Teranode

OUR INDUSTRY PERSPECTIVE · MAY 2026

Powering Trust in Agentic AI

Know Your Agent: Identity and Payments in Agentic Economies

FOREWORD

Trust, Rebuilt for the Age of Autonomous Agents.

The financial industry spent decades engineering trust into payments. Now AI agents are beginning to act autonomously on behalf of users, institutions, and merchants. This report examines the structural gap that emerges when non-human actors move money and sets out what we believe an AI-ready identity and control model should look like.

IN THIS REPORT

Who, exactly, is the AI agent — and can we prove it?

EXECUTIVE SUMMARY

The Accountability Gap — and Five Ideas to Close it.

Autonomous AI agents are starting to initiate transactions on behalf of humans, businesses, and other agents. The identity and accountability frameworks built for humans cannot carry that weight. We argue that AI agent identity is becoming critical payment infrastructure for any institution that originates, processes, or handles payments. In Europe, this shift will land first on regulated payment institutions, especially Payment Service Providers (PSPs).

01

An accountability gap exists today

Most deployed AI agents cannot be cryptographically identified, attributed to a named human, or bound to specific scope of action.

02

Identity is a chain of five components

A trustworthy agent identity requires a unique identifier, developer authority, code integrity, human delegator, and authorized intent — verifiable together.

03

Static secrets were never built for agents

API keys and human-owned credentials cannot survive an era where software makes consequential decisions and moves money on its own.

04

Hardware-based identity is the trust anchor

A hardware-bound secret that even the developer cannot access creates non-repudiation for every action an agent takes.

05

PSPs need an AI-ready control tower

Under PSD3/PSR, DORA, and the EU AI Act, PSPs and other institutions that make or handle payments will need identity, delegation, policy enforcement, and audit visibility into AI-initiated payment actions.

The strategic frame

The opportunity is not a new rail, but an AI-ready control layer above existing rails that unifies agent identity, delegation, and policy enforcement. For European PSPs and other institutions that handle payments, running this layer turns autonomous agents from a risk into a new trust service.

SECTION 01

AI Agent Identity: The Missing Foundation of Autonomous AI

The structural question at the center of every autonomous AI deployment today.

The financial industry has spent decades building trust in payments: verifying the identities of humans, encrypting data in transit, and locking down the servers and networks that process transactions. However, as AI agents begin acting autonomously on behalf of users, transacting, delegating, and deciding without direct human involvement, a foundational question has gone largely unanswered: **who, exactly, is the AI agent, and can we prove it?**

Most existing identity frameworks were designed for humans and adapted, imperfectly, for machines. That mismatch leaves a critical accountability question unanswered in every autonomous AI deployment.

> There is a structural accountability gap at the heart of every autonomous AI deployment today.

Why this matters now

Agent-initiated commerce is no longer hypothetical. Booking, purchasing, subscribing, and settling actions are already being delegated to software that reasons, plans, and transacts. The missing layer is not more computation or more data; it is a way to answer, with cryptographic certainty, three questions every regulated counterparty will eventually have to ask.

THE THREE QUESTIONS THAT MATTER MOST

01 Is this agent authentic — is it the one we registered?

02 Is it acting within its authorized scope right

03 Has it been tampered with or compromised since

SECTION 02

What AI Agent Identity Actually Requires

A trustworthy AI agent identity is not one credential, it is a composable chain of five verifiable components, each answering a distinct question.

1. **Unique Identifier**
   A stable, cryptographic reference that cannot be spoofed or reassigned — e.g. signed digital credentials as used in Visa TAP and W3C decentralized identity standards.
2. **Developer Authority**
   A verified organizational identity confirming who built and registered the agent, established through KYB/KYC verification.
3. **Code Integrity**
   Cryptographic confirmation that the agent's code is untampered since the moment it was registered. Any modification invalidates the identity.
4. **Human Delegator**
   A cryptographic link connecting a named human to the agent's right to act on their behalf, within an explicitly defined scope.
5. **Authorized Intent**
   A tamper-evident record of the specific actions, amounts, merchants, and timeframes the agent is permitted to execute.

> **Design principle**
>
> Identity is not a single document, badge, or API key. It is a chain of signed assertions each issued by a distinct authority, independently verifiable, and revocable without breaking the rest of the chain.

Together, these five components answer the three questions that matter most, turning an AI agent from a convenient piece of automation into a legitimate counterparty to a regulated transaction.

SECTION 03

Where the Industry Falls Short

The gap between what is being deployed and what consequential AI actually requires.

Almost every autonomous agent in production today relies on some combination of API keys, OAuth scopes, and service accounts. These mechanisms were designed for delegated authorization: a human grants a system a fixed scope at setup time, and that scope persists unchanged. They have no way to verify that an agent's runtime behavior remains within the spirit of the original grant. The mismatch creates silent failure modes that only become visible when something goes wrong.

