Nick Clark · Published June 28, 2026 · PDF

Transit Cognitive State

Transit of cognitive state between execution substrates is a secondary inventive step of the Cross-Patent Architecture disclosed in United States Patent Application 19/647,395: agent cognitive state is frozen during transport between substrates while the lineage field accumulates transit events, with confidence evaluation upon arrival.

Transit Cognitive State

Agent cognitive state frozen during transport between substrates while lineage accumulates transit events, with confidence evaluation upon arrival. The transit cognitive state arises from the interaction of the execution mechanisms of Application [6] and the transport mechanisms of Application [3], and applies when a semantic agent is between execution substrates. During transit, the agent is neither executing, since no substrate provides computational resources, nor in non-executing cognitive mode, since no compute is available for speculative reasoning, nor dormant, since the agent's state is actively in transport.

The transit cognitive state freezes the agent's cognitive domain field values at their pre-transit levels while the lineage field continues to accumulate transit events such as departure timestamp, transport path, and arrival validation. Upon arrival at a destination substrate, the confidence governor evaluates whether the transit duration, transit path characteristics, and destination substrate capabilities warrant a confidence adjustment before resuming execution.

Disclosure Scope

This article describes subject matter disclosed in United States Patent Application 19/647,395. Statements regarding the transit cognitive state, the confidence governor, and the carrying of cognitive domain fields across state-preserving transport reflect that disclosure and are not separate guarantees of system behavior.

Cross-Patent Architecture (/cross-patent-architecture)

All 40 steps → (/inventive-steps)

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

Chapter 1 (/patents/19-647395/chapters/foundation)

PRIMARY TECHNICAL DISCLOSURE

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

SECONDARY TECHNICAL

- **Transit Cognitive State (/articles/cross-patent-architecture/transit-cognitive-state)**
- Substrate Identity Revocation During Active Cognition (/articles/cross-patent-architecture/substrate-identity-revocation)
- Policy Freshness Across Asynchronous Execution (/articles/cross-patent-architecture/policy-freshness-asynchronous-execution)
- Governance Authority Evaluation via 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)
- Unified Substrate for Governed Information Acquisition (/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)
- World-as-Model Systems: Navigating the Physical World, Cognition, and Discovery as One Governed Model (/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)
- 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)
- 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)
- 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)

APPLICATIONS · SPECIFIC

- Palantir Foundry and AIP (the ontology-based data/operations platform plus its AI orchestration layer) vs a cross-tier governed architecture: where does end-to-end action attribution live? (/articles/cross-patent-architecture/palantir-foundry-aip)
- Microsoft's integrated AI stack (Azure AI Foundry, Microsoft Fabric, Entra, and Copilot) vs a single cross-domain governance architecture: how do coherence and one governance chain differ from an integrated product suite? (/articles/cross-patent-architecture/microsoft-ai-stack)
- Amazon Web Services' integrated AI/data stack (Bedrock, SageMaker, and surrounding data/identity services) vs a unified cross-tier governed agent architecture (/articles/cross-patent-architecture/aws-ai-stack)
- NVIDIA's full-stack AI platform (NVIDIA AI Enterprise, NIM microservices, and the CUDA/hardware-to-software stack) vs a substrate-independent governance architecture (/articles/cross-patent-architecture/nvidia-ai-enterprise)
- Databricks Data Intelligence Platform (lakehouse plus Mosaic AI, Unity Catalog governance, and agent tooling) vs an agent-resident cross-patent architecture: where governance lives (/articles/cross-patent-architecture/databricks-data-intelligence)

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