



[Home](#) [Licensing](#) [Patents](#) [Articles](#)

Six-Field Canonical Agent Schema: Structural Definition of Autonomous Semantic Agents

by [Nick Clark](#) | Published March 27, 2026 | [PDF](#)

Structural definition of a semantic agent as a persistent object comprising intent, context, memory, policy reference, mutation descriptor, and lineage fields that collectively govern autonomous execution. Within the cognition-native execution platform, this capability operates as a structural primitive at the execution substrate level. It is not an optional enhancement or a configurable plugin but a mandatory architectural property that every participant encounters. The result is a system where six-field canonical agent schema is enforced by construction rather than by convention, policy, or external oversight.

What It Is

Structural definition of a semantic agent as a persistent object comprising intent, context, memory, policy reference, mutation descriptor, and lineage fields that collectively govern autonomous execution. This is a structural mechanism within the cognition-native execution platform that operates at the execution substrate level. It is not advisory, not configurable at the discretion of individual participants, and not dependent on external enforcement infrastructure.

Every interaction within the system encounters this mechanism as a mandatory constraint. The behavior it produces is deterministic: given the same inputs and the same system state, the outcome is identical regardless of which node evaluates it, when the evaluation occurs, or what substrate hosts the computation.

Why It Matters

Conventional execution platforms address this problem through orchestration layers, container schedulers, and external control planes. These approaches function adequately under controlled conditions but introduce structural fragility when the orchestrator becomes unavailable or the system scales beyond centralized coordination. The underlying assumption that a central scheduler can maintain consistent state across all execution instances becomes a liability precisely when reliability matters most.

Six-field canonical agent schema removes this fragility by embedding the relevant capability directly into the execution substrate layer. There is no external dependency that can fail independently, no middleware that can be misconfigured, and no trust assumption that can be violated by a single compromised participant. The guarantee is structural.

How It Works

The mechanism operates through deterministic evaluation embedded in the cognition-native execution platform. When a relevant operation is initiated, the system evaluates the applicable structural constraints against the current state. This evaluation consults the fields, policies, and lineage records that travel with the objects themselves rather than relying on external state that may be stale, unavailable, or compromised.

The outcome of each evaluation is recorded in an append-only lineage structure. This record is cryptographically committed, ensuring that the complete history of decisions, transitions, and state changes remains auditable and tamper-evident. No evaluation outcome can be retroactively altered without breaking the cryptographic chain.

Because the evaluation logic and the data it operates on travel together, the mechanism functions identically across network partitions, substrate migrations, and administrative boundaries. There is no central evaluation point that must be available for the system to operate correctly.

What It Enables

With six-field canonical agent schema as an architectural primitive, systems built on this foundation can operate autonomously while maintaining the structural guarantees that centralized architectures achieve through oversight. The capability is not a tradeoff between autonomy and governance but a resolution of the apparent conflict between them.

This enables deployment across centralized cloud infrastructure, federated multi-party environments, fully decentralized networks, and edge installations with intermittent connectivity. The structural guarantees hold regardless of deployment topology because they are properties of the objects and protocols themselves, not properties of the infrastructure that hosts them.

[Execution Platform All 21 steps →](#)

The complete runtime for governed, persistent agents.

Patent

[US 19/230,933](#) · filed

Primary Technical Disclosure

[◦ A Cognition-Native Execution Platform for Distributed, Stateful, and Governable Agents](#)

