

# **End-to-End Lineage and Audit: Reconstructing Any Agent Action Across Every Tier of the Stack**

Regulated operators cannot answer a basic question about autonomous AI agents: when an agent made a decision three substrates and two weeks ago, what was its exact state, which policy bound it, and who authorized it. Most stacks lose that thread the moment an agent crosses a tier boundary, leaving compliance teams to reconstruct behavior from disconnected logs. This capability is built on the Cross-Patent Architecture, disclosed in United States Patent Application 19/647,395, which threads a single deterministic lineage record through every tier the agent touches: substrate execution, the cognition platform, content discovery, state-preserving transport, and identity. It draws on sibling portfolio inventions for the execution substrate, the canonical agent schema with its lineage field, the transport layer, cryptographic governance, and trust-slope identity.

---

## **What This Application Specifies**

This application specifies a single provenance and governance thread, the lineage field, that travels with an autonomous agent as a canonical part of the agent itself rather than as an external log scattered across the systems the agent happens to run on. In the architecture disclosed in United States Patent Application 19/647,395, every semantic

agent carries a fixed set of canonical fields (intent, context, memory, policy reference, mutation descriptor, and the lineage field), extended with cognitive domain fields such as integrity, confidence, affective state, personality, and capability. Each cognitive domain field is independently tracked with a current value and a trajectory over time, and every change to any of them is written into the same lineage chain as all other agent state transitions.

The defining property is deterministic reconstructibility. The specification records, in the lineage field, each proposed mutation, each composite admissibility determination, and each cognitive domain field update, such that the complete behavioral trajectory of the agent is reconstructible from the lineage field alone. Because the agent carries the complete cognitive state, an execution substrate that hosts the agent validates proposed state transitions without retaining authority over the agent's cognitive state. The audit thread is therefore a property of the agent, not of any one host, and it does not break when the agent moves.

For a domain operator, that means the answer to "what was this agent's exact state at the moment of a disputed decision" is not assembled from correlated logs of varying fidelity. It is replayed. The specification describes reconstructing an agent's state at any historical point by replaying the deterministic update function over the sequence of recorded observations in the lineage, producing the exact state vector that existed at the queried timestamp, without requiring persistent storage of moment-to-moment values.

## **Why It Matters**

Autonomous agents in production do not stay in one place. A single workflow can begin on a centralized substrate, migrate to an edge or embodied substrate, traverse a content index to gather information, and negotiate with other agents along the way. Each of those tiers, in a conventional deployment, keeps its own records in its own format with its own clock. When a regulator, an auditor, or an incident responder asks what

happened, someone has to stitch those records together and hope the seams line up. The seams rarely do, and the gaps at tier boundaries are exactly where accountability fails.

The cost of that gap is concrete in regulated settings. Financial services, healthcare, and other supervised domains increasingly expect operators to explain automated decisions after the fact and to demonstrate that the system operated within its authorized policy at the time. Emerging governance expectations for high-impact AI, including risk-management and record-keeping obligations of the kind reflected in frameworks such as the EU AI Act and the NIST AI Risk Management Framework, push in the same direction: keep records sufficient to reconstruct and review automated behavior. A provenance thread that survives across tier boundaries is precisely what makes such reconstruction tractable instead of forensic guesswork.

## **How It Composes With the Domain**

The architecture's distinguishing contribution is that the lineage thread is maintained continuously through the structural interactions between tiers, not just within any single tier. Several composition mechanisms in the specification keep the thread intact at exactly the boundaries where audit trails normally fragment.

State-preserving transport across substrates. When an agent migrates between substrates, the transport layer carries the complete agent state, including cognitive domain fields and cryptographic signatures, across the network. The agent's state is not reconstructed at the destination but carried intact, and the destination substrate validates the agent's lineage continuity on arrival. For an operator, a migration is a recorded, verifiable continuation of the same audit thread rather than a handoff to a fresh log.

A transit state that keeps recording. The specification defines a transit cognitive state that applies while an agent is between substrates, when no host is providing compute. During transit the agent's cognitive field values are frozen at their pre-transit levels, while the lineage field continues to accumulate transit events: departure timestamp, transport path, and arrival validation. The window between tiers, normally an audit blind spot, becomes part of the record.

