Nick Clark · Published April 25, 2026 · PDF

Audit-Grade Agent Identity for Regulated Finance and Healthcare: Continuity-Proof Lineage Across the Agent Lifecycle

Agents deployed in finance and healthcare must answer two questions that today's stateless inference services cannot: who was this agent, continuously, across every model upgrade and hardware migration, and what exactly did it do, in a record no one can quietly rewrite. This application is built on the Agent-Resident Execution Substrate, disclosed in U.S. Provisional Application No. 64/070,239, which carries the agent's identity, append-only lineage, and governance policy as substrate-resident state preserved across the agent's entire lifecycle. It draws on sibling portfolio inventions for memory-native identity and cryptographically enforced governance.

What This Application Specifies

Regulated industries do not buy capability. They buy accountability. A model that recommends a treatment, approves a loan, flags a transaction, or drafts a disclosure has to be attributable to a stable, identifiable actor whose complete history of decisions can be reconstructed and verified after the fact, often years later and often by an outside examiner. Conventional inference services are built the opposite way: the model is shared across all users, holds no persistent identity, and discards its context between requests, so there is no continuous entity to audit and no internally guaranteed record of what was decided.

This application places a regulated-industry agent on the Agent-Resident Execution Substrate disclosed in U.S. Provisional Application No. 64/070,239. On the substrate, the agent is a persistent execution entity that carries its own identity field, an append-only lineage field, a cognitive state field, and a governance policy field as substrate-resident state. The inference models the agent uses are subordinate managed components, registered in a tool registry and subject to governed lifecycle operations, rather than the agent itself. The agent's identity does not depend on any particular model artifact and is preserved across replacement, retraining, archival, or removal of any subordinate model, across substrate-runtime updates, and across migration to new hardware. For a compliance function, this is the load-bearing property: the actor under examination remains one continuous, identifiable actor even after every model it was running has been upgraded.

The substrate records each operation as an append-only lineage record. The spec specifies inference dispatches, their outcomes, integrity-signal feedback, every tool lifecycle operation, governance policy updates, counterparty encounters, and off-device disclosure events as recorded events, each verifiable against its predecessor under a continuity proof. The lineage is structured so that no prior record can be modified or deleted without producing a detectable continuity break, optionally enforced by cryptographic chaining in which each record references its predecessor. Any party with access to the lineage can verify that the recorded history has not been altered. That is an audit trail with an internal integrity guarantee, not a log file a privileged process can rewrite.

Why It Matters

Financial services and healthcare share a structural demand: a regulated decision must be reconstructable and the reconstruction must be trustworthy. A bank examiner reviewing adverse-action decisions, a health authority investigating a clinical recommendation, or an internal model-risk function validating a deployed system all need to establish, after the fact, exactly which actor decided what, under which

governing policy, using which model version. The hard part is not storing logs. The hard part is binding those logs to a continuous identity and guaranteeing they were not edited.

Three failure modes recur with conventional deployments. First, identity discontinuity: every model upgrade silently produces a different decision-maker, so the entity that made last quarter's decisions no longer exists to be examined. Second, mutable records: the audit trail lives in infrastructure that the operator can rewrite, which is precisely the trust gap a regulator is there to close. Third, uncontrolled disclosure: sensitive records (patient data, transaction histories, model parameters trained on protected information) can leave the boundary without an authoritative record of what left and under whose authority.

The substrate answers each of these as a structural property rather than an operational promise. Identity is preserved by continuity guarantee across the entire lifecycle. The record is append-only under continuity proof. And, as specified in the privacy invariant, lineage records, model artifacts, training corpora, model parameters, and counterparty identity records are not transmitted off the device except under an explicit disclosure policy object, with every off-device disclosure recorded as a deterministic disclosure event that the spec expressly contemplates being audited by a regulatory authority. The framing to a specific regulation is external context, but the mechanisms the spec discloses map directly onto what regulated examinations actually demand.

How It Composes With the Domain

A regulated deployment runs the substrate in its enterprise-server or industrial configuration, where the operating institution, rather than an individual person, is the user identity governing the substrate, and organizational policy objects are the dominant authority. The spec explicitly lists a hospital and a laboratory among

industrial deployments and contemplates organizational deployments whose counterparty records correspond to clients, vendors, employees, and contractual partners, with persistence promotion governed by organizational policy.

