



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

Serverless Execution Without Cold Starts or State Loss

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

Serverless computing promised execution without infrastructure management. It delivered execution without state. Every Lambda invocation, every Cloud Function trigger starts from zero, reconstructing context from external databases and caches before doing useful work. Memory-resident execution offers a structural alternative: persistent semantic objects that carry their own execution state, self-evaluate readiness, and resume where they left off without cold starts, external state stores, or orchestration frameworks.

The state problem in serverless

Serverless platforms achieve scalability by treating functions as stateless. Each invocation receives a clean execution environment. Any state the function needs must be fetched from external storage: DynamoDB, Redis, S3. This external state fetch is the cold start penalty that serverless users accept as a cost

of scalability.

For simple request-response workloads, the penalty is tolerable. For AI agent workloads that accumulate context across interactions, maintain conversation history, track long-running tasks, and update their understanding of the world, the penalty is architectural. Every invocation must reconstruct the agent's entire cognitive state from serialized storage before the agent can think. The agent starts every interaction with amnesia and must read its own diary before responding.

Durable execution frameworks like Temporal and Durable Functions attempt to solve this by persisting execution state in external stores and replaying event histories. This solves durability but introduces its own overhead: event sourcing complexity, replay latency, and a tight coupling between the execution logic and the persistence framework.

Why warm containers do not solve the state problem

Cloud providers offer warm container pools and provisioned concurrency to reduce cold start latency. These keep execution environments alive between invocations, preserving in-memory state. But warm containers are a caching optimization, not a structural solution. The platform makes no guarantee about which container handles the next invocation. State affinity is best-effort. For stateful AI agents, best-effort state preservation means unpredictable behavior.

The fundamental issue is that serverless architectures locate execution state outside the executing entity. The function is ephemeral. The state is in a database. Bringing them together requires coordination that grows more expensive as state complexity increases.

How memory-resident execution addresses this

Memory-resident execution inverts the serverless model. Instead of ephemeral functions fetching external state, persistent semantic objects carry their own state and execution logic as intrinsic properties. The object does not start from zero. It resumes from its last evaluated state, because that state is part of the object itself.

Each memory-resident object carries an execution cycle that it manages independently: evaluate current state, determine if conditions warrant action, execute if eligible, update state, and enter dormancy until the next evaluation trigger. No external orchestrator schedules these cycles. The object self-evaluates based on its own wake triggers and semantic backoff parameters.

Dormancy is a first-class state, not a terminated process. A dormant object retains its full memory and governance state. When a wake trigger fires, the object resumes execution from its dormant state without reconstruction. There is no cold start because the object never fully stopped. It simply was not actively executing.

For AI agents, this means the agent's conversation history, accumulated context, learned preferences, and in-progress tasks persist as part of the agent object. When the agent is needed, it resumes with full context. When it is not needed, it enters dormancy at near-zero resource cost. The execution model matches the cognitive model: agents that persist and resume rather than functions that start and stop.

What implementation looks like

An enterprise deploying memory-resident execution replaces stateless function invocations with persistent agent objects. Each agent carries its own state, governance, and execution logic. The deployment infrastructure provides substrate for persistence rather than scheduling for execution.

For customer service platforms, this means each customer interaction is handled by an agent that remembers the full relationship history. The agent does not look up the customer's record in a database. The agent is the record, carrying the interaction history as part of its own memory state.

For workflow automation, memory-resident objects replace multi-step orchestrated pipelines with self-executing objects that carry the workflow state through each step. A purchase order object that needs three approvals manages its own approval sequence, entering dormancy between approvals and waking when the next approval arrives. No orchestrator tracks the workflow. The object tracks itself.

[Memory-Resident Execution All 21 steps →](#)

Persistent objects that execute without orchestration.

Patent

[US 19/538,221](#) · filed

Primary Technical Disclosure

[◦ Memory-Resident Execution: Persistent Semantic Objects Without Orchestration](#)

Secondary Technical

[◦ Six-Action Execution Evaluation Cycle: Parse, Evaluate, Select at Every Node](#)

[◦ Cognition-Authority-Execution Separation: Reasoning Cannot](#)

[Authorize Action](#)

[◦ Dormancy as First-Class Execution State: Valid Suspension Without Failure](#)

[◦ Semantic Backoff: Retry Pacing From Execution](#)

[Outcomes Rather Than Fixed Timers](#)

[◦ Wake Triggers for Dormancy Exit: Explicit Reentry Conditions in Memory](#)

[◦ Persistent Polling Behavior:](#)

[Autonomous Condition Evaluation Without Schedulers](#)

[◦ Intent Refinement During Execution: Adaptive Objectives Without Re-Instantiation](#)

[◦ Compositional Execution Through Recursive Delegation: Parent-Child Lineage Tracking](#)

[◦ Negative Capability Signals: Recording What Cannot Be Done](#)

[as Structured Constraint](#)

[◦ Swarm-Based Execution Emergence: Coordinated Behavior Without Centralized Control](#)

[◦ Latency and Failure as Semantic](#)

[Signals: Structured Inputs From Adverse Conditions](#)

[◦ LLM as Advisory Execution Node: Inference Without Authority Over Agent State](#)

[◦ Append-Only](#)

[Memory Field: Complete Execution Lineage Through Immutable Records](#)

Applications (General)

[• Serverless Execution Without Cold Starts or State Loss](#)

[◦ Long-Running Autonomous Workflows Without External Orchestration](#)

[◦ Drone Operations](#)

[Surviving Disconnection](#)

[◦ Deep Space Agent Execution Without Ground Control](#)

[◦ Underwater Robotic Operations Without Connectivity](#)

[◦ Rural](#)

[Healthcare Agents Surviving Intermittent Connectivity](#)

[◦ Operations in Infrastructure-Destroyed Environments](#)

[◦ Offline Financial Transaction Agents](#)

Applications (Specific)

[◦ Cloudflare Durable Objects Made State Local. The Objects Still Need Orchestration.](#)[◦ Azure Service Fabric Actors Are Addressable. They Are Not Autonomous.](#)[◦ Akka Perfected the Actor Model. Actors Still React Instead of Self-Execute.](#)[◦ Orleans Made Virtual Actors Practical. The Actors Still Execute on Request.](#)[◦ Dapr Provides a Sidecar Runtime for Microservices. The Services Still Need External Orchestration.](#)[◦ wasmCloud Runs WebAssembly Actors. The Actors Wait for Messages.](#)[◦ Spin Made WebAssembly Serverless. The Functions Are Still Trigger-Based.](#)[◦ Fermion Built the WebAssembly Cloud. The Cloud Hosts Functions, Not Self-Executing Objects.](#)[◦ Fly Machines Made Micro-VMs Fast. The VMs Still Need External Orchestration.](#)[◦ Railway Simplified Application Deployment. The Applications Still Depend on External Execution Triggers.](#)
[Memory-Resident Execution 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