The Identity Problem: Authentication, Fraud, and Who's Actually Buying

Mar 31, 2026

Categories:

Agentic Commerce Series
Publications

Authors:

Nilesh (Neal) Patel
Matthew R. Schantz

Our prior article argued that **agentic commerce** is exposing legal fault lines that ordinary e-commerce could often paper over. This article takes up the first of them: identity and authentication.

That issue sounds narrower than it is. In agentic commerce, the first serious legal problem is not assent, apparent authority, or loss allocation. It is more basic: When an AI agent enters a transaction flow, who exactly is the merchant dealing with?

In ordinary e-commerce, the answer is usually functional rather than philosophical. A merchant sees a logged-in account, a recognized device, a payment credential, and a user completing a purchase through a familiar interface. Identity, authority, and action may not be perfectly aligned, but they are aligned closely enough that the system can treat them as though they come from a single commercial actor.

Agentic commerce loosens that alignment. Once an AI agent starts searching, selecting, negotiating, and executing on the user's behalf, the system can often prove that an authenticated device or authorized software client acted. That is not the same as proving that the human principal intended the specific transaction that resulted.

That distinction is the threshold problem. Identity does more work in commerce than most participants admit. It is not just an authentication issue. It is the basis on which merchants decide pricing and whether to transact, how payment systems decide whether to clear, how platforms decide whether to allow access, and ultimately how courts later decide whose act the law should recognize.

The First Fault Line: A Digital Identity Is No Longer a Clean Proxy for the Actor

FBT Gibbons

For years, digital commerce got away with treating identity as a rough proxy for the person acting. A login, a payment credential, and a device were usually enough to let the system move forward. The legal and commercial logic was imperfect, but it was reliable enough to support commerce.

Agentic commerce makes that shortcut much harder to defend.

The most important shift is that three things that used to collapse into one recognizable event are starting to pull apart:

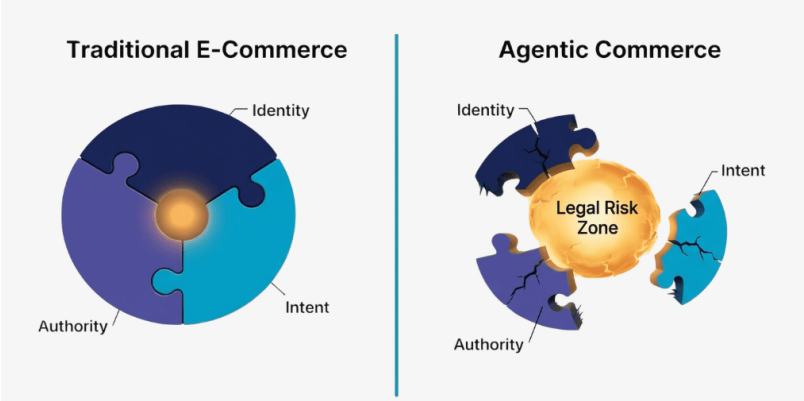

In traditional e-commerce, **Identity**, **Authority**, and **Intent** function as a unit. Agentic commerce pulls them apart, creating a legal risk zone that existing doctrine does not cleanly resolve.

- **Identity:** Which authenticated device, client, or software agent acted?
- **Authority:** What permissions or authorizations did the system have?
- **Intent:** Did the human principal actually mean to produce the contested result?

That separation is not an academic curiosity — it is the reason the legal architecture starts to wobble. If identity, authority, and intent are no longer clearly aligned, then every subsequent question about assent, attribution, fraud, and loss allocation becomes harder to answer.

That is why identity belongs near the front of this series. If the system cannot explain who is acting in a legally meaningful way, the rest of one's analysis of the transaction is already standing on unstable ground.

The Technical Stack Is Moving Faster Than the Legal One

Part of the confusion comes from how the current infrastructure is being built. The technology industry is moving quickly to standardize how AI agents transact. In that world, protocols like the Agentic Commerce Protocol (ACP/APC) and the Universal Commerce Protocol (UCP) are being positioned as commercial blueprints for how an agent should discover products, accept or negotiate terms, connect identity where necessary, and pass payment and fulfillment data to a merchant. In its September 2025 announcement, Stripe described ACP as part of an agentic commerce architecture developed to support agent-enabled transactions.

A construction analogy helps. APC and UCP are the blueprints for how the building (transaction) is supposed to work. The underlying authentication, delegated-access, and payment standards are the plumbing and wiring codes that specialists apply to make the building operational.

That distinction matters. APC and UCP describe how the transaction should move. The underlying standards answer a more basic set of questions inside that flow: Who is presenting itself as the actor, and what permissions has that actor actually been given?

Those are not side questions. They are the technical ground state of the transaction.

