The Gatekeepers of Agentic Commerce

Apr 28, 2026

Categories:

Agentic Commerce Series
Publications

Authors:

Nilesh (Neal) Patel
Katelyn E. Moody

By the time lawyers get to questions on identity, authority, assent, and loss allocation in agentic commerce, they are already inside the transaction. Those questions matter, and they are where this series on agentic commerce has focused so far. However, those questions assume something foundational that may be increasingly open to challenge: *that the agent was allowed to participate in the transaction at all*.

That assumption is becoming more difficult to rely on.

The next major question in agentic commerce is not just whether an agent-initiated transaction is valid once it occurs, but whether the actors controlling the transaction rails (merchants, marketplaces, platforms, payment providers, and protocol operators) will allow agents to transact on meaningful terms in the first place. That conflict is already in federal court. In *Amazon.com Services LLC v. Perplexity AI, Inc.*, Amazon sought to prevent Perplexity from using its infrastructure to power an AI shopping agent that would stand between Amazon's commercial surface and the end user. After the district court issued a preliminary injunction, the U.S. Court of Appeals for the Ninth Circuit stayed that ruling pending appeal, and the merits remain unresolved. While one case will not settle the field, the structural question *Amazon v. Perplexity* raises — *whether a dominant commercial surface can restrict or resist a new AI-mediated layer trying to insert itself between merchant infrastructure and the consumer* — is precisely the question this article addresses. In many cases, the first law of agentic commerce may be written less by appellate doctrine than by the integration rules, access restrictions, eligibility decisions, and private infrastructure choices that determine who gets through the gate.

The reality is that a legally intelligible transaction can still be commercially blocked. A system might authenticate the principal, define the agent's authority, preserve a record of assent, and still fail to matter if the merchant refuses autonomous buyers, the platform requires a privileged integration path, the payment layer

demands approved trust credentials, or the relevant commercial surface decides that only certain agents may enter at all. At that point, the controlling issue is not validity alone. It is participation.

The Market Is Already Building Controlled Participation Paths

Importantly, this is not a hypothetical future architecture. It is already visible in the way agentic commerce is being built.

OpenAI's commerce materials provide a clear example. Product-feed onboarding in ChatGPT is currently available only to approved partners. Merchants are directed into a documented integration path, required to provide structured catalog data, and are informed that OpenAI may remove products or ban sellers from being surfaced if platform policies are violated. Those are not minor implementation details. They are participation rules. They determine who may appear inside the shopping experience, through what integration path, and under what continuing conditions.

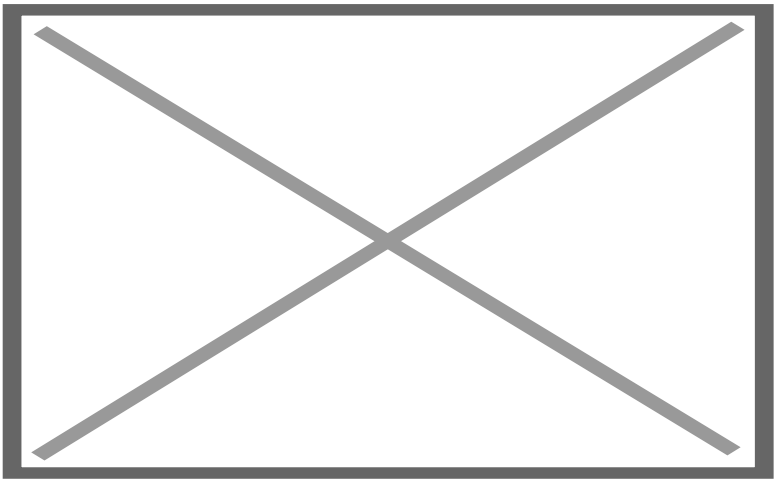

Figure: The Agentic Commerce Access Stack — a legally valid transaction can still be commercially blocked at any layer by a different private actor.