Governance authority resolved against the agent's own integrity trajectory. When an agent encounters a governance policy signed by an authority its trust-slope history does not recognize, it evaluates the governance claim against its own integrity trajectory, the accumulated pattern of normative consistency recorded in its lineage, rather than relying on signature validation alone. The decision to accept or question an authority is itself grounded in, and recorded against, the lineage thread.

Policy freshness under asynchronous execution. Agents in this architecture can suspend and resume across asynchronous intervals. On resume, if the agent detects that the policy in force at suspension has been superseded, the confidence governor treats stale policy as a confidence input; a sufficiently significant policy change can drop the agent into a non-executing cognitive mode and generate an inquiry for the current policy before execution continues. The audit record therefore captures not only what the agent did but whether it was operating under current governance when it did it.

Substrate identity revocation during active cognition. If a hosting substrate's device-hash identity validation fails mid-execution, the capability envelope reclassifies the substrate as unverified, readiness drops, and the agent transitions to non-executing cognitive mode pending re-validation or migration. Because cognitive state is carried by the agent, the lineage thread is preserved across the revocation event rather than lost with the compromised host.

Discovery that records every step. When the agent performs discovery over the content index, it traverses as a schema-conformant agent carrying its own governance, identity, and cognitive state, and it records each traversal step in its own lineage. Information gathering, often opaque in conventional pipelines, becomes a governed and reconstructable sequence within the same thread.

## **What This Enables**

The practical payoff is a single reconstructable account of an agent's life across the whole stack. An operator can take any action an agent took, anywhere in the system, and replay the deterministic update function over the recorded lineage to recover the exact state, the policy in force, the admissibility determination, and the cross-tier events that surrounded it. Several concrete capabilities follow.

End-to-end incident reconstruction. After a disputed or anomalous decision, an investigator can trace the agent backward and forward across substrate migrations, transit windows, discovery traversals, and policy transitions without leaving the lineage thread or correlating external systems.

Compliance review without state retention overhead. Because state at any timestamp is replayable from the lineage and the update function, the architecture supports post-hoc regulatory review without storing moment-to-moment state, which keeps audit feasible at scale.

Policy-currency assurance. The policy-freshness mechanism makes it possible to demonstrate, from the record, that an agent either operated under current governance or correctly paused to obtain it, which is a frequent demand in supervised domains.

Cross-tier accountability for multi-agent work. Because each agent carries its own lineage and integrity trajectory, interactions and delegations between agents are anchored to verifiable records on both sides rather than to a single orchestrator's view.

Embodiments may apply the same thread across centralized, federated, decentralized, and embodied substrates, so the audit guarantee holds whether the agent runs in a data center, at the edge, or in a physical device.

## **Boundary Conditions**

The grounding here is the cognition application, United States Patent Application 19/647,395, which discloses the cross-domain coherence and the cross-tier composition described above. The sibling tiers it composes with (the execution substrate, the canonical agent schema and its lineage field, the state-preserving transport, the cryptographic governance, the adaptive content index, and trust-slope identity) are referenced by category as co-pending portfolio work; their internal specifics are not claimed by this filing and should not be read into it.

Determinism is a precondition, not a free lunch. Reconstruction is exact only to the extent the update functions are deterministic and the relevant observations were recorded in the lineage; non-deterministic external effects outside the recorded observation sequence are outside the reconstruction guarantee. The lineage thread also presumes agents that conform to the disclosed schema and carry their own state; legacy components that do not participate in the schema fall outside the thread and must be bridged at their boundary. Cryptographic provenance is only as strong as the key management and signing discipline operating around it. None of the domain, market, or regulatory framing in this article is a representation about any specific deployment's certification status; it describes an enabling implementation of the disclosed technology, not a compliance guarantee.

