

HUBFORUM 2026 · AGENTIC ENTERPRISE SERIES

VERIFIABLE AUTHORITY

Agentic commerce and enterprise
trust infrastructure

Takeaways

1 **AGENTIC COMMERCE HAS SHOWN UP IN THE NUMBERS, NOT YET IN OPERATIONS.**

AI-referred traffic to U.S. retail sites grew 393% year over year in Q1 2026 and now converts 42% better than traditional search, after converting 38% worse a year earlier. Over the same period, Shopify saw traffic multiply eightfold and orders nearly thirteenfold. Yet this traffic remains below 1% of retail visits, and of 80 major U.S. and European merchants, only four accept an agent-initiated payment – all of them American, with no new entrants since January.

2 **THE BOTTLENECK IS NOT INTELLIGENCE, IT IS AUTHORIZATION.**

Three questions replace the question of model performance: who the agent is, what it is permitted to do, and what can be proven about what it actually did. Identity is the easiest, and many players have tackled it. Authorization is rarely addressed seriously, because a rule written into the software the agent controls is not a guarantee. Proof is almost always overlooked – and it is precisely what auditors, insurers, and acquirers demand.

3 **THE INCIDENTS ARE NO LONGER THEORETICAL.**

In May 2026, an instruction hidden in Morse code was enough to divert roughly \$174,000 in a single action. In February, an agent exposed its own private keys despite being instructed not to. In April, 26 of the 428 LLM routers tested were injecting malicious tool calls. The entry vector is mundane, and the architectural corollary is unequivocal: an agent that knows a secret will eventually reveal it.

4 **LEDGER'S ANSWER: A DETERMINISTIC LAYER ON TOP OF A PROBABILISTIC SYSTEM.**

A language model cannot promise never to make a mistake; what can be done is to place the signing boundary inside a secure element the model does not control. The agent submits an access request; nothing executes without explicit human approval or compliance with the organization's policy. MoonPay and Shisa AI have put it into production, and the Agent Stack has been battle-tested across more than a thousand agents.

A deterministic layer

A probabilistic model cannot promise never to make a mistake. It can, however, be bounded by a component that does not make mistakes: that is the whole principle.

The access-request model

The agent holds nothing: it submits an access request. That request succeeds in only two cases – explicit human approval, or compliance with the policy defined by the organization. Read operations run without friction and never touch the device; any movement of value stops at the trusted display, driven by the secure element rather than the host machine. What is displayed is therefore what will be signed.

From policy engine to governed autonomy

Approving every single action does not scale to a fleet. So the human signs a policy – limits, authorized counterparties, operation types – which a hardware module then enforces on every action. The agent operates at machine speed within that envelope; anything outside it is denied a signature and escalated for human arbitration. Autonomy and security, rather than one at the expense of the other.

THREE DEPLOYMENTS BACKED BY LEDGER HARDWARE

MoonPay, March 2026 · First production example: the agent identifies opportunities and prepares the operations; every transaction requires a physical press on the device.

Shisa AI, Tokyo · The open-source shisad security daemon uses Ledger as its trusted-display signer for its highest-risk approval tier.

Ledger Agent Stack, July 2026 · Open-source toolkit battle-tested on more than a thousand agents over six weeks of private preview before public release.

The architecture of trust

Four layers, only one of which lies beyond the reach of the agent's software. That asymmetry, and it alone, is what makes a commitment enforceable.



The three upper layers are interchangeable and competitive: this is where the protocol battle is being fought, and no winner has been declared. The bottom layer is stable and agnostic – whichever standard prevails above it, the secure component beneath remains the same. It does not replace any of the other three: it makes their promises verifiable.

What Ledger makes available

1 LEDGER ENTERPRISE, AVAILABLE TODAY.

Hardware-enforced signing policies, governed approval workflows, trusted-display multisig, and full approval traceability. Deploys immediately on treasury and outbound payments.

2 THE POLICY ENGINE.

The human signs limits, authorized counterparties, and operation types; the hardware enforces them on every action. What it unlocks: an agent that works at machine speed without every move requiring sign-off.

3 ASYNCHRONOUS APPROVALS.

Next step on the roadmap: human validation is inserted at the right moment without blocking the execution chain. What it unlocks: use cases where waiting for a human cost more than the risk it avoided.

4 PROOF OF HUMAN AND PROOF OF YOU.

Attest that a hardware device was actually operated by a person, without handling biometrics or ID documents. What it unlocks: proof that stands up to an auditor or an insurer, without creating an identity database waiting to be stolen.

5 LEDGER AGENT STACK, OPEN SOURCE.

Device Management Kit, wallet and Enterprise command-line interfaces, distributed in AGENTS.md and SKILL.md formats. What it unlocks: connecting an existing agent to the hardware in a matter of days, with no re-architecture.

Three sectors, real-world use cases

Financial services, utilities & energy, retail & CPG: three sectors where delegation is already documented, and three ways of running into the same authorization question.

Financial services

This is the most advanced sector, because the risk there can be quantified. In March 2026, Santander and Mastercard executed the first European agent-initiated payment, within a regulated banking framework and under predefined limits. In July, BBVA, CaixaBank, and Banca Transilvania did the same with Visa, meeting strong customer authentication requirements. Regulators are following: in July 2026, the UK Treasury opened a consultation on consent and liability.

Utilities & energy

The sector delegates at scale, but on operations rather than purchasing: in August 2026, Iberdrola reported nearly 500 agents deployed or in development across grids, renewable generation, and customer relations, while specifying that responsibility for decisions remains with people. The tipping point will be trading: an agent that buys flexibility commits real money.

RETAIL & CPG – THREE DOCUMENTED CASES

Nestlé · Retail execution agents deployed with effie.ai carry out every task of a store visit except physical merchandising.

L'Oréal, Unilever, Mars, Beiersdorf, Reckitt · In March 2026, these five groups adopted the Agentic Merchant Protocol to retain control of their catalogs and avoid dependence on any single AI provider.

Unilever · A five-year partnership with Google Cloud to build a digital foundation geared toward agentic commerce, marketing, and data.

Five decisions

To make within the next hundred days

1 INVENTORY THE FLOWS, NOT THE AGENTS

Map not the agents but the flows where value or a commitment changes hands. One named owner per flow.

2 SEPARATE READ FROM WRITE

Allow without friction whatever only reads; isolate whatever commits. The least costly trade-off in the whole setup, achievable within a few weeks.

3 WRITE THE AUTHORIZATION POLICY

Limits, counterparties, human-escalation thresholds. A versioned policy, reviewed like code – and a clearly identified enforcement point.

4 INSTRUMENT THE PROOF

Decide what must be provable, to whom, and for how long. Target: produce proof of an authorization in under twenty-four hours.

5 APPOINT AN OWNER

A single role carries human agency – policy, proof, incidents. Ledger created this role at executive committee level in April 2026.

These five decisions run on existing budget lines – fraud, payment compliance, access management. Diagnosis and analytical framework: HUB Institute. Architecture and execution building blocks: Ledger.

Authority is proven.

HUBFORUM PARIS · September 30, 2026

HUB Institute × Ledger white paper, September 2026.

Ledger: more than 8 million signers across 165 countries, CC EAL6+ certified secure element. Dated public sources: Ledger, Mastercard, Visa, Iberdrola, Nestlé, Adobe Analytics, Shopify, Salesforce.