The protocol landscape forming around agentic commerce shows the same pattern at a broader level. There is no single gatekeeper, but there are multiple layers at which participation can be conditioned. Commerce protocols govern how products and orders are structured. Payment and trust protocols condition how

delegated credentials are recognized, what limits are attached to their use, and whether an agent-initiated transaction may be permitted and processed — a function already being built into Visa's Token Service and Mastercard's Delegated Authentication framework, both of which embed credentialing and authorization conditions at the network level before any individual merchant sees the request. Infrastructure protocols determine whether agent-to-system interactions are recognized as legitimate at all. A merchant might be willing to list products but unwilling to permit fully autonomous checkout. A platform might allow discovery but only through an approved commerce schema. A payment actor might support delegated credentials only when trust and scope conditions are satisfied. The effect is cumulative. The market is not building one gate. It is building a stack of them.

Consider what happens when a consumer delegates grocery purchasing to an agent with standing instructions to find the lowest price on a weekly list. The agent does not visit a single merchant's app. It queries multiple surfaces (whether through approved integrations, permitted automated access, or other mechanisms whose legitimacy is itself contested), selects the optimal combination of price and delivery, routes the transaction through a preferred payment credential, and notifies the consumer only after the order is placed. The merchant whose products were selected never controlled the search, the comparison, the checkout presentation, or the moment of decision. It received an order. Everything upstream of that order (the context in which its products were evaluated, the criteria by which they were chosen, and the customer relationship that might have been built) belonged entirely to the agent. That is a fundamentally different commercial relationship than the one merchants built e-commerce infrastructure to support. And it is exactly why some of the most important actors in the current market are trying to shape the conditions under which agents can participate before that shift becomes structural.

The Architecture Will Be Built Before the Law Arrives

The early rules of agentic commerce will not emerge all at once through legislation, regulatory activity, or through a coherent body of case law. They will emerge through terms of access, technical specifications, merchant onboarding standards, trust requirements, product eligibility rules, and the practical conditions imposed by the private actors whose systems an agent must traverse to complete a transaction. Those choices will decide what kinds of agents are tolerated, what authority signals are required, what commerce flows are acceptable, and when exclusion is justified.

That does not make the law irrelevant. It means, as with other waves of technology, the law will arrive after the commercial architecture has already been built and will have to work with the structure it finds rather than the structure it might have designed. As *Amazon v. Perplexity* already demonstrates, that toolkit (legal authority, platform terms, technical blocking, and access control) is already in use.

A merchant deciding whether to permit autonomous buyers is making an access-policy decision with downstream legal consequences. A platform deciding whether only approved partners may be surfaced in an agentic shopping flow is making a participation decision with competitive and contractual implications. A payment or trust provider deciding what credentials, limits, or verification conditions an agent must satisfy before its request will be honored is making a gatekeeping decision that may determine which models of agentic commerce can scale. Importantly, the legal significance of those choices does not depend on how the actors characterize them internally.

That timing gap is not merely structural — it is visible in where regulatory attention is currently focused. The Federal Trade Commission and other agencies are actively scrutinizing AI systems for deceptive practices, false capability claims, and undisclosed paid recommendations. Those are legitimate enforcement priorities. But they address the surface layer of the problem (whether an AI system is accurately described to consumers) rather than the architectural layer this series is mapping. Right now, regulators do not seem to be asking who bears the loss when a genuine delegated system produces a misauthorized transaction, whether platform

gatekeeping conditions constitute anticompetitive access control, or how consumer disclosure obligations function when no human is in the contracting moment. The regulatory gaze is on the edges of agentic commerce. The structural questions are at its core. By the time enforcement doctrine catches up to the architecture, the architecture will already have decided most of the answers.

The Relevant Actors Do Not Want the Same Thing

Part of what makes this issue so commercially consequential is that the major participants in agentic commerce are not aligned. That misalignment will drive the access-control disputes.

Merchants and platforms share a common instinct: they want the upside of agentic demand without surrendering control over how their products are presented, how customer relationships are mediated, or how much discretion a third-party system can exercise inside the purchase flow. A merchant may accept agent traffic while insisting that checkout occur only through an approved pathway, under merchant-defined conditions, with bounded use of product data. A platform may want standardization, governable integrations, and the ability to administer which merchants, products, and agents are surfaced in its environment. Both are trying to avoid being drafted into a transaction architecture they did not design.

