe and OpenAI's Agentic Commerce Protocol uses a Shared Payment Token scoped strictly to a specific merchant and cart total, isolating the consumer's payment credentials from the AI model entirely. Google's Agent Payments Protocol requires an Intent Mandate and a Cart Mandate, which are cryptographic signatures binding the transaction to the consumer's stated parameters and countering prompt-injection attacks that could alter the agent's purchasing instructions during the API handoff.

These are not consumer-protection statutes or judicial doctrines. They are private infrastructure choices that will determine how agentic transactions are classified, authenticated, and disputed long before any court or legislature resolves the underlying legal questions.

The Reg E Gap: Where Consumer Protection Breaks Down

The Electronic Fund Transfer Act, implemented by the Consumer Financial Protection Bureau (CFPB) as Regulation E, provides the primary federal consumer protection for disputed electronic payments. It defines an unauthorized transfer as one initiated by a person other than the consumer without actual authority to initiate the transfer. The definition was designed for a world in which a consumer either authorized a transaction or did not, collapsing the question of authority, error, and disputed authority into a binary unauthorized or authorized framework. Agentic commerce exposes the fragility of that collapse, because many disputed transactions will involve authorized access combined with disputed execution, rather than the absence of authority altogether.

Thus, agentic commerce creates a category Reg E did not contemplate: a standing authorization granted to an agent to initiate

transactions on a consumer's behalf, without transaction-specific consent. If a consumer instructs an agent to purchase a replacement filter and the agent purchases the wrong one at double the price, the consumer technically authorized the transaction by providing the access device and defining the general scope of delegation. Current legal interpretation suggests the transaction may not be unauthorized (thus, authorized) within the meaning of Reg E because the consumer authorized the agent's access. That leaves the consumer without the dispute rights they rely on for every other electronic transaction, not because of a deliberate policy choice but because the statutory framework was designed for a two-party world.

While the CFPB may have considered this in the abstract, it has not addressed it directly. Its August 2025 Advance Notice of Proposed Rulemaking under Section 1033 of Dodd-Frank is reconsidering who qualifies as a representative authorized to access consumer financial data, and whether that status should require a fiduciary duty to the consumer. The resolution of that question will affect which AI platforms can legally orchestrate agentic commerce. However, it does not resolve the underlying Reg E gap on misauthorized transactions (see part six in this series, "The Gatekeepers of Agentic Commerce") — those where the agent acted within configured rules but outside the consumer's real expectations — which is likely to emerge as one of the most important dispute zones in agentic commerce.

For financial institutions, the operational consequence is immediate. When a consumer disputes an agent-initiated transaction as unauthorized, the institution must classify the dispute under existing frameworks without clear guidance on which framework applies or how to treat a delegation that was real but may have been exceeded. As discussed earlier in the series, the institutions that develop internal frameworks for this classification problem before regulators or courts resolve it will be in a materially stronger position than those waiting for formal guidance. For a deeper examination of liability allocation and proof burdens when an AI-initiated purchase goes wrong, see part five, "When the AI Purchase Goes Wrong, Who Pays and Who Can Prove It?"

FBT Gibbons

The Stablecoin Gap: Where the Entire Governance Framework Disappears

One further gap deserves notice. Everything described above assumes the agent-initiated transaction runs on card rails. That assumption is not guaranteed. Some agentic commerce protocols are explicitly designed to be payment-rail-agnostic, and stablecoin settlement is a realistic near-term execution pathway for certain agent-initiated transactions.

A concrete example is already in the market: Coinbase's x402 protocol, which activates the long-reserved HTTP 402 Payment Required status code to enable automatic, on-chain stablecoin payments within standard website flows. The x402 protocol is explicitly designed for agentic commerce — allowing an AI agent to receive a payment request, submit a signed stablecoin payment, and complete the transaction without human involvement. It settles instantly. Unlike card network transactions, it contains no mechanism for reversals or chargebacks initiated by consumers, issuers, or merchants as a matter of right.

When an AI agent routes a transaction through x402 or a similar stablecoin rail rather than a card network, the entire governance architecture this article has been describing does not apply in a limited or modified form. It does not apply at all. There is no CIT/MIT classification. There is no 3DS authentication requirement. There is no chargeback mechanism. There is no Reg E unauthorized transaction protection, potentially reallocating risk away from intermediaries and entirely onto consumers.

However, this is not to say that stablecoin systems are incapable of such mechanisms. Many stablecoin platforms are built on reserves and retain issuer-level controls that allow transactions to be frozen or unwound in the event of fraud or misconduct. Currently, most stablecoin platforms require a government or law enforcement request to negate a transaction pursuant to court orders, law enforcement requests, or other forms of extraordinary governmental process. Those controls, however, operate outside the transaction flow itself. They are discretionary, ex post, and typically unavailable to consumers as a standardized dispute right.

They do not function as a reversal framework comparable to card network chargebacks, nor do they provide a defined evidentiary, timing, or liability allocation structure for resolving ordinary commercial errors or excesses of delegated authority.

