




The Authorization Gap Report

Why Agentic Commerce Can't Scale Without an Identity Layer

Contents

01 A trillion-dollar question nobody has answered

03 Liminal's perspective

05 Why the gap cannot fill itself

07 What the issuance layer has to look like

09 Conclusion

02 Agentic commerce is scaling faster than its infrastructure

04 The protocols that built the rails

06 What the regulatory environment is telling us

08 The layer now exists

SECTION 01 A TRILLION-DOLLAR QUESTION

A trillion-dollar question nobody has answered

The agentic commerce wave is already here. AI agents are browsing, comparing, booking, and buying on behalf of hundreds of millions of people. The infrastructure to move money at scale exists. The infrastructure to prove who authorized it does not.

Every major agentic commerce protocol has launched with serious engineering behind it, serious institutional support, and a clear picture of how authorization records should be structured and verified.

What none of them has defined is who issues the credentials that make those authorization records trustworthy. Each protocol openly acknowledges this.

Google's Agent Payment Protocol puts it plainly in its specification: The protocol's security relies on information being cryptographically signed by the user and merchant. But this raises a crucial question: how are the public keys that verify these signatures distributed, and why should anyone trust them? Establishing these 'roots of trust' is a critical area for innovation.

That gap, the identity issuance layer that every protocol assumes but none has specified, is the authorization gap. This report explains why it exists, why it cannot close on its own, what filling it actually requires, and why the infrastructure to fill it is available today.

Agentic commerce is scaling faster than its infrastructure

What agentic commerce actually is

Agentic commerce is what happens when an AI agent acts on your behalf, browsing, deciding, and transacting, without you approving each individual step. You set the parameters once. The agent operates within them.

The everyday version is already familiar. A consumer instructs an AI agent to monitor prices and automatically reorder household goods when supplies run low, or to find and book travel that fits within a set budget. The agent handles research, comparison, and checkout. The consumer never sees the individual transactions unless they choose to look.

The higher-stakes version is something genuinely new. Imagine instructing an AI agent to actively manage your investment portfolio, not just execute preset trades the way today's automated investing tools do, but to interpret market signals, read financial news, assess risk, and make judgment calls about when to buy, sell, or reallocate your assets.

The agent operates within parameters you have defined, but within those parameters it exercises real discretion.

It might shift your portfolio away from tech stocks after a bad earnings report, or move into bonds ahead of an anticipated rate change. These are decisions with long-term financial consequences, made by a system that acts in your name.

Both use cases raise the same foundational question for every merchant, bank, and financial institution involved: did the real person actually authorize this?

The scale of what's already in motion

ChatGPT reached [900 million weekly active users](#) as of early 2026, more than double the user base from a year prior. The commercial pressure to monetize that reach through agentic transactions is significant and growing. Google, Visa, Mastercard, Stripe, and OpenAI have all launched agentic commerce protocols. Sixty-plus organizations have joined Google's ecosystem alone. The protocols are moving from specification to deployment faster than most anticipated.


The fraud problem is moving at the same speed.

The FBI's [2024 Internet Crime Report](#) recorded \$16.6 billion in cybercrime losses, the highest figure in the report's history, driven primarily by fraud.

Deloitte's Center for Financial Services [projects that generative AI-enabled fraud](#) will grow from roughly \$12 billion in 2023 to \$40 billion by 2027, a 32% compound annual growth rate, fueled by AI agents that can open accounts, impersonate legitimate cardholders, and execute fraudulent transactions continuously, at machine speed, at near-zero marginal cost.

The core question that none of the emerging protocols has answered: when an AI agent initiates a transaction on your behalf, who authorized it, and how can anyone prove that?

EXHIBIT 1 · GENAI-ENABLED FRAUD LOSSES, US\$B



32% CAGR. 2023 and 2027 are published endpoints; intervening years interpolated.

The dispute problem hiding in plain sight

When a human makes a purchase, the transaction carries implicit proof of authorization. A physical card at a chip terminal, a biometric authentication on a phone, a logged session with device fingerprinting. The evidence is imperfect, but it exists and it works well enough that disputes, while common, are resolvable.

When an AI agent makes a purchase, the transaction carries a log of what the agent did. What it does not carry, under any current protocol, is a cryptographically signed record proving that a specific, verified human authorized this specific action. The agent left a trail. Just not the one that matters.