Secondary Technical

[◦ Six-Field Canonical Agent Schema: Structural Definition of Autonomous Semantic Agents](#)◦ [Semantic Nest Instantiation: Dynamic Execution Environments From Agent Density and Entropy](#)◦ [Trust Zone Overlay Governance: Logical Policy Domains Independent of Network Topology](#)◦ [Scoped Quorum Mutation Validation: Independent Validators With Meta-Policy Escalation](#)◦ [Meta-Policy Override Resolution: Higher-Level Governance for Local Quorum Decisions](#)◦ [Semantic Router: Schema-Aware Propagation Replacing Address-Based Forwarding](#)◦ [Dynamic Agent Hash Derivation: Deterministic Identity From Memory and Mutation History](#)◦ [Dynamic Device Hash Derivation: Substrate Identity From Device-Local Entropy](#)◦ [Content Anchor Hash Derivation: Perceptual Identity for Non-Executing Digital Content](#)◦ [DAH-DDH Slope Entanglement: Binding Agent Identity to Host Device Lineage](#)◦ [Trust Slope Validation Across Zone Migration: Continuity Verification With Quarantine](#)◦ [Pseudonymous Propagation: Recognition by Slope Rather Than Global Identifier](#)◦ [Alias Slope-Band Indexing: Symbolic Resolution Through Trust-Slope Pathfinding](#)◦ [Fallback Rehydration: Recovering Partial Agents Through Contextual Policy Inference](#)◦ [Structural Validator With Fallback Routing: Schema Verification Before Execution](#)◦ [Execution Graph Manager: Structured Lineage of Agent Reasoning and Transformation](#)◦ [Full and Partial Agent Interoperability: Cross-Boundary Semantic Exchange Under Policy](#)◦ [Cross-Topology Substrate Deployment: Identical Agent Structure Across All Substrates](#)

Applications (General)

[◦ Multi-Cloud Agent Orchestration Without Centralized Schedulers](#)◦ [Autonomous Fleet Coordination Through Self-Governing Agents](#)◦ [Enterprise Workflow Without Orchestration Servers](#)◦ [Smart Contract Execution Without Blockchain Latency](#)◦ [Distributed Scientific Computing With Governed Agents](#)◦ [Supply Chain Autonomous Agents](#)◦ [Energy Grid Management Through Autonomous Agents](#)◦ [Disaster Response Coordination Without Central Command](#)

Applications (Specific)

[◦ Kubernetes Orchestrates Containers. It Does Not Know What They Are Doing.](#)◦ [Temporal Solved Durable Workflows. The Workflows Have No Semantic Identity.](#)◦ [Apache Airflow Orchestrates DAGs. The Tasks Inside Them Are Ungoverned.](#)◦ [Prefect Made Data Workflows Pythonic. The Execution Model Is Still Task Scheduling.](#)◦ [AWS Step Functions Made Serverless Orchestration Visual. The Steps Have No Semantic State.](#)◦ [Azure Durable Functions Made Stateful Serverless Possible. The State Has No Governance.](#)◦ [Nomad Schedules Any Workload. It Does Not Know What Those Workloads Are.](#)◦ [Docker Swarm Simplified Container Orchestration. The Containers Are Still Opaque.](#)◦ [Apache Mesos Managed Datacenter Resources.](#)

[The Resources Had No Semantic Governance.](#)◦ [Argo Workflows Orchestrates Kubernetes-Native Pipelines. The Pipeline Steps Have No Governance.](#)◦ [Dagster Made Data Pipelines Software-Defined. The Pipeline Has No Governance Substrate.](#)◦ [Luigi Defined Task Dependencies for Data Pipelines. The Tasks Execute Without Governance.](#)◦ [Camunda Orchestrates Business Processes. The Process Engine Has No Semantic Agent Governance.](#)◦ [Zeebe Scaled Workflow Orchestration Horizontally. Governance Did Not Scale With It.](#)
[Execution Platform overview →](#)

AQ

deterministic

autonomy

Legal

Subject to one or more pending U.S. and international patent applications, see [Patents](#) for the current list and status. No license, express or implied, is granted. Any use requires a separate written agreement—see [Licensing](#). Patent applications referenced on this site are pending. Claim scope, if any, is subject to examination and may issue in altered form or not at all. See [Legal](#) for terms and conditions.

Adaptive Query™ is a trademark of Nicholas Clark. U.S. federal registration is pending. federal registration. AQ™, AQ Inside™, Adaptive Index™, Adaptive Network™, Semantic Agent™, @AQ™, AQID™, and Adaptive Coin™ are used as trademarks in connection with the Adaptive Query platform and brand. Other names may be trademarks of their respective owners.

Platform operated by Adaptive Query LLC, which provides patent and trademark licensing services. Copyright © 2025-2026 Nicholas Clark. All rights reserved.

Last updated: 2026-03-03



-
- [Inventive Steps](#)
- [Licensing](#)
- [Patents](#)
- [Articles](#)
- [Legal](#)
- [Opportunities](#)
- [Sitemap](#)



-
- nick@qu3ry.net
- 72 28 14 36 01



[Invented by Nick Clark](#) | Founding Investors: Devin Wilkie