Several substrate primitives compose into a regulated-industry implementation:

- **Continuity-proof lineage as the system of record.** Every inference dispatch, outcome, and integrity signal is appended to the lineage with its input descriptor, the endpoint that served it, the output, a timestamp, and the governing policy version. Because the recorded policy version travels with each operation, an examiner can later verify that a given decision was admissible under the policy in effect at the time, even after the policy has since been revised. The replaced policy object is archived in association with its successor, so the governing rule for any historical operation remains recoverable.
- **Identity preserved across model upgrades.** When a clinical or financial model is retrained or replaced, the lifecycle controller performs a governed substitution: the updated artifact is validated and promoted, or rolled back on failure, with the event recorded in the lineage. The agent's identity, cognitive state, and lineage are untouched. The institution can therefore upgrade models continuously without ever creating a new, un-auditable actor.
- **Hardware-bound identity and governed migration.** The agent's identity can be cryptographically bound to a hardware security element (a secure enclave, trusted platform module, or hardware security module) whose private key material is not extractable. The bound identity is not transferable to another device except under a governed migration operation carrying attestations from both the originating and destination hardware elements, recorded in the lineage. A data-center hardware refresh becomes a recorded, attested migration of the same identity rather than a break in the chain. The substrate can further quarantine agent execution if the hardware element fails attestation.

- **Counterparty records for the parties an audit touches.** The spec's counterparty identity records explicitly include regulatory inspectors among the entities an industrial substrate encounters, alongside operators, suppliers, and customers. Each record carries a counterparty scope object specifying the admissible categories of inference, the admissible categories of lineage disclosure, and the procedural requirements for interactions, all evaluated and recorded at each encounter.
- **Egress-controlled disclosure with a verifiable trail.** Producing records for an examiner is itself an off-device disclosure event. It runs through the disclosure policy object (recipient, scope, authorization attestation, retention requirement, revocation mechanism) and is recorded so the full set of disclosures originating from the substrate is itself auditable. Enforcement options the spec describes include a runtime egress filter over outbound traffic and signed disclosure-policy preconditions checked before any transmission key is released.
- **Scope partitioning under one identity.** A single institutional agent can maintain named scopes, each with its own corpus policy, tool subset, and lineage partition, while preserving one continuous identity across all of them. A regulatory-domain-specific scope is one of the spec's enumerated examples. This lets, for instance, a payments line of business and a lending line of business operate under distinct admissibility policies and isolated lineage tails without fragmenting the agent's overall continuity proof, which spans the unified lineage regardless of scope.

What This Enables

Built this way, a regulated-industry deployment supports capabilities that are awkward or impossible on stateless services:

- **Reconstruction on demand.** Because the complete operational history is deterministically reconstructible from the append-only lineage, an institution can reproduce any past decision: the inputs, the model version and publisher, the

governing policy version, and the downstream outcome, without depending on the integrity of external logging that an examiner has no reason to trust.

- **Continuous model improvement without losing the audit subject.** The substrate's lineage-derived retraining lets subordinate models improve from real downstream outcomes (acceptance, revision, execution success or failure, integrity-signal feedback) under a corpus policy that filters for admissibility, while the governing agent identity and its lineage persist unbroken. Models get better; the auditable actor stays the same.
- **A personal corpus model inside the governance boundary.** The substrate can hold a model whose parameters internalize an institution's own designated body of work without retrieving over that corpus at inference time, fine-tuned only on artifacts admissible under the corpus policy. For an institution holding protected data, internalizing knowledge within the governance boundary, rather than shipping records to an external retrieval service per query, is a meaningful posture.
- **Federated identity across an institution's devices.** Multiple substrate devices can operate under one federated agent identity verified through cross-device attestations, exchanging lineage records (not raw model weights) under federation policy, with each federation event recorded on every participating device. The institution presents one continuous, attestable agent identity across its fleet while each device retains local responsibility for its own artifacts and lineage.
- **Compensation and provenance for third-party models.** Installed models retain publisher signatures and attribution metadata across their lifecycle, and the substrate can attribute dispatch and retraining events to a model's publisher under a compensation policy, useful where an institution must demonstrate the provenance and licensing posture of every model in its decision pipeline.

Boundary Conditions

The home invention is disclosed in a U.S. provisional application; it is an early-stage filing describing an architecture and its embodiments, not a certified or shipping compliance product. The substrate provides structural mechanisms (continuity-preserving identity, append-only lineage under continuity proof, governed disclosure). It does not, and cannot, on its own establish legal compliance with any particular regulatory regime. Mapping the substrate's recorded events onto the specific evidentiary requirements of a given financial or healthcare regulator is integration and validation work external to the disclosure, and regulatory acceptance of any audit mechanism is a matter for the relevant authority, not a property of the technology.

The continuity and integrity guarantees are as strong as the cryptographic primitives, hardware security elements, and signed policy objects underpinning them; a compromised hardware element triggers quarantine, but the trust model is only as good as its key custody and attestation chain. The lineage guarantees that records were not altered after the fact; it does not by itself guarantee that the original recorded decision was correct. And general machine-learning, cryptographic, and audit-logging techniques referenced as building blocks are prior art the application builds upon, not claimed inventions. The novel contribution is the substrate architecture that binds identity, lineage, governance, and disclosure into one continuity-preserving entity across the agent lifecycle.

