Subject: The Identity-to-Execution Gap: Primary Barrier to Confident AI Agent Adoption in Financial Services
Date: Sunday, March 29, 2026 at 4:08:44 PM Central Daylight Time
From: Vincent Nijjar <vincent@zlar.ai>
To: caisi-events@nist.gov <caisi-events@nist.gov>

Submitted to: NIST Center for AI Standards and Innovation (CAISI) — Listening Sessions on Sector-Specific Barriers to AI Adoption

Sector: Financial Services

Date: March 29, 2026

The Barrier

Financial institutions cannot safely deploy AI agents for autonomous action because no standard governs what agents do after authentication.

The credential boundary is largely addressed. OAuth, OIDC, SPIFFE, and SCIM help authenticate agents and manage identity. The execution boundary is not addressed. There is no standard that provides deterministic, per-action enforcement of what an authenticated agent is allowed to do, or cryptographic proof that enforcement occurred.

This is the core structural gap preventing confident agent deployment in regulated financial services.

The Evidence

Cloud Security Alliance / Strata Identity (February 2026) found that only 28% of organizations can trace agent actions back to a human sponsor, and only 21% maintain a real-time inventory of active agent identities.

The FIFAI II Report (March 23, 2026), co-sponsored by OSFI, the Canadian Department of Finance, FINTRAC, FCAC, and the Bank of Canada, found that 44% of workshop participants identified autonomous AI systems as the primary source of AI-related systemic risk.

Gartner projects that more than 50% of successful agent attacks by 2029 will exploit access control weaknesses. The primary gap is authorization, not identity.

Rubrik Zero Labs reports an 82-to-1 ratio of non-human identities to human identities, and estimates that a single agent requires 15 to 20 distinct non-human identities.

Major financial institutions continue to rely on manual constraints. JPMorgan Chase has adopted an internal-first rollout approach to de-risk agent deployment. Morgan Stanley maintains a humans-press-the-button model. Goldman Sachs has highlighted the frontier challenge that emerges when agents begin acting like Goldman employees. BlackRock chose a supervised agentic architecture for Aladdin Copilot, covering approximately $25 trillion in assets under management, emphasizing reliability and testability.

The Regulatory Friction

Four overlapping frameworks create compounding compliance friction at the execution boundary.

SOX Sections 302 and 404 require CEOs and CFOs to certify internal control effectiveness. Agent-influenced financial decisions therefore require verifiable audit trails that bind actions to human authorization.

SR 11-7 and OCC 2011-12 require pre-deployment validation, ongoing monitoring, and complete documentation. These requirements become difficult to satisfy in environments involving multi-model chaining and dynamic tool invocation.

OCC 2023-17 imposes third-party risk management requirements for every external API an agent invokes at machine speed.

BCBS 239 requires complete data lineage from origin to final use. Agent-based aggregation and processing chains create lineage complexity that many legacy systems cannot reliably track.

Each of these frameworks requires capabilities that do not currently exist at the execution boundary: deterministic enforcement, tamper-evident audit trails, and action-level accountability.

The Treasury's Financial Services AI Risk Management Framework (February 2026), developed with more than 100 institutions, establishes control objectives across risk categories where execution enforcement and auditable evidence artifacts are directly relevant.

The SEC's 2022 amendments to Rule 17a-4 (Release No. 34-96034) explicitly recognize cryptographic audit trails, including hash chains and digital signatures, as a valid alternative to WORM storage. This creates a regulatory pathway for cryptographic governance mechanisms.

What the Standards Gap Prevents

Without execution-boundary governance, financial institutions cannot:

Demonstrate to examiners that agent actions were authorized by policy at the time of execution under SOX 404.

Produce tamper-evident records that satisfy SEC Rule 17a-4's audit-trail alternative through hash-chained, digitally signed records.

Provide evidence that supervision occurred for every agent action under FINRA Rule 3110.

Reconstruct complete data lineage through agent processing chains under BCBS 239.

Enable independent third-party verification of governance records for PCAOB AS 2201 examinations.

A Solution Exists

ZLAR is an open-source execution-boundary governance system that is already built, deployed, and operating in production.

Its architecture is simple:

An agent attempts an action, such as opening a file, calling an API, or executing a trade.

A gate intercepts the action before execution. Every action is intercepted, not sampled and not reviewed after the fact.

The gate evaluates the action against codified policy rules. This evaluation is deterministic and does not rely on AI reasoning. Safe actions are allowed, prohibited actions are blocked, and actions requiring judgment are routed to a human for approval in real time.

The human decision-maker approves or denies the action, for example from a mobile device. The agent cannot proceed without that approval when policy requires it.

A cryptographic proof is generated for every decision. Allow, deny, and override events are signed with Ed25519 and hash-chained into a tamper-evident audit trail. These records are durable and independently verifiable.

The gate sits between the agent and the world. Nothing passes through without enforcement.

Why This Is Different

The gate itself has no intelligence. It cannot be persuaded, prompt-injected, or manipulated through reasoning because it is not a model. It is a deterministic enforcement point.

This is the key design principle: the mechanism governing the agent cannot itself be the kind of system the agent can influence.

What This Solves for Financial Services

For SOX 404, every agent action is bound to policy evaluation and, where required, human authorization at the moment of execution. The audit trail proves it.

For SEC Rule 17a-4, hash-chained and Ed25519-signed records support the cryptographic audit-trail alternative recognized in Release No. 34-96034.

For FINRA Rule 3110, human supervision is structurally enforced. The agent cannot act where policy requires oversight unless a human approves.

For BCBS 239, the system creates complete action-level lineage from the original agent request through policy evaluation to the human decision, all cryptographically sealed.

For SR 11-7, the policy itself is a human-authored, version-controlled, and signed artifact. The gate enforces that policy deterministically. This makes model-risk governance more auditable because the governance layer is not itself a model.

Standards Alignment

ZLAR maps to NIST SP 800-207 as the Policy Decision Point and Policy Enforcement Point for agent actions.

Policy evaluation targets Cedar, a formally verified language currently in CNCF Sandbox.

Cryptographic signatures use FIPS 186-5 through Ed25519, with a migration path to FIPS 204 using ML-DSA-44 for post-quantum readiness.

ZLAR is open source, available now, and already running in production to govern AI agents.

The barrier described in this submission is real, measurable, and solvable. The solution already exists.

GitHub: https://github.com/ZLAR-AI/ZLAR
Website: https://zlar.ai

Vincent Nijjar
Founder, ZLAR

vincent@zlar.ai

About the Submitter

Vincent Nijjar has 25 years of experience in financial services and is submitting a concurrent public comment on the NCCoE concept paper on AI agent identity and authorization.

He is available to participate as a working group member or technical contributor to the CAISI initiative.

https://www.linkedin.com/in/vincentnijjar/