

Data Privacy, Security, Safety & Risk Management

October 13, 2026

Who Said 'Buy'? Agentic Payments and the Evidence Problem

By [Amy S. Mushahwar](#), [Macauley, Venora](#), and [Tricia Y. Wagner](#) CIPP/US, CISSP, CISA

As AI agents begin to shop, select, pay, and move digital assets on others' behalf, payment disputes will increasingly turn on what the principal authorized the agent to do. With AI-driven fraud accelerating and trusted and hostile agents looking alike at checkout, the organizations best positioned in a dispute may be those that can reconstruct what happened with precision. General counsel should be reviewing fraud posture, contracts, internal controls, records retention, and insurance now, before an incident forces the issue.

Imagine a consumer tells an AI assistant to find the best deal on running shoes. The agent researches, compares products, fills a cart, and completes a purchase. The consumer receives shoes they never would have chosen. The consumer disputes the charge. Was the transaction authorized?

Now imagine an agent purchasing on behalf of a company whose instructions were altered before a payment went out. Who can show what the company approved?

Our clients have one thing in common: They need to get paid, and they need that money to move securely. That is why we will be at Money20/20 USA this month. Tricia Wagner is a former PCI QSA. Amy Mushahwar is a former CISO who has handled common point of purchase (CPP) investigations, card brand forensic audits, payment card incidents, and payment disputes. Together, we are going to listen, learn, and compare notes with the people building the rails our clients depend on.

At Money20/20, we will be listening for how networks, issuers, acquirers, and agent platforms are approaching dispute allocation; whether common standards for agent identity and intent records are emerging; how fraud teams are adapting as legitimate and hostile agent traffic converge; and how digital asset platforms are approaching privacy-preserving compliance, including zero-knowledge proofs, as agents begin to transact and trade in tokens. Money 20/20 includes sessions on agentic commerce guardrails, agent identity, security in open finance, and digital assets.

One topic is at the top of our list: agentic payments, where AI agents shop, select, and pay on someone else's behalf. For decades, payment disputes have generally begun with a straightforward question: Was the transaction authorized? Agentic commerce changes the shape of that question. When an AI agent is authorized to search, compare, and purchase, the harder question may be whether anyone can prove what the consumer authorized the agent to do.

For general counsel, that is more than a technology question. It includes questions about disputes, contracts, records, insurance, and incident response, and it is arriving sooner than most risk registers reflect. Retailers may see it first in their returns and chargebacks, if they are not seeing it already; digital asset platforms may see it in agentic trading and on-chain payments.

What General Counsel May Want To Ask Now

Agentic payments affect both sides of the ledger. A company may be a merchant receiving agent-initiated orders; a deployer of agents that purchase, pay invoices, move treasury funds, or trade digital assets on its behalf; or both. Either way, several questions are worth putting on the agenda before an incident forces them there.

Fraud Posture. Do your fraud and bot-management teams have a policy for agent traffic, and does that policy distinguish verified agents from unverified automation?

Contracts. How do your agreements with agent platforms, payment processors, and procurement vendors allocate responsibility when an agent acts outside its instructions or is compromised? Do they give you access to the agent logs and intent records you would need in order to support or contest a disputed transaction?

Internal Controls. If agents are initiating disbursements, do your approval thresholds, segregation-of-duties expectations, and audit trails account for a nonhuman initiator? Controls built around a human clicking “approve” may need to be redesigned rather than simply extended. If agents act in brokerage accounts, how will the firm meet applicable supervision and recordkeeping duties?

Records and Legal Holds. Are agent logs and mandate records covered by your retention schedule? Would your litigation hold process reach them, and can they be preserved in a form that supports authentication? Who can suspend a compromised agent’s credentials, and how quickly?

Insurance. Have you reviewed how your cyber, crime, and social engineering coverage would respond to a loss involving an agent-initiated transaction?

The Fraud Groups Are Not Waiting

All of this is unfolding while card fraud groups are becoming more aggressive and more automated. INTERPOL’s March 2026 Global Financial Fraud Threat Assessment found that AI-enhanced fraud is 4.5 times more profitable than traditional methods. Industry data points the same way: Signifyd reports that AI-driven card-testing attacks increased 175 percent from January through April 2026 compared with the same period in 2025. Why? Attackers may also benefit from the enhanced efficiencies of AI.

The tooling is also evolving. AI card-testing agents can use real browser sessions and residential proxies to mimic human checkout behavior.

This creates a dilemma for merchants and the institutions behind them. Legitimate agent-driven shopping traffic is arriving on commerce sites too, and blocking anything that looks automated will cost real conversions. In practice, the same channel that carries a customer’s trusted shopping agent may also carry a fraud ring’s testing agent, and the two can look very similar at checkout.

For anyone who has worked a CPP investigation, there is a further concern. As agent platforms aggregate payment credentials and delegated authority across many consumers and merchants, they may become attractive targets and potentially a new kind of common point of compromise.

The Rails Are Being Built Quickly

The infrastructure is moving fast. Major payment networks and technology providers have released frameworks for agentic commerce. These frameworks share a design goal: letting agents transact while limiting risk through tokenization, scoped credentials, and agent verification.