CURRENT STATE

- Agents identified by shared API keys that any process can present.
- Developer authority inferred from the key's environment, not cryptographically verified.
- No runtime guarantee that the agent's code has not been modified since deployment.
- Human delegation is implied, not signed, scope lives in documentation, not proofs.
- Audit trails record API calls, not intent; attribution to a named actor is brittle.

MISSING GUARANTEES

- A hardware-bound, non-exportable secret that uniquely identifies one agent instance.
- A verifiable chain from agent → developer organization → KYB-verified legal entity.
- Code integrity enforced at runtime; modified agents are automatically locked out.
- A signed link from a named human to the specific scope an agent may act within.
- A tamper-evident record of authorized intent, settable against every transaction.

The continuous enforcement problem

Even where identity checks exist today, they tend to be performed once at provisioning, at login, or at first call. Autonomous agents run continuously, update themselves, and are redeployed across environments. Point-in-time checks cannot detect drift, tampering, or silent substitution between checks.

> **Today, attestation happens once. Autonomous agents need it to hold every time they act.**

SECTION 04

A New Standard: The KYA-Ready Agent Identity

From human-owned credentials to agent-owned, hardware-based cryptographic identity.

The industry is converging on a new acronym: **KYA — Know Your Agent**. Where KYC and KYB verify humans and businesses, KYA establishes that a non-human actor can be trusted to transact. For that guarantee to be meaningful, the agent itself (not its developer, not its runtime, not its host) must hold a secret that proves its identity.

> **The core problem**
>
> Software-only credentials ultimately anchor to a human principal or operator, which makes it hard to hold the agent itself accountable. For high-stakes, regulated transactions, you need an identity the agent owns and proves through a hardware-based secret.

The question that must be answered

How can an agent hold a secret that **even its developer cannot access**? Without that property, the developer, or anyone who compromises the developer, can impersonate the agent, rewrite its scope, or replay its credentials in a different context. Solving this is the difference between an agent that can be trusted with money and an agent that can merely be told to move it.

What changes under KYA

Before	After (KYA-ready)
Static API keys issued to the developer	Hardware-based secret bound to the agent
Human-owned credentials reused across agents	Agent-owned credentials, unique per instance
Identity proven once at provisioning	Identity proven continuously at every action
Scope documented in policy	Scope signed, verifiable, enforced at runtime
Attribution via environment metadata	Cryptographic attribution to a named human

SECTION 05

Hardware-Based, Agent-Owned Identity

A hardware-bound secret the agent alone can use and even its developer cannot extract.

Any credible answer to "Know Your Agent" has to satisfy one hard constraint: **the agent, not its operator, must hold the secret that proves its identity**. That property is what makes the five-component identity chain trustworthy. Every action an agent signs is cryptographically attributable to that specific agent, running unmodified code, acting on behalf of a named human, inside an explicitly authorized scope.

Five verifiable components, anchored in an agent-held, non-extractable secret.

Hardware-based identity grounds an agent's identity and execution integrity in a verifiable root of trust, proven through attestation at runtime rather than through static credentials. In practice, that root lives in dedicated hardware security components, not in software asserting its own trustworthiness.

Crucially, hardware-based identity is not about pinning an agent to one environment; it is about anchoring it in a tamper-resistant foundation that can be re-established across environments and runtime configurations. The agent proves itself each time through attestation, instead of replaying long-lived secrets. Where such roots of trust exist, they deliver much stronger integrity and authenticity guarantees than purely software-based credential models, which remain vulnerable to extraction, replay, and spoofing.

SECTION 05

Non-repudiation by design

Hardware-based identity does more than authenticate the agent. It produces evidence that cannot be plausibly denied. Every signed action becomes a permanent, verifiable record of who acted, under whose authority, and within which scope.

The non-repudiation guarantee

For every action an agent takes, the recipient (e.g., a PSP, a merchant, a network) can verify, without relying on the agent's operator, that the action was produced by a specific, unmodified agent, acting on behalf of a specific, verified human, inside a specific, signed scope.

> **A modified agent is automatically locked out, not because policy says so, but because the cryptography does.**

Agents as verifiable counterparties

The practical consequence is that an AI agent, like a human counterparty, acquires a verifiable reputation for its actions. Recipients no longer need to rely on the operator behind the agent; they can independently verify, transaction by transaction, exactly which agent acted, which human authorized it, and within which scope. That is the foundation on which an AI-ready payment system can be built: one where accountability is a property of the cryptography, not an afterthought in policy.

Takeaway: what the cryptography has to guarantee

01

Delegation integrity

A named human's authority flows to the agent through signed delegation, with an explicit, revocable scope. The agent cannot act outside that scope, and the chain from action to agent to human is always verifiable.