Every fraud dispute in the agentic era ends the same way unless this changes: someone claims their agent went rogue, and no record can definitively resolve it.

EXHIBIT 2 · THE PROTOCOL LAND RUSH, 2025–2026

Google AP2 · UCP 60+ ecosystem partners incl. Shopify, Target, Walmart	Mastercard Verifiable Intent SD-JWT VC blueprint, designed to sit alongside AP2	Visa Trusted Agent Protocol Agent registration and signature verification	OpenAI Agentic Commerce Protocol Shared Payment Token scoped by merchant, time, amount	Stripe Machine Payments Protocol Machine-to-machine payments, no human in the loop
---	--	---	--	--

SECTION 03 LIMINAL'S PERSPECTIVE



Liminal's perspective


Liminal's research into the identity and fraud infrastructure underpinning agentic commerce points to a consistent gap between capability and trust.

Per Liminal's Buyer Survey (March 2026), 58.4% of buyers view agentic identity as one of the technologies that will have the greatest impact on identity and fraud solutions over the next two years. Also, 77.2% of buyers are concerned about synthetic identities as a threat to their organization, the same threat vector that makes unverified agent transactions risky.

Liminal frames the solution around three pillars detailed in its Seminal Research Report on Identity Authorization Networks: authentication (proving an agent's identity before it acts), authorization (scoping what it's allowed to do, down to spend limits and merchant boundaries), and verification (confirming after the fact that its actions matched user intent).

Liminal's market modeling in that same report values the identity-authorization-network space (the infrastructure category agentic commerce depends on) at up to \$15.5B by 2028, growing at a 60.3% CAGR, underscoring how much buyer investment is about to move into exactly this layer.

EXHIBIT 3 · WHAT IDENTITY BUYERS SAY



Liminal Buyer Survey, March 2026

PILLAR 01

Authentication

Proving an agent's identity before it acts.

PILLAR 02

Authorization

Scoping what it's allowed to do, down to spend limits and merchant boundaries.

PILLAR 03

Verification

Confirming after the fact that its actions matched user intent.

Traditional identity checks answer one question: is this really John? Agentic commerce forces a second one: did John authorize this agent to act, and within what limits? That gap between proving identity and proving authorization is the real blocker. Agentic AI collapses authentication, fraud detection, and authorization into a single continuous decisioning problem, and most organizations are not built for it yet.

Standards like x401 matter because they give the authorization layer a common shape. The businesses that treat authentication and authorization as one foundational job, not as a bolt-on after a pilot, are the ones that will open the agentic channel safely.

Filip Verley

Chief Innovation Officer, Liminal

SECTION 04 THE PROTOCOLS THAT BUILT THE RAILS

The protocols that built the rails

The protocols defining agentic commerce are serious, well-designed, and well-resourced. Understanding where the identity gap sits requires a basic picture of what they actually do, because they do not all do the same thing.

Commerce: discovery, checkout, fulfillment	Google UCP, OpenAI ACP
Trust and authorization: who authorized what	Google AP2, Mastercard Verifiable Intent, Visa TAP
Payment execution: how money moves	Stripe MPP

The identity issuance gap sits at the trust and authorization layer, where the credentials that give authorization records their legal and evidentiary weight need to come from somewhere.

THE SHARED GAP

Across the protocols designed around consumer authorization, the same gap recurs. Each can describe the structure of an authorization record and how it should flow through commerce infrastructure. None defines how a consumer's identity is established before they can produce a trusted signature, or which authorities are qualified to issue the credentials that make that signature meaningful.

AP2 and Verifiable Intent get closest. Both use verifiable credential frameworks and both flag the issuer trust question explicitly as open. The others rely on platform trust, network-managed key infrastructure, or treat consumer identity as optional.

What the ecosystem is waiting for is a neutral, qualified identity provider and Certificate Authority that the protocols can point to as a trusted root.

Google AP2 and the Universal Commerce Protocol

[AP2](#) is the foundational authorization protocol in the emerging agentic commerce stack.

Google's Universal Commerce Protocol, built with Shopify, Target, Walmart, and over 20 partners, delegates payment authorization to AP2 while handling the broader commerce experience of discovery, cart management, and fulfillment.