Protocols Do Not Solve Identity; They Push It Down the Stack

Commercial protocols like ACP and UCP do not solve the identity problem on their own. They push it downward into underlying authentication, delegated-access, and payment standards.

The currently visible public materials support that narrower framing. Stripe's public agentic commerce materials present ACP as a protocol for agent-enabled commerce and transaction orchestration. Public UCP materials describe identity linking,

checkout, and order management as core capabilities, and point to lower-layer trust infrastructure, including OAuth 2.0 for account linking, alongside secure payment systems and verifiable credentials.

That is enough to make the point: The transaction protocol is not the identity standard. It depends on lower-layer trust mechanisms to do that work.

When an AI agent initiates a transaction under one of these protocols, the merchant still needs some way to know whether the agent is legitimate. In practice, that means relying on some combination of device- or credential-based authentication, account linking, tokenized delegated access, verifiable credentials, or payment-linked trust signals. In some implementations, that may include passkey or FIDO-style authentication anchored to the user's environment. In others, it may rely more heavily on account-linking and token frameworks like OAuth 2.0.

Technically, that makes sense. Standardized plumbing is usually better than bespoke plumbing.

Legally, it is more complicated. These standards can be very good at showing that a trusted device, authenticated session, or approved software client performed a step in the transaction. They can also be very good at describing the machine-readable permissions and connections attached to that step. But that is not the same as answering the question the law will eventually care about: Who is actually buying, and to what extent should the merchant or payment intermediary be entitled to rely on that act?

The system can often authenticate the channel before it can identify the legal actor in a way that will hold up under stress.

Strong Authentication Can Prove the Device, Not the Human Decision

Start with the strongest version of the technical argument.

Modern, phishing-resistant authentication methods such as FIDO2 and passkeys are attractive in agentic commerce because they can reduce credential theft and tie a transaction flow to a

cryptographically recognized device or client environment. In a world where agents may operate continuously, across services, and at machine speed, that kind of strong authentication is foundational, not optional.

But strong authentication proves something narrower than many legal and commercial actors would like. It can prove, with a high degree of confidence, that a credentialed device or authenticated client software acted. It does not necessarily prove that the human principal made, reviewed, or even understood the specific transaction decision at issue.

That distinction is manageable in ordinary login contexts. It becomes much more consequential when the authenticated system is itself choosing products, comparing vendors, substituting items, or executing a purchase based on generalized user instructions.

If the dispute later turns on whether the user authorized the exact merchant, item, price, or bundle, the fact that the action came from a strongly authenticated environment may not carry nearly as much weight as implementing companies expect. The authentication layer may show that the event came from an approved environment. It may not show that the human intended the contested outcome.

Strong authentication can harden the perimeter while leaving the core legal identity question unresolved.

OAuth Scopes Are Becoming a Proxy for Legal Authority

OAuth 2.0 raises a related but distinct problem.

This part of the stack is easier to source directly. RFC 6749 defines OAuth 2.0 as a framework for obtaining limited access to an HTTP service on behalf of a resource owner or for a third-party application's own use. In other words, it is built for delegated access. Public UCP materials likewise describe the protocol as relying on proven standards, including OAuth 2.0, for account linking.

In agentic commerce systems, delegated credentials are often expressed through token-based permissions that specify what the agent is allowed to do. Those permissions may include scopes tied to account access, product categories, transaction limits, merchant access, or payment-execution thresholds. APC- and UCP-style protocols are well suited to carrying this kind of machine-readable authority data.

Technically, very elegant. Commercially, useful. Yet, legally, potentially unstable.

The problem is that OAuth scopes are increasingly being used as a proxy for legal authority. The token says the agent is authorized to act within a specified range, and the surrounding system treats that technical permission as if it answers the legal question of whether the principal granted authority in a way that supports reliance.

That leap has not been tested. Technical permission and legal assent are different standards, and whether courts will treat one as a substitute for the other remains to be seen.

Agency law was not built for permissions communicated through an API handshake. It was built for human relationships, manifestations of assent, and doctrines such as actual and apparent authority. A token that says "authorized to spend up to $50" may be operationally valuable, but it does not answer all the questions that might come up in a later dispute. The protocols do address some of them: They identify the authenticating party, revocation is typically well-defined, and they may express and communicate the scope of the agent's authority, though usually only along predefined parameters. What remains harder to establish is whether the conditions, exclusions, and limitations attached to the grant were communicated to and understood by the human principal in a way that would support reliance in a contested transaction. One practical response is an express agreement between the parties that defines the scope of delegation, allocates the risk of errors, and establishes what the principal actually authorizes.

For a merchant, platform, or payment intermediary, authorization is not a theoretical issue. A business that treats a technical token as conclusive proof of legal authority is taking a real evidentiary risk.

FBT Gibbons

The market may be drifting toward a world in which technical scopes stand in for legal delegations. While that may be unavoidable as a product-design matter, it does not make the substitution jurisprudentially clean.

