
Proof of Intent Is Not Proof of Trust

Oct 13, 2026 | by Tim Brown



Why Agentic commerce needs risk decisions it can prove, before the dispute rules are written

Emerging standards can show what a consumer told an AI agent to buy. They cannot show whether that consumer was who they claimed to be, or whether the purchase should have been approved.

The industry has moved fast. Tokenized agent credentials, checkout protocols and network programs are live, and standards such as Mastercard's Verifiable Intent and Google's AP2 now record what a consumer authorized. What no standard captures is whether that authorization deserved to be trusted: whether the account had been taken over, whether the agent behaved as expected, and whether the purchase fit the customer.

That gap becomes a cost the moment a purchase is disputed. Without agent-specific rules, those disputes fall back to card-not-present rules written for a human pressing a buy button, and merchants absorb the loss by default. The 2026 holiday season will be the first high-volume test.

Kyance's position: **decide, then prove**. Every agent purchase should be approved on a real assessment of risk: whether the customer is legitimate, whether the agent is trusted, and whether the purchase fits what was authorized. That decision and its basis should then travel with the transaction in a portable, verifiable record. Networks, regulators and courts will eventually assign liability for agent purchases, and they will do it on the evidence the ecosystem already produces. The industry should agree on that evidence now.

The gap: payment rails shipped, adjudication did not

Every layer of the agentic stack has moved in the past eighteen months except one. Visa's Trusted Agent Protocol, Mastercard Agent Pay and Verifiable Intent, Google's AP2 mandates, and EMVCo's draft agentic payments framework all identify agents or record consent. As of October 2026, none of them sets a new dispute or liability rule ([Chargeflow](#)).

Four facts define the gap:

- **Liability defaults to the merchant.** The Agentic Commerce Protocol keeps chargebacks with the merchant and its payment provider, not the agent platform ([ShopAppy](#)). No US, EU or UK payments law is yet written for agent-made purchases, and whether Regulation E or Regulation Z authority tests cover an AI agent is untested.
- **The rules will not arrive in time for peak.** Industry analysis expects no binding Visa or Mastercard agent-dispute rule before the 2026 holidays, with April 2027 the earliest likely cycle ([ShopAppy](#)). Mastercard's own August 2026 report still treats liability and dispute resolution as problems the industry must solve together.
- **Legacy evidence no longer proves anything.** IP address, device fingerprint and session behavior describe the agent, not the person who authorized it. They cannot show whether a purchase fell inside the authority a consumer delegated.
- **Agent checks miss the human.** Visa's Trusted Agent Protocol and Mastercard Agent Pay verify the agent. Neither shows whether the person behind it is legitimate. A fraudster who takes over a real account, or opens a synthetic one, can delegate to a fully trusted agent and pass every agent check.

Scenario

Consider Tina, a hypothetical but entirely plausible shopper. She asks her AI agent to buy a laptop, and it completes a \$1,400 purchase at a mid-market merchant. Weeks later she disputes the charge, saying she never authorized it. The merchant fulfilled a valid order, but it has no record of the agent or what the agent was told to do. Nor can it answer the questions that matter most. Was it really Tina who set up the agent, or had her account been taken over? Did a \$1,400 laptop fit what she asked for and how she normally buys? Is she now disputing a purchase she did make? Without those answers, the dispute is effectively unwinnable, and the \$1,400 is gone.

Agent purchases are still a small share of commerce. Mastercard's own research puts consumer willingness to let an agent buy fully on its own at about 10% ([ShopAppy](#)). That argues for acting now, not later: the first rules will be written on whatever evidence exists when volume arrives.



Decide, then prove

Evidence alone only explains a loss after it happens. A stronger model assesses risk before a purchase is approved, then records what was checked. At the moment of purchase, three questions need answers:

1. **Is the customer legitimate?** Is this a verified person, and does the account show signs of takeover or fraud?
2. **Is the agent trusted?** Is it a known agent, run by an accountable platform, behaving as expected?
3. **Does the purchase fit the authority?** Does it match what the customer delegated and how they normally buy?

The answer is a risk decision with stated reasons. When those reasons travel with the transaction, an issuer reviewing a dispute sees not only that a purchase was authorized, but what was checked before it was approved. The record works in both directions: it makes a false “my agent did it, not me” claim hard to sustain, and it proves when an agent went beyond what the customer allowed.