The consumer who delegated purchasing authority to an agent that settled in stablecoins has none of the dispute rights they would have had if the same transaction ran through Visa or Mastercard. The GENIUS Act and related federal stablecoin legislation moving through Congress as of 2026 would, if enacted, impose reserve and redemption requirements on stablecoin issuers. None of this legislation creates a consumer dispute framework equivalent to Reg E. While similar safeguards could be embedded into agentic commerce rails, the challenge then becomes defining appropriate parameters to determine how and when those safeguards are enacted and may require resolving threshold governance questions such as when reversals are available, who may invoke them, what proof is required, and how disputed claims are settled.

For financial institutions and payment counsel advising clients on agentic commerce deployments, the payment rail the agent uses is not a technical detail. It is a threshold question that determines whether any of the existing consumer protection and dispute infrastructure applies at all.

What Financial Institutions and Payment Counsel Should Be Doing Now

The payment infrastructure issues in agentic commerce are not theoretical. They are arriving now as consumer adoption of AI shopping agents accelerates and as the first disputed transactions work their way through bank operations and dispute teams.

- **Audit your dispute classification framework for agent-initiated transactions.** Identify how your institution currently classifies a disputed transaction where the consumer provided credentials to an AI agent but disputes the specific purchase the agent made. Is it treated as an unauthorized transfer under Reg E? A billing dispute under Reg Z? A

chargeback under network reason codes? The answer matters because the consumer's rights, the institution's liability exposure, and its remediation obligations differ materially across those frameworks. Conduct this audit before your first volume of agent-initiated disputes arrives.

- **Assess your authentication infrastructure against the 3DS frictionless flow failure pattern.** If your institution operates as an issuer, evaluate how your Access Control Server handles transactions with device telemetry consistent with AI agent execution: data center IP addresses, absent behavioral biometrics, headless browser signatures. If your ACS defaults to challenge flow for these transactions and the challenge cannot be completed, you are declining a growing category of potentially legitimate agent-initiated transactions. Develop a policy for how to treat these transactions and whether to adopt 3RI or Secure Corporate Payments exemption treatment for recognized agent flows.
- **Review your platform and partner agreements for agentic commerce exposure.** If your institution sponsors merchants, acquirers, or payment service providers that are accepting agent-initiated transactions, your downstream agreements should address three questions: (1) who bears the loss when an agent-initiated transaction is disputed as unauthorized; (2) how is liability allocated among the issuer, the acquirer, the AI platform, and the merchant; and (3) what logging and evidence preservation standards apply? Network rules will govern the baseline, but contractual allocation within the network framework can materially affect where the first-hit loss sits. Institutions that address this in agreements now will be in a stronger position than those relying on silence when significant disputes arise. One structural asymmetry compounds this exposure: non-bank AI platforms operating as agent providers are generally not subject to Model Risk Management requirements that apply to regulated financial institutions. Banks contracting with these platforms cannot assume the AI developer has validated its models, tested its controls, or assessed its consumer protection risks at the standard the bank itself must meet. Third-party agreements

with non-bank AI providers should require model transparency, audit rights, and indemnification structures that reflect this disparity rather than assuming symmetrical risk-management obligations on both sides.

Bottom Line

The payment infrastructure layer of agentic commerce is not waiting for the law. Visa has moved from indirect governance to express rule text. Mastercard is governing through product programs while its public operating rules remain silent. The Agentic Commerce Protocol, the Trusted Agent Protocol, and the Agent Payments Protocol are establishing the private protocol architecture that will determine how agent-initiated transactions are classified, authenticated, and disputed. Reg E's unauthorized transaction framework has a gap that no regulator has yet addressed.

These are not abstract future problems. They are the rules of the current market, and the parties who understand them will have a significant advantage over those waiting for legislative clarity that may arrive long after the commercial architecture has hardened.

And once the payment infrastructure is understood as governance rather than mere plumbing, the next question becomes unavoidable: whether the parties controlling that infrastructure, and the conditions they impose on who may use it, are exercising market power in ways that antitrust doctrine will eventually need to address.

This article was prepared with the assistance of generative AI tools. The analysis, conclusions, and legal positions are the authors' own.

Agentic Commerce Series

- Part 1 — AI Agents Are Starting to Act Inside the Transaction, and Commerce Law Is Not Ready
- Part 2 — The Identity Problem: Authentication, Fraud, and Who's Actually Buying

- Part 3 — The Authority Problem: When Does an Authorized Agent Become an Unauthorized Buyer?
- Part 4 — Contracting by Agent: When the Agent Clicks, Who Assents?
- Part 5 — When the AI Purchase Goes Wrong, Who Pays, and Who Can Prove It?
- Part 6 — The Gatekeepers of Agentic Commerce
- Part 6 (A) — The Payment Infrastructure Layer: How Network Rules Already Govern Agentic Commerce
- Part 7 — Agentic Commerce, Price Control, and the New Antitrust Questions
- Part 8 — What Epic Reveals About Agentic Commerce

FBT Gibbons