Think of AP2 as building a chain of signed receipts that proves, step by step, what a consumer authorized and what their agent actually did. AP2 calls these records mandates.

The Intent Mandate is signed by the consumer upfront and captures their instructions along with the parameters the agent must stay within.

The Cart Mandate is the critical binding moment: when the agent identifies a specific product and price, the consumer signs a Cart Mandate confirming they approve this specific transaction. The Payment Mandate chains to the Cart Mandate and connects the payment method to the confirmed cart.


In real-time mode, the consumer reviews and signs the Cart Mandate before the purchase goes through.

In delegated mode, the consumer signs the Intent Mandate upfront with defined constraints, and the agent acts automatically when those conditions are met.

AP2's framework for building this chain of signed records is well-designed. What the protocol explicitly does not define is how the consumer establishes their identity before receiving the credential they use to sign. That step is deferred to identity providers.

That is precisely where the gap sits.

EXHIBIT 4 · THE AP2 MANDATE CHAIN AND ITS MISSING FIRST LINK



Mastercard Verifiable Intent

[Mastercard's Verifiable Intent](#) is best understood as the implementation blueprint for what AP2 describes conceptually. Where AP2 articulates the mandate chain concept, Verifiable Intent shows how to execute it using a specific, privacy-preserving credential format: the SD-JWT VC (Selective Disclosure JSON Web Token for Verifiable Credentials). It is designed to work alongside AP2, UCP, ACP, and other protocols rather than competing with them.

The selective disclosure aspect solves a real problem in commerce: different parties in a transaction need to see different things. A merchant needs to know what was authorized in the cart. The payment network needs to know the payment details.

Neither should see the other's information. Verifiable Intent enforces these privacy boundaries by design, making selective disclosure a mathematical property of the credential rather than a policy aspiration.

At its core, Verifiable Intent creates a chain of signed records that travels with a transaction and answers the fundamental authorization questions in a way any party can verify independently.

A credential issuer verifies the consumer's identity and issues a credential linking who they are to the signing key they will use. The consumer uses that credential to sign their instructions to the agent. The agent then uses its bound authorization to produce the payment and checkout records that the network and merchant respectively need to see.

The result is a complete, independently verifiable chain from “here is who this person is” through “here is what they authorized” to “here is what the agent actually did.”

THE OPEN QUESTION The first link in that chain, the credential issuer who verifies the consumer’s identity, is what the protocol has not yet defined.

Visa TAP, OpenAI ACP, and Stripe MPP

[Visa’s Trusted Agent Protocol \(TAP\)](#) focuses on a specific, practical problem: helping merchants tell the difference between a legitimate AI agent and a malicious bot.

Built on existing internet infrastructure standards, TAP requires every agent to include a cryptographic signature with every request, verified against Visa’s published keys. TAP is effective at establishing that a Visa-registered agent is operating legitimately.

Its gap, relative to the identity problem, is that the consumer’s identity in TAP is context passed alongside the agent’s signature, not cryptographically tied to a signed record of what that consumer specifically authorized.

OpenAI’s Agentic Commerce Protocol (ACP) is built around the Shared Payment Token: a payment credential issued by the buyer’s payment provider (such as Stripe) with built-in restrictions scoped by merchant, time, and amount.

The merchant trusts OpenAI’s word that the user exists and has authorized the transaction. This works within the OpenAI ecosystem, but it produces an authorization record that cannot be independently verified without trusting OpenAI as the intermediary.

Stripe’s Machine Payments Protocol (MPP) is designed for scenarios where there may be no human in the loop at all: machine-to-machine transactions where an agent pays for a service directly, without a consumer approval step.

Human identity is an optional extension in MPP, appropriate for the automated service-to-service use cases it targets. It also illustrates that the agentic commerce stack serves many purposes, and the protocols built around consumer authorization are all the more important to get right.

Visa TAP	OpenAI ACP	Stripe MPP
Agent signatures verified against Visa keys; consumer identity as context data	Institutional trust via OpenAI	Human identity as optional extension

SECTION 05 WHY THE GAP CANNOT FILL ITSELF

Why the gap cannot fill itself

The protocols have already grappled with who could serve as that trusted authority. Each credible candidate, examined honestly, carries a structural conflict that disqualifies it from the role.