This suggests a workable principle for future rules: **liability should follow verifiable evidence**. Whoever approves an agent purchase should be able to show what was checked. Networks have used this logic before. The EMV liability shift placed counterfeit fraud losses on the party that had not adopted the stronger, checkable standard.

What a dispute-grade record must prove

A dispute asks one question: did this person authorize this purchase? For an agent transaction, answering it credibly means linking six facts into one verifiable record.


1. **The human principal.** A verified customer stands behind the agent, not just an account or a device.
2. **The agent.** A specific agent, operated by a specific platform, carried out the purchase.
3. **The delegated authority.** The scope the person granted: merchants, categories, spending limits and time window, captured when it was granted.
4. **The transaction.** The actual items, price and merchant, bound to that authority at the moment of purchase.
5. **The risk decision.** What was assessed about the customer, the agent and the purchase, and the outcome.
6. **Validity at that moment.** The delegation had not been revoked or exceeded when the agent acted.

Delegation chain · four facts and a risk decision bound into one record

Each link in the chain is captured when it happens, then bound into one record that any party can check on its own.

The record also needs four properties to hold up in a dispute weeks or months later:

- **Signed and tamper-evident**, so no party can alter it after the fact.
- **Independently verifiable** by an issuer, acquirer or arbiter using published keys, without calling the party that created it.
- **Selectively disclosed**, so a merchant sees proof of authority without receiving more personal data than the dispute requires.
- **Durable** for the full dispute window, not just the checkout session.



None of this needs new science. Verifiable credentials, signed mandates and tokenized authority already exist. What is missing is agreement that they should be assembled into one artifact that every party accepts as evidence.

Why the record must be protocol-neutral

Merchants will receive agent orders through ACP, UCP, AP2, Visa Intelligent Commerce and Mastercard Agent Pay, each carrying its own consent and identity data. A record that works across all of them gives merchants one evidence playbook, gives issuers and networks one artifact to write rules against, and gives consumers proof when an agent overstepped.

Identity has done this before: authentication scaled on the internet when FIDO standardized one proof that any relying party could check.

What the industry should do now

The rules will come. The question is whether they are written on evidence or on guesswork. Each party can act before the holiday peak:

- **Merchants:** tag every agent-initiated order, keep whatever consent and agent-identity data arrives with it, and review dispute terms with your acquirer now.
- **Acquirers and payment providers:** accept a signed proof-of-authority record as compelling evidence in representment, and offer a defined handling standard for agent disputes.
- **Issuers and networks:** pilot a common evidence format in the 2027 rule cycle, so liability rules can reference one artifact rather than many.
- **AI platforms:** issue verifiable delegation records by default, so users and merchants can both prove what an agent was allowed to do.
- **Standards bodies:** converge existing work on intent, mandates and agent identity into a single, portable evidence profile.



Kyance's Contribution

Kyance assesses the customer, the agent and the purchase together at the moment of purchase, then captures the decision in a **Trust Receipt**. A Trust Receipt is a signed SD-JWT verifiable credential that binds the verified customer, the agent, the delegated authority, the risk decision and the transaction. Any party can verify it against Kyance's published keys and check whether the delegation was revoked, while selective disclosure limits what each party sees. It is designed to serve as evidence when a purchase is disputed.

We offer the Trust Receipt format as a starting point, not a finished standard. We invite networks, payment providers, merchants and standards bodies to test it, challenge it and help shape a shared format.

Sources

- [Why agent chargeback rules likely miss the 2026 holidays](#), ShopAppy, updated September 13, 2026
- [Agentic Commerce Regulation in 2026: What Merchants Need To Know](#), Chargeflow, updated October 11, 2026
- [AI Agent Chargeback Liability: Who Pays When Agentic Commerce Goes Wrong](#), Chargeflow, July 7, 2026



About Kyance

Kyance builds trust and identity infrastructure for agentic commerce. Our products verify the agent and the customer behind it, assess the risk of each agent purchase, and produce evidence designed to hold up in disputes.

Kyance's founding team brings together both sides of the problem. **Tim Brown, CEO**, has 30 years in identity and biometrics, holds multiple patents in the field, and has served in board and working group leadership at the FIDO Alliance. The team also brings 30 years in payments, fraud and dispute management, spanning banking, fraud prevention and chargeback industry consortiums.



Contact: info@kyance.com · www.kyance.com