The most consequential position belongs to vertically integrated actors — i.e., the platforms that are simultaneously building agents and controlling the commercial surfaces on which agents compete. Amazon illustrates the dynamic precisely. In *Amazon v. Perplexity*, Amazon sought to prevent a third-party AI agent from standing between its commerce surface and the end user, characterizing that intermediation as unauthorized appropriation of its infrastructure. At the same time, Amazon is actively developing Buy for Me, an AI agent designed to stand between consumers and other merchants' commerce surfaces. The legal rationale for each position may be defensible in isolation. Taken together, they reveal the underlying commercial objective: Amazon wants to control the agent layer that sits above other merchants' surfaces while preventing competing agent layers from sitting above its own. That

is not a legal position. It is a market-structure strategy and the kind of vertically integrated conduct that invites scrutiny under Sherman Act Section 2's leveraging doctrine and, if the foreclosure effects are sufficiently systemic, further scrutiny under essential facilities analysis.

Payment and trust providers see a different set of problems. Their interest is in whether the agent is acting under real delegated authority, whether payment credentials are being used within permitted bounds, and whether dispute handling remains operationally manageable. For these actors, gatekeeping often appears as trust infrastructure — credentialing standards, authorization limits, and transaction verification requirements embedded at the network level. But trust infrastructure still decides who gets through.

Agent developers, by contrast, have reason to prefer broader participation and fewer proprietary chokepoints. A developer building an agent that can compare merchants, optimize purchasing, and transact across multiple surfaces does not want to negotiate a separate commercial relationship with every actor in the stack. It wants interoperability, predictable access, and architectures that do not force it into a patchwork of bespoke approvals. From that perspective, private gatekeeping looks less like safety and more like exclusion.

If agentic commerce develops into a major transaction channel, these participants will not only dispute who bears downstream loss when transactions go wrong; they will also dispute upstream who sets the conditions of entry, whose integration path becomes standard, whose trust layer becomes mandatory, and whose control points become the default architecture of the market.

The Law Has Seen Versions of This Fight Before

The law has already seen disputes in which automated actors sought to interact with a digital surface, and the operator responded by asserting legal and technical control over access. The instructive example is *hiQ v. LinkedIn*, a case that arose from automated access to public profile data, not autonomous

purchasing. While it does not directly govern agentic commerce, *hiQ* exposed a recurring structural conflict: *whether a platform owner could treat automated participation on its surface as something subject to exclusion, technical blocking, and contractual restriction, even where the doctrinal path was contested*.

What *hiQ* establishes for present purposes is narrow but important. Even if automated access to public data may avoid certain criminal-law consequences under the Computer Fraud and Abuse Act, contractual and civil constraints remain highly relevant. Legal possibility does not eliminate private gatekeeping. A platform can still shape access through a combination of platform control, technical countermeasures, and contract. And courts will engage with those tools on their own terms rather than treating access as categorically open or categorically closed.

Agentic commerce raises a higher-stakes version of that pattern. The automated actor is no longer merely collecting information. It is attempting to participate in a transaction via payment credentials, triggering fulfillment, and potentially binding the principal to terms. When a new AI-mediated layer tries to position itself between users and an existing commercial surface, the incumbent actor may respond not just with technical blocking but with the full toolkit of access control: terms of service, API restrictions, eligibility requirements, and forced-pathway design. Agentic commerce is likely to produce the same gatekeeping instincts demonstrated in scraping disputes, only closer to the center of the transaction itself.

Three Forms of Access Control, and Why the Subtlest May Matter Most

The first and most visible form of access control is **outright exclusion**. Can a merchant or platform decide that autonomous buyers are simply not welcome unless separately approved? Can a marketplace refuse third-party agents while permitting only its own preferred commerce pathways? Can a platform declare that access by an agent outside designated interfaces violates its terms or undermines the safety and integrity of the commercial environment? Those questions may arrive dressed in technical or contractual language, but they are fundamentally questions about

participation.

The second form is **conditioned participation**. A merchant or platform may not exclude agents entirely but may require that they identify themselves, use approved credentials, comply with a prescribed protocol, submit only through an authorized API, respect product and merchant restrictions, or operate under platform-imposed trust and logging requirements. Those conditions may be commercially reasonable in isolation. The harder question — one that antitrust doctrine will eventually have to answer — is whether they become the mechanism by which powerful actors with interests in the agent layer shape which competing agent models are permitted to survive.