AI platforms	AI platforms authenticate users and issue the credentials that authorize agent actions. But an AI platform vouching for whether its own agent acted within bounds is the equivalent of asking a contractor to write their own inspection report. The entity whose agent is acting and the entity vouching for the authorization cannot be the same party. Merchants and financial institutions cannot rely on a self-attested authorization record, and they should not be asked to.	
Merchants	Merchants have a direct financial interest in transaction outcomes. An authorization record issued by a merchant that shifts dispute liability away from that merchant is structurally incapable of being a neutral document, regardless of how carefully it is designed.	
Card networks	Card networks are building real trust infrastructure. Visa's TAP and Mastercard's Verifiable Intent are substantive contributions to the ecosystem. But each network is developing its own key infrastructure independently, creating fragmentation rather than a shared root of trust. Network-controlled key infrastructure also concentrates visibility into consumer authorization scope across the ecosystem, raising both privacy and competitive concerns that are difficult to design around.	
OS and platform vendors	OS and platform vendors see the authorization event and the full scope of what the agent is permitted to do. Many also operate advertising businesses whose economics are influenced by consumer purchasing behavior. That conflict is difficult to architect away.	
Banks	Banks are largely absent from the working groups defining these protocols. Requiring a bank login at every agent authorization event is a user experience that will not scale to the volume agentic commerce requires.	
A disinterested third party	What the issuance role requires, and what disqualifies every candidate above, is a party whose only business is verifying who someone is and issuing the credentials they need to sign their own authorizations. No stake in the commerce. No interest in purchasing behavior. No network economics to protect.	

The disinterested third party requirement is not an idealistic preference. It is the logical consequence of the authorization problem itself. An authorization record is only worth the paper it is printed on if the entity that anchored the identity behind it has nothing to gain from the outcome being disputed.

What the regulatory environment is telling us

The law already supports agent authorized actions

Courts and legislatures dealt with the question of computer-executed transactions decades before AI agents existed. Cases from the early days of computing established that actions taken by computers, without immediate human direction, can bind the party who chose to transact via computer. That case law was codified at the turn of the century through two statutes that together provide federal and all-states coverage.

The [Electronic Signatures in Global and National Commerce Act \(ESIGN\)](#), which serves as the foundational rule set at the federal level, that provides the legal basis for the validity and enforceability of electronic signatures, and the [Uniform Electronic Transactions Act \(UETA\)](#), the state-level counterpart, both explicitly recognize that transactions can be conducted by electronic agents, including contracts formed without any individual reviewing or even being aware of the agent's actions.

The person who deploys the agent owns the results.

The remaining question under this framework is attribution: how do you establish that a specific, identifiable human authorized a specific agent to act? ESIGN requires that an agent's actions be legally attributable to the person to be bound. UETA specifies that attribution can be established through "the efficacy of any security procedure applied to determine the person to which the electronic record or electronic signature was attributable."

This is the key insight from a regulatory standpoint. The law has provided the framework and intentionally left the "how" of attribution open to technology and standards. Without a trustworthy, independent security procedure for establishing that attribution, the legal clarity UETA and ESIGN provide cannot be put to practical use in agentic commerce disputes.

Financial regulators are moving in the same direction

FinCEN has been openly working toward modernizing its Customer Identification Program (CIP) requirements to accommodate third-party identity verification. In early 2024, FinCEN issued a [Request for Information on CIP modernization](#), seeking input on whether banks should be permitted to collect and verify customer identification information through trusted third-party sources rather than directly from the customer.

By mid-2025, FinCEN and federal banking agencies had [clarified that banks may collect taxpayer identification numbers from trusted third parties](#), a signal that the regulatory direction of travel is toward third-party identity verification infrastructure, not away from it.

The broader AML/CFT modernization agenda, which FinCEN has described as building toward programs that are "effective, risk-based, and reasonably designed," points toward the same conclusion: the 2003-era model of collecting identity information directly from customers at account opening is not adequate for an environment where AI agents initiate transactions at machine speed on behalf of people who may not be present.

The global identity standards landscape is converging

The European Union's [eIDAS 2.0 regulation](#) requires all 27 EU member states to issue certified digital identity wallets to citizens by the end of 2026, with mandatory acceptance by regulated private sector organizations including banking, healthcare, and large online platforms required by late 2027.