02

Tamper evidence

Any modification to the agent's code invalidates its hardware-bound attestation. A modified agent is automatically locked out: it cannot sign a legitimate action, and any action it does produce is immediately detectable as unauthorized.

SECTION 06

Where this Sits in the Emerging Agent-Identity Payment Stack

How hardware-based identity complements the protocols already in market.

The race to define agent identity as critical payment infrastructure has begun. Visa's Trusted Agent Protocol, Mastercard's Agent Pay work, Skyfire's KYAPay protocol, and Trulioo's identity verification rails each address pieces of the problem, from intent signaling to agent authentication to payer identification. What has been missing is a verifiable trust anchor: proof at runtime that the agent making a request is the registered agent, acting within its authorized scope, and has not been modified. Hardware-based identity is the strongest available foundation, providing cryptographic proof of an agent's identity and integrity that can be verified at runtime across any deployment context.

Who	What they address
Visa TAP	Trusted intent signaling between agents and recipients
Mastercard Agent Pay	Agent-scoped network tokens and intent records
Skyfire	Agent identity and payment authorization layer
Trulioo	Identity verification infrastructure for regulated actors
Teranode Group	Hardware-based trust anchor beneath the agent itself

Why a trust anchor matters more than a new protocol

Protocols describe how parties talk. Trust anchors decide whether what they say can be believed. Every agent-payment protocol in the market assumes, at some layer, that the agent is who it claims to be. A hardware-based identity closes that assumption, giving the entire stack a foundation that regulators, networks, and PSPs can audit with the same rigor they apply to cardholder authentication today.

> **Agent identity is becoming critical payment infrastructure. The race is on to define who anchors it.**

Once you can prove the agent at runtime, the next question is who operates the controls around that proof.

SECTION 07 · CONTINUED

From Trust Anchor to Control Layer: Why European PSPs Carry this First

Identity is one layer in a larger control stack. In Europe, running that stack will fall, first, on Payment Service Providers.

Anchoring agent identity is necessary, but not sufficient. For an AI-initiated payment to clear under European rules, identity has to sit inside a broader control layer. One that binds delegation, policy, and audit to every action an agent takes, and does so with the same evidentiary rigor that PSPs already apply to human-initiated payments.

That shift, from who is the agent to what is the whole accountability chain around the agent, is where the regulatory weight lands. And in Europe, the first institution positioned to carry that weight is the PSP.

CONCEPTUAL MODEL · AI-READY PAYMENT CONTROL LAYER

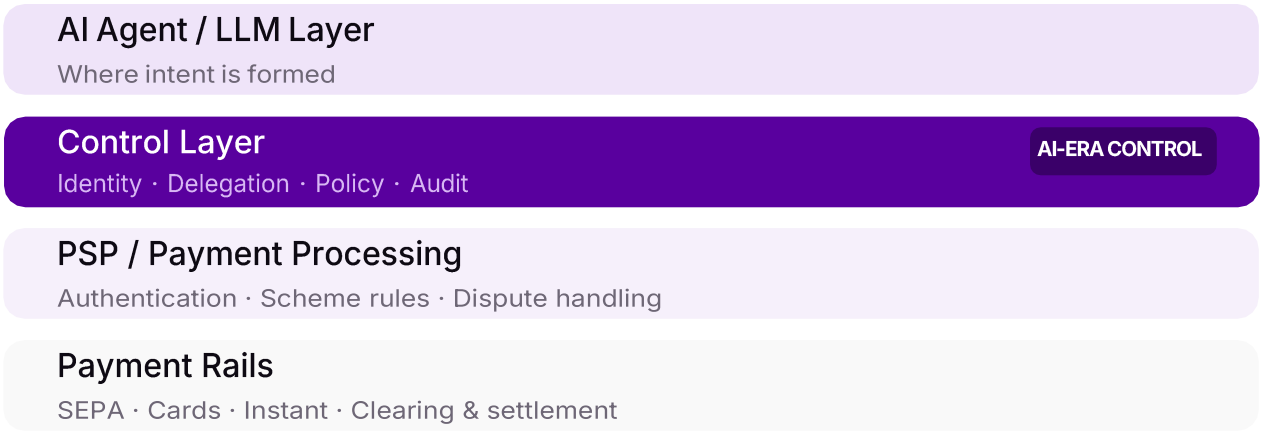

The control layer is not a replacement for the PSP or the rails beneath it. It sits between the agent and the payment institution, translating autonomous behavior into evidence a regulated counterparty can act on: which agent, on whose behalf, inside which scope, under which policy, with what proof. Hardware-based identity is the anchor at the bottom of that layer. Delegation, policy enforcement, and auditability are what make the layer usable by a PSP under existing law.