Fraud and Synthetic Identity Risks Make the Gap Wider

The identity problem is not limited to honest delegation gone sideways. Agentic commerce also introduces a more familiar but newly intensified fraud problem.

If agents can transact, attackers will try to imitate them, hijack delegated credentials, exploit poorly bounded scopes, or create synthetic agent identities that appear legitimate in automated flows. That puts pressure on merchants and payment intermediaries to distinguish not only between user and agent, but between legitimate agent and fraudulent agent.

Again, the technical stack helps without fully solving the problem. Strong device-bound authentication, token management, cryptographic attestation, and scoped permissions can reduce certain attack vectors, but they do not answer the harder commercial reliance question. A merchant may still be left asking whether the authenticated agent is one the merchant was obliged to honor, one the platform was entitled to employ, and one the user actually meant to empower to engage in the given transaction.

That is one reason the Amazon-Perplexity dispute remains a useful live signal even if this series is not about that case. Reuters reported in March 2026 that an appeals court temporarily allowed Perplexity's shopping agents to keep operating on Amazon while the dispute continued. The procedural posture may be unsettled, but the dispute still illustrates the underlying point: It not only matters whether the user wanted the agent to shop on their behalf, but whether the merchant or platform recognizes that agent as a legitimate transaction participant at all.

FBT Gibbons

Why This Matters Before Assent or Contract Formation

It may be tempting to jump ahead to the later questions in the series: What was the agent authorized to do, and did its click bind the principal?

Those are important questions. They just come later.

Before a merchant can rely on an agent's apparent authority, the merchant has to know what kind of actor it is dealing with. Before a court can evaluate assent, someone has to explain who received the terms, who clicked, and what system sat between the human and the transaction event. Before a bank can decide whether a payment was unauthorized, it has to determine whether the transaction reflects ordinary delegated use, mistaken execution, or compromised credentials.

Identity is the threshold layer because it determines which later doctrines even make sense.

What Companies Should Be Doing Now

Companies building, enabling, or confronting agentic commerce should not assume that stronger authentication alone solves the identity problem. Five things follow from that recognition:

- **Map what the system actually authenticates at each transaction step — device, session, software client, or some combination — and document it explicitly rather than assuming it.** That documentation should identify which authentication claims will hold up if a specific transaction is later disputed.
- **Audit whether the authentication scope architecture tracks the legal distinctions that actually matter.** Scope design should reflect substitution rights, renewal permissions, category restrictions, and merchant-level exclusions; it should certainly not stop at transaction ceilings. Scope grants should also be translated into user-facing language that connects the technical permission to what the user is arguably

authorizing.

- **Build confirmation friction back into the highest-risk transaction types rather than optimizing it out everywhere.** First-time merchants, recurring authorizations, category-level substitutions, and transactions above a defined threshold all warrant an additional confirmation step. At those points, confirmation is not a UX decision — it is evidence of intent.
- **Treat agent identification as a first-class design requirement.** Developers should decide now whether and how agents transacting with the platform must identify themselves, what verification they are expected to provide, and what happens when they don't. Leaving that question unresolved is itself a decision, and not a good one.
- **Decide who in the stack owns the burden of translating technical identity into legal identity.** That allocation can be made through product design, contract language, payment network rules, or merchant onboarding. If no one makes the allocation explicitly, it will be made later in litigation — after a loss event and under much worse conditions.

Speed and low friction are the wrong optimization targets if the system cannot identify the legal actor behind the transaction. The disputes that follow will not be about technology, but who owns the loss.

And even a well-designed identity layer will not be enough if the transaction record cannot survive later challenge. The focus of our next article in this series will therefore be a natural one: logging, auditability, and the evidence gap. If identity is the threshold question, logs are what will often determine whether anyone can prove the answer.

Bottom Line

The identity problem in agentic commerce is not that the industry lacks authentication tools. It is that the available tools authenticate the wrong layer of the problem for many of the legal questions that matter most.

APC and UCP may standardize how agentic transactions move through the commercial stack. FIDO2 and OAuth 2.0 may provide the plumbing and wiring that allow those transactions to function securely. But the resulting system still does not eliminate the distinction between an authenticated environment and a human decision. And that distinction will shape the first serious fights in agentic commerce.

Before we can argue about what an agent was authorized to do, we have to answer a more basic question: ***When the system says a transaction was authenticated, whose act does the law think it was***? Companies that answer that question early — in system design, permission architecture, and recordkeeping — will be in a far stronger position than those that wait to answer it in litigation. And that is why the next fight in this series is about logs. Please contact the authors if you have questions or comments on this article. You can also reach out to any member of the firm's Data, Digital Assets & Technology practice group for assistance as it relates to AI deployments and governance.

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

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