## **Disclosure Scope**

The mechanisms described here (the lineage field as a canonical carried-state field, deterministic reconstruction of agent state by replay, the transit cognitive state that continues recording across substrate boundaries, integrity-trajectory governance

authority evaluation, policy-freshness handling on asynchronous resume, substrate identity revocation during cognition, and lineage recording across discovery traversal) are disclosed in United States Patent Application 19/647,395. References to sibling portfolio tiers (execution substrate, agent schema, transport, cryptographic governance, content index, and identity) describe co-pending work by category and are not claims of this filing. The domain framing (regulated operators, incident response, and references to governance expectations such as the EU AI Act and the NIST AI Risk Management Framework) is external context provided to illustrate a faithful enabling implementation; it is not part of the patent claims and is not a legal or regulatory representation about any particular system.

---

## **Cross-Patent Architecture** (</cross-patent-architecture>) [All 40 steps → \(/inventive-steps\)](/inventive-steps) **itecture)**

Cross-cutting architectural principles that compose every primitive into a coherent platform.

[Chapter 1 \(/patents/19-647395/chapters/foundation\)](/patents/19-647395/chapters/foundation).

### **PRIMARY TECHNICAL DISCLOSURE**

- [Cross-Patent Architecture, Articles \(/articles/cross-patent-architecture\)](/articles/cross-patent-architecture)

### **SECONDARY TECHNICAL**

- [Transit Cognitive State \(/articles/cross-patent-architecture/transit-cognitive-state\)](/articles/cross-patent-architecture/transit-cognitive-state)
- [Substrate Identity Revocation During Active Cognition \(/articles/cross-patent-architecture/substrate-identity-revocation\)](/articles/cross-patent-architecture/substrate-identity-revocation)
- [Policy Freshness Across Asynchronous Execution \(/articles/cross-patent-architecture/policy-freshness-asynchronous-execution\)](/articles/cross-patent-architecture/policy-freshness-asynchronous-execution)
- [Governance Authority Evaluation via Integrity Trajectory \(/articles/cross-patent-architecture/governance-authority-integrity-trajectory\)](/articles/cross-patent-architecture/governance-authority-integrity-trajectory)
- [Discovery Agent as Schema-Conformant Index Traverser \(/articles/cross-patent-architecture/discovery-agent-schema-index-traverser\)](/articles/cross-patent-architecture/discovery-agent-schema-index-traverser)

- [Unified Substrate for Governed Information Acquisition \(/articles/cross-patent-architecture/cross-tier-navigation-world-as-model\)](/articles/cross-patent-architecture/cross-tier-navigation-world-as-model).

## APPLICATIONS · GENERAL

- [One Governed Platform, Not Four Integrated Systems: A Unified Architecture Spine for Agent Execution, Cognition, Content, and Spatial Tiers \(/articles/cross-patent-architecture/unified-governed-platform\)](/articles/cross-patent-architecture/unified-governed-platform).
- [World-as-Model Systems: Navigating the Physical World, Cognition, and Discovery as One Governed Model \(/articles/cross-patent-architecture/world-as-model-systems\)](/articles/cross-patent-architecture/world-as-model-systems).
- **[End-to-End Lineage and Audit: Reconstructing Any Agent Action Across Every Tier of the Stack \(/articles/cross-patent-architecture/end-to-end-lineage-and-audit\)](/articles/cross-patent-architecture/end-to-end-lineage-and-audit)**
- [Moving Governed AI Agents Across Clouds and Vendors Without Losing Identity: Substrate Portability via the Cross-Patent Architecture \(/articles/cross-patent-architecture/portability-across-substrates\)](/articles/cross-patent-architecture/portability-across-substrates)
- [Cross-Patent Architecture: Why a Coherent AI Platform Needs a Shared Governance Authority at the Foundation, Not as a Feature \(/articles/cross-patent-architecture/ai-platform-foundation\)](/articles/cross-patent-architecture/ai-platform-foundation)
- [Regulated Cross-Domain Deployment: One Governance Authority and Policy-Freshness Model Across Every Tier of an End-to-End System \(/articles/cross-patent-architecture/regulated-cross-domain-deployment\)](/articles/cross-patent-architecture/regulated-cross-domain-deployment)

---

[Cross-Patent Architecture overview → \(/cross-patent-architecture\)](/cross-patent-architecture)