Against the fraud backdrop, agent verification may prove to be the most consequential of those features. If merchants cannot reliably tell a trusted agent from a hostile one by behavior alone, the ability to verify an agent’s identity and the scope of its authority becomes even more critical. Agentic or next-generation identity solutions, like C1, should be considered to ensure identity anomalies can be examined at machine speed.

The Authorization Question Has Changed Shape

Payment law has traditionally treated authorization as something close to binary: The customer approved the transaction, or they did not. Agents complicate that framework.

The running-shoe purchase illustrates the problem: The agent completed a transaction within the consumer's budget, but did it follow the consumer's instructions? The answer may depend on the payment method, the authority delegated, and the facts. Regulation E governs covered consumer electronic fund transfers; card network rules and commercial payment agreements may present different questions.

Now add a fraud ring to the scenario. A dispute may involve a consumer's legitimate agent that misread instructions, a compromised agent acting on a criminal's instructions, or a criminal's agent impersonating a legitimate one. Each of those fact patterns could lead to very different outcomes, and early in a dispute they may be hard to tell apart.

We do not expect these questions to be resolved quickly. In the meantime, the organizations best positioned in a dispute may be those that can reconstruct what happened with the greatest precision.

Where Evidence Architecture Comes In

Agentic payment systems generate a growing body of records: intent mandates, cart mandates, delegated authority scopes, agent identity credentials, transaction artifacts, and activity logs. The practical issue is not whether those records exist. It is whether they can be authenticated, preserved, linked to a specific transaction, and produced when someone asks who told the agent to buy.

Agentic commerce may require a parallel evidentiary framework, one that shows not only that a transaction occurred but also what authority was delegated and what the agent did as a result. That record may distinguish a misbehaving legitimate agent from a hostile one.

Four Lenses, One Set of Facts

We look at agent-initiated transactions through four related lenses:

Privacy. What consumer information did the agent collect, where did it travel, and did the permission granted cover that use?

Security. How do you distinguish a stolen card from a valid credential in the hands of a compromised agent, or a hostile agent posing as a trusted one? Does your detection approach treat them differently?

Information Governance. Who owns the intent records, cart mandates, and agent logs, and can they be produced when a dispute or investigation arises?

AI Safety. How do you know when an agent is faithfully executing instructions and when it is drifting beyond what its principal intended?

Each lens raises its own legal and operational questions. All four depend on the same foundation: the ability to reconstruct what happened and support it with evidence.

Privacy and Security in Digital Assets

These issues extend to digital assets. Privacy and security are increasingly central as platforms consider whether they can show that a transaction meets a compliance condition without disclosing more personal information than necessary. Zero-knowledge proofs—an emerging token standard—may allow a party to demonstrate a limited fact, such as completion of required screening, without sharing its underlying personal details with every participant. Compliance tooling built on this technique is emerging, including standards and tools that generate compliance proofs on the user's device for use across blockchain networks. In December 2025 remarks, the SEC chair recognized the potential for such tools to prove compliance while protecting financial privacy. That observation is not an endorsement of a particular token standard or a substitute for applicable sanctions, identification, or recordkeeping duties. For dispute and evidence purposes, the open question is whether a zero-knowledge proof will be accepted as a substitute for the records that regulators, networks, and courts have traditionally expected.

Relatedly, privacy-focused digital assets have been one of the notable trends of the year. Some use the same zero-knowledge techniques to support shielded transactions that hide the sender, receiver, and amount from public view, and on-chain data suggests a growing share of holders are choosing those shielded features.

Agents are also starting to transact in tokens rather than cards, from staking tokens to access AI services to paying for data and compute on-chain. And agentic trading is following a similar path to agentic commerce: Agents are being deployed to execute trades, rebalance portfolios, and move digital assets on a principal's behalf, often around the clock and at machine speed. These models raise the same questions asked about card payments, applied to new rails: what authority the agent held, what it consumed or traded, and whether the record can be authenticated and produced later.

If you will be at Money20/20 and want to compare notes on agentic payment security, fraud and incident readiness, dispute management, digital assets, or building an evidentiary record for agent-driven transactions, we would welcome the conversation. Reach out here or by email, and we will find time at The Venetian.

Contacts

Please contact the listed attorneys for further information on the matters discussed herein.

AMY S. MUSHAHWAR

Partner

Chair, Data Privacy, Security, Safety & Risk Management

T: 202.753.3825

amushahwar@lowenstein.com

MACAULEY VENORA

Counsel

T: 202.753.3761

twagner@lowenstein.com

TRICIA Y. WAGNER CIPP/US, CISSP, CISA

Counsel

T: 202.753.3658

twagner@lowenstein.com

NEW YORK

PALO ALTO

ROSELAND

SALT LAKE CITY

SAN FRANCISCO

WASHINGTON, D.C.

WILMINGTON

This article has been prepared by Lowenstein Sandler LLP to provide information on recent legal developments of interest to our readers. It is not intended to provide legal advice for a specific situation or to create an attorney-client relationship. Lowenstein Sandler assumes no responsibility to update the Alert based upon events subsequent to the date of its publication, such as new legislation, regulations and judicial decisions. You should consult with counsel to determine applicable legal requirements in a specific fact situation. Attorney Advertising.