The third form is forced **pathway design**, which may prove the most consequential of the three. Instead of asking whether agents are allowed to transact in the abstract, the market asks through which pathway they are allowed to transact. A platform may say the agent cannot use the ordinary consumer surface but must come through the approved commerce protocol. A payments actor may say the transaction cannot proceed under general delegated access but must use a trusted-agent credentialing regime. A merchant may say autonomous ordering is acceptable only where the platform preserves particular records, exposes certain authority signals, or limits discretion in specified ways. In each case, the formal answer is not no. It is yes, but only this way.

That kind of conditional admission may prove more important than outright exclusion. It allows the market to adopt agentic commerce while preserving strong control over how the architecture evolves. It also lets powerful intermediaries, including those with their own agents in the market, determine which forms of agent participation are commercially viable. In practice, that may matter more than whatever broad legal principle a court eventually announces.

What Companies Should Be Doing Now

Companies building, enabling, or confronting agentic commerce should not wait for courts to resolve these questions before deciding how participation will work.

- **Create an agent-access map.** Produce a one-page document, owned jointly by product and legal, that identifies every point in the company's stack where agent participation can currently be admitted, denied, or conditioned — product interfaces, APIs, checkout flows, trust layers, payment authorization points, credential provisioning, and dispute-handling pathways. For each point, record the current default, the decision-maker, and whether the choice was deliberate or inherited. In many organizations, those decisions are already being made without coordination: product teams through API design, security teams through trust requirements, payments teams through authorization logic, and commercial teams through platform agreements. The map's purpose is to surface where real decisions are already happening and whether anyone owns them.
- **Convene a cross-functional access review.** Bring product, legal, payments, security, and commercial owners together around a narrow agenda: (1) does the business want open, controlled, or approved-pathway participation; (2) how are considerations relating to payment authorization, credential scope, and loss allocation guiding that choice; and (3) where and how is that choice currently being implemented in practice? A business that wants agents to transact freely should design for it. A business that wants agent access only through structured, trusted pathways should make those pathways explicit. A business that does not want autonomous buyers on critical surfaces should treat that as a deliberate policy choice, not a technical default.
- **Publish or internally adopt a written agent participation policy with minimum specified fields.** The policy should state at minimum: agent identification requirements, permitted transaction categories, spending scope and escalation thresholds, credential and trust conditions, dispute pathway, and the rationale for any exclusion or condition. If access depends on merchant approval, protocol compliance, trusted credentials, bounded payment authority, or preserved transaction records, those requirements should be visible. The

market will be harder to govern if the real conditions of participation remain implicit, and the legal risk of implicit conditions is higher than the legal risk of stated ones.

- **Distinguish between safety rationales and competitive rationales in a written decision memo.** Some participation controls will be justified by real concerns about fraud, user protection, loss allocation, and transaction integrity. Others will function more like commercial gatekeeping designed to preserve control over surfaces, data, or customer relationships. Many decisions will reflect both. Businesses should be explicit about which rationale is carrying the weight. The legal risk is not only in exclusion itself — it is also in characterizing a market-power decision as purely technical or purely protective. Under Sherman Act Section 2's leveraging doctrine, that characterization will matter.
- **Assume the architecture chosen now will be read later as the company's first real policy statement on agentic commerce.** The actors that define participation early (merchants, platforms, payment providers, and protocol operators) will shape the field long before courts produce settled doctrine. If those decisions are not made consciously, the market will still make them through fragmentation, default settings, and conflict after the fact.

Bottom Line

The next question in agentic commerce is not just over whether agent-initiated transactions are valid. It is over who gets to stand at the gate.

As agentic commerce grows, the actors controlling commercial surfaces, protocols, trust layers, and payments infrastructure will all try to shape who may participate and on what terms. The resulting rules may look technical, contractual, or operational. In substance, they will decide how much of the market is open, how much is controlled, and whose architecture becomes the path through which agentic transactions must run.

That is why the law of agentic commerce will not be written only in disputes over identity, authority, assent, and loss allocation after a transaction occurs. It will also be written upstream, in the quieter but more consequential decisions about access, integration, and exclusion that determine which transactions are allowed to happen in the first place.

And once the conditions of participation are set, the next question becomes unavoidable: whether those conditions themselves constitute a competitive problem, and who has the legal tools to challenge them. That is where we will turn our attention next in this series exploring the six fault lines of agentic commerce. If you haven't followed us up to this point, you can check out the first five articles below.

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