Disclosure Scope

The technology described here, the agent-resident execution substrate with continuity-proof lineage, hardware-bound and lifecycle-preserved identity, governed model lifecycle operations, egress-controlled disclosure, and federated agent identity, is disclosed in U.S. Provisional Application No. 64/070,239. This application also references mechanisms disclosed in related portfolio filings, including memory-native identity and authentication and cryptographically enforced governance for autonomous

agents, each incorporated by reference in the underlying disclosure for its corresponding subject matter. The finance and healthcare framing, the description of examiner and model-risk workflows, references to regulatory inspectors and audit obligations, and any characterization of what regulators require are provided as external domain context to illustrate a faithful enabling implementation; they are not part of the patent claims, do not constitute legal or compliance advice, and do not assert that the substrate satisfies any specific regulatory standard. Named regulatory concepts are referenced only as real-world domain facts, and any commercial alternatives are referred to by category rather than by product.

Agent-Resident Execution Substrate (/agent-resident-execution-substrate)

All 40 steps → (/inventive-steps)

Persistent execution environment carried by the agent, not the host — identity, state, and lineage across power cycles, devices, and upgrades.

Provisional application

PRIMARY TECHNICAL DISCLOSURE

- Agent-Resident Execution Substrate, Articles (/articles/agent-resident-execution-substrate)

SECONDARY TECHNICAL

- Persistent Semantic Agent (/articles/agent-resident-execution-substrate/persistent-semantic-agent)
- Managed Inference Tool Registry (/articles/agent-resident-execution-substrate/managed-inference-tool-registry)
- Agent-to-Tool Dispatcher (/articles/agent-resident-execution-substrate/agent-to-tool-dispatcher)
- Lineage-Derived Training Signal (/articles/agent-resident-execution-substrate/lineage-derived-training-signal)
- Identity Preservation Across Upgrades (/articles/agent-resident-execution-substrate/identity-preservation-across-upgrades)
- Cognitive State-Conditioned Dispatch (/articles/agent-resident-execution-substrate/cognitive-state-conditioned-dispatch)

- Governed Tool Lifecycle (/articles/agent-resident-execution-substrate/governed-tool-lifecycle)
- Continuity-Proof Lineage (/articles/agent-resident-execution-substrate/continuity-proof-lineage)
- Substrate Runtime Continuity (/articles/agent-resident-execution-substrate/substrate-runtime-continuity)
- Personal Corpus Model Training (/articles/agent-resident-execution-substrate/personal-corpus-model-training)
- Heterogeneous Inference Endpoints (/articles/agent-resident-execution-substrate/heterogeneous-inference-endpoints)
- Atomic Lifecycle Substitution (/articles/agent-resident-execution-substrate/atomic-lifecycle-substitution)
- Integrity Signal Feedback (/articles/agent-resident-execution-substrate/integrity-signal-feedback)
- Hardware-Bound Identity (/articles/agent-resident-execution-substrate/hardware-bound-identity)
- Cognitive State Append-Only Invariant (/articles/agent-resident-execution-substrate/cognitive-state-append-only-invariant)
- Counterparty Identity Records (/articles/agent-resident-execution-substrate/counterparty-identity-records)
- Privacy Egress-Controlled Disclosure (/articles/agent-resident-execution-substrate/privacy-egress-controlled-disclosure)
- Federated Cross-Device Agent Identity (/articles/agent-resident-execution-substrate/federated-cross-device-agent-identity)

APPLICATIONS · GENERAL

- Personal AI Agents That Survive Device Loss: One Continuous Identity and a Private Corpus Across Every Device (/articles/agent-resident-execution-substrate/personal-cross-device-agents)
- Enterprise Agent Fleets: Stable Agent Identity and Governed Tool Access Across Model Upgrades and Infrastructure Migration (/articles/agent-resident-execution-substrate/enterprise-agent-fleets)
- **Audit-Grade Agent Identity for Regulated Finance and Healthcare: Continuity-Proof Lineage Across the Agent Lifecycle (/articles/agent-resident-execution-substrate/regulated-industry-agents)**
- Edge and On-Device Agents: Hardware-Bound Identity Across Heterogeneous Inference Endpoints (/articles/agent-resident-execution-substrate/edge-and-on-device-agents)
- Agent-to-Agent Commerce With Counterparty Identity Records and Egress-Controlled Disclosure (/articles/agent-resident-execution-substrate/agent-to-agent-commerce)
- Governed Tool Lifecycles for Managed Inference-Provider Ecosystems: A Substrate Approach to Owning, Routing, and Retiring AI Tools (/articles/agent-resident-execution-substrate/managed-tool-ecosystems)

– Proving Unbroken Continuity in Long-Lived Autonomous Systems Across Substrate Migration and Atomic Model Substitution (/articles/agent-resident-execution-substrate/long-lived-autonomous-systems)

Agent-Resident Execution Substrate overview → (/agent-resident-execution-substrate)