Selective disclosure is a core design principle of the EUDI Wallet framework: consumers can prove specific attributes, such as being over 18 or being a verified resident, without exposing underlying personal data. The technical standard for this is the same SD-JWT VC format that Mastercard Verifiable Intent uses.

The NIST SP 800-63 digital identity guidelines, the OpenID Foundation's evolving verifiable credential standards, and the FIDO Alliance's authentication specifications are all converging on the same interoperable infrastructure.

The identity layer for agentic commerce will be built on these standards.

The question is who operates the issuance infrastructure that makes them work.

The identity standards that will govern agentic commerce are being finalized now. The protocols that will define how authorization records flow through the commerce stack are already in production.

Whoever is embedded in those standards as a trusted issuer by the time they harden becomes the default. The window to establish that position is open. It will not stay open indefinitely.

EXHIBIT 6 · THE CONVERGENCE TIMELINE

EARLY 2024	MID-2025	END 2026	LATE 2027
FinCEN issues its Request for Information on CIP modernization.	Banks may collect taxpayer identification numbers from trusted third parties.	All 27 EU member states must issue certified digital identity wallets.	Mandatory acceptance by regulated private sector organizations.
27 EU member states required to issue wallets	SD-JWT VC Shared credential format across EUDI and Mastercard Verifiable Intent	3 Standards bodies converging: NIST, OpenID Foundation, FIDO Alliance	

SECTION 07 WHAT THE ISSUANCE LAYER HAS TO LOOK LIKE

What the issuance layer has to look like

TECHNICAL REQUIREMENTS

IAL2 identity proofing

The identity verification process has to meet [NIST SP 800-63A's Identity Assurance Level 2 standard](#), the federally recognized benchmark for establishing that a person online is who they claim to be.

In practice, this means government document validation, database checks against credible sources, and biometric verification, designed to resist fraud at both individual and automated scale.

Certificate Authority qualification

The credential issuer has to be a qualified Certificate Authority, audited under [WebTrust for Certification Authorities](#), the same standard used to qualify CAs trusted by major browsers and operating systems worldwide.

This roots the consumer's key binding in PKI infrastructure, not a self-generated key pair, giving relying parties the ability to trace the key behind a signature back through auditable CA infrastructure to a verified identity.

Transaction-bound signatures

The credential has to be capable of producing Key Binding JWTs: cryptographic signatures tied to a specific transaction, amount, payee, and timestamp, not just proof that a user was authenticated at some point. A session token proves someone logged in. A Key Binding JWT proves a specific person authorized a specific action.

Selective disclosure

The credential format has to support selective disclosure, the ability to prove specific attributes without revealing others. When a consumer authorizes an agent to make purchases up to \$500 at electronics retailers, the merchant verifying that authorization should not need to see the consumer's date of birth or home address. They only need to confirm the authorization is valid and the constraints are as represented. The SD-JWT VC format makes this technically enforceable rather than a policy aspiration.

TECHNICAL REQUIREMENTS, CONTINUED

Revocable agent delegation

The consumer has to be able to revoke an agent's authorization without re-doing the underlying identity verification. The credential architecture has to separate identity proofing (who you are) from agent delegation (what this specific agent is permitted to do on your behalf), so that revoking an agent does not require reverifying the person.

LEGAL REQUIREMENTS

Authorization records in agentic commerce will be used as evidence in disputes. The issuance layer has to produce records that hold up in that context.

ESIGN and UETA require that agent-authorized transactions be legally attributable to an identifiable human through a reliable security procedure. The credential chain has to satisfy that requirement: from IAL2 identity proofing through CA-issued certificate through transaction-bound signature, the chain of attribution has to be auditable, non-repudiable, and legible to examiners and courts. The format matters as much as the process, because an authorization record that cannot be read by the systems that need to rely on it provides no practical protection.

TRUST REQUIREMENTS

This is the constraint that eliminates most candidates. The issuance layer has to be operated by an entity with no stake in the outcome of the transactions it is helping to authorize.

No interest in the purchasing behavior it is processing. No platform or network economics that create an incentive to weight authorization outcomes in any direction.

Independence is not a feature. It is the condition under which authorization records can be trusted by parties who were not present when they were created.