Why the zoom-out matters now

AI is moving closer to the payment decision itself. Until recently, AI lived at the edges of payments: fraud scoring, personalization, customer support. It is now moving into the moment of payment, proposing the purchase, initiating the transaction, negotiating with the merchant, and in some deployments deciding whether to proceed. That migration creates an evidentiary problem European payments law has not yet fully named, and it is the reason a control layer, rather than a new rail, is the durable answer.

A subtle regulatory issue is emerging

Under PSD2 and the proposed PSD3 / PSR framework, payment initiation and strong customer authentication rest on the assumption that a *human* consented. DORA adds operational resilience and third-party oversight obligations. The EU AI Act layers in governance and transparency requirements for high-risk AI systems. Each regime makes sense on its own. Friction appears at the intersection where an AI system shapes or initiates a payment action and the chain of accountability has to hold for PSPs, merchants, and any institution under payments regulation.

The evidentiary question

When an AI system influenced a payment, who decided? Can the PSP prove it with the same rigor it uses to prove customer consent today?

Why the pressure may fall on PSPs first

Three dynamics concentrate the pressure on PSPs rather than on AI providers.

The visibility gap

AI providers see the prompt and the output; they do not see the payment. The PSP is the only party with a complete view of the transaction, the counterparties, and the rails used. Under PSD2 and the proposed PSD3/PSR, that is where liability and strong customer authentication duties already sit.

DORA obligations

Operational resilience and third-party risk rules land on the regulated financial entity. Because PSD2 and PSD3/PSR already define the PSP as the core regulated payment institution, DORA's operational and third-party risk requirements naturally attach to them when AI becomes part of the payment chain.

The direction of AI governance

The EU AI Act focuses on documentation, oversight, and human accountability for high-risk systems. Those expectations map most naturally onto institutions already supervised under PSD2 and the future PSD3/PSR framework, which means PSPs become the default control tower for AI-driven payments.

The likely control gap

PSPs have mature controls for human-initiated payments: authentication, consent capture, scheme-level rules, dispute handling, and audit. The control layer around AI-initiated payment actions, where an agent proposes, decides, or initiates, is far less mature, both for PSPs and for other institutions that make or handle payments. It is also where most of the regulatory ambiguity will concentrate.

Five questions every firm will need to answer

For any AI-initiated payment action, a PSP and any firm that makes or handles payments should be able to answer the following with signed, auditable evidence.

01

Who is the agent?

Which specific, registered agent instance produced the action, verifiable to a hardware-bound identity, not inferred from an API key.

02

Who is accountable?

Which named human or legal entity is ultimately responsible for this agent's actions, verified through KYC / KYB.

03

What authority was delegated?

What scope, limits, and conditions were signed over to the agent and by whom.

04

What policy constraints apply?

Which runtime constraints (amount, merchant, geography, time, category) were enforced at the moment of the action.

05

What evidence exists?

A tamper-evident record linking the action, the agent, the human, the scope, and the policy enforcement, producible on demand.

What an AI-ready PSP model may need

Translating the five questions into capabilities yields a minimum viable control tower, one that composes with existing rails rather than replacing them, and that other payment-handling institutions can adopt alongside PSPs.

Capability	What it delivers
Agent identity	Verifiable, hardware-based identity for every AI agent touching a payment action.
Delegated authority	Signed, revocable delegation from a named human or entity to the agent, with explicit scope.
Policy enforcement	Runtime enforcement of amount, merchant, geography, and category limits; not just documentation.
Approval logic	Deterministic rules for when human-in-the-loop approval is required, with cryptographic capture.
Auditability	Tamper-evident, queryable records mapping actions to agents, humans, scopes, and policies.
Responsibility boundaries	Clear contractual and technical split between PSP, AI provider, merchant, and end user.

STRATEGIC TAKEAWAY

Not a new rail, an AI-ready control model.

The temptation, when new technology meets payments, is to imagine a new rail. The more durable opportunity is different: an AI-ready control model above existing rails that brings agent identity, delegated authority, and policy enforcement into the same fabric that already handles human payments for PSPs, schemes, merchants, and any institution that makes or handles payments.

> **The opportunity for PSPs is not a new payment rail. It is the AI-ready control layer that makes every existing rail and every payment-handling institution AI-ready.**

Where we go from here

Teranode Group is building both the hardware-based trust anchor and the AI-ready control layer that sits above existing rails. We are working with PSPs, payment networks, merchants, marketplaces, and enterprise payment teams to define what an AI-ready control layer looks like in practice, and to anchor it in an identity model that regulators, networks, and institutions can audit with the same rigor they apply to today's payments.

Start the conversation

We welcome discussions with PSPs, schemes, regulators, and enterprise stakeholders on AI-ready control layers, KYA, and hardware-based agent identity.
Contact: agentic-ai@teranode.group · teranode.group

COLOPHON

Published by Teranode Switzerland AG · Know Your Agent: Identity and Payments in Agentic Economies · May 2026.