INTEROPERABILITY REQUIREMENT

The issuance layer cannot be purpose-built for one protocol and require the others to adapt. It has to produce credentials in the format the leading protocols are already built around, work across AP2, Verifiable Intent, TAP, ACP, and MPP without requiring protocol-specific integrations, and remain compatible as the standards continue to evolve.

A proprietary identity standard that requires the ecosystem to change its infrastructure to accommodate it is not a neutral trust layer. It is a competitive play dressed up as infrastructure.

SECTION 08 THE LAYER NOW EXISTS

The layer now exists

Proof is a neutral, qualified identity provider and Certificate Authority with no stake in the commerce being authorized. Proof performs the identity verification and issues the digital certificates and verifiable credentials that consumers and organizations use to sign their own authorization records.

This is not a competing protocol. It is the identity and key trust infrastructure the existing protocols acknowledge a need for and have not yet defined.


Proof ID is a cryptographic digital identity credential issued after Kantara IAL2-certified proofing, signed by Proof's WebTrust-audited Certificate Authority, and formatted as an SD-JWT VC: the same format Mastercard Verifiable Intent specifies and that AP2's credential framework is built to accommodate.

It supports selective disclosure, produces Key Binding JWTs for transaction-bound signatures, and establishes the clear chain of attribution required by ESIGN and UETA. The credential is reusable across protocols and relying parties.

The Proof API gives developers three actions: Enroll, Verify, Sign. Organizations that want to require Proof-issued identity behind their agent authorization flows can integrate through a single API without building their own CA infrastructure or proofing pipeline.



x401 is an open HTTP protocol that extends standard web infrastructure with a cryptographic identity challenge-response layer, allowing any server or merchant to require verified identity as a condition of processing an agent request, without building a custom integration to every individual protocol.



Where Proof connects across the current protocol stack

PROTOCOL	IDENTITY MECHANISM	WHERE PROOF CONNECTS
AP2	Verifiable credential mandate signing; identity establishment deferred to identity providers	Proof issues the digital certificate and VC the consumer uses to sign the Cart and Intent Mandates
Mastercard Verifiable Intent	A chain of signed credentials rooted in an issuer who verifies identity	Proof is that issuer; the full chain of signed records traces back to Proof's identity verification
UCP	OAuth session identity; defers to AP2 for transaction-level proof	Proof's role flows through AP2
Visa TAP	Agent signatures verified against Visa keys; consumer identity as context data	Proof-issued credentials elevate consumer identity from asserted context to cryptographically verified
OpenAI ACP	Institutional trust via OpenAI	For implementations requiring open, independently verifiable identity, Proof provides a portable alternative
Stripe MPP	Human identity as optional extension	Proof provides the identity layer for consumer-facing MPP deployments where authorization matters

9,000

Organizations served


1M+

Identity-proofed humans

\$640B

Secured transactions

Proof is an active contributor to the Kantara, the OpenID Foundation, the FIDO Alliance, and W3C working groups. The infrastructure is in production.



SECTION 09 CONCLUSION

Agentic commerce will happen.

The protocols are too far along, the commercial incentives are too strong, and the demand for agents that can act on our behalf is too real for the question to be whether this infrastructure gets built. The question is whether it gets built on a foundation that merchants, financial institutions, regulators, and the people delegating authority to these agents can actually rely on.

The gap Google's AP2 identified openly, who issues the trusted public keys and why should anyone trust them, has a clear answer. A neutral, qualified identity provider and Certificate Authority, performing rigorous and independent identity verification, issuing the credentials that consumers and organizations use to sign their own authorization records.

- The legal framework governing what those signatures mean already exists.
- The technical standards the protocols are converging on already exist.
- The issuance infrastructure exists today.

What remains is for the ecosystem to require that the identity layer meet the same standard of independence and rigor that the protocols themselves have established for everything else. Self-attested authorization, where the party whose agent acted also vouches for the authorization, is not a trust anchor. It is a liability exposure for everyone in the chain.

The authorization gap is closeable. The tools to close it are available now.

[Talk to Proof](#)

About Us

Proof is the identity authorization company that verifies people and organizations, issues high-assurance cryptographic credentials and creates tamper-evident evidence of what they authorize—whether they act directly or through AI agents. Proof has secured more than \$640 billion in transactions for 10,000 organizations, including 16 of the Fortune 100. Learn more at proof.com.