

Memory Layers for Agents: Why Memo, Zep, and Letta Get Close

A whole category, Memo, Zep, and Letta, formerly MemGPT, has formed around giving language-model agents persistent memory. They get the diagnosis right: stateless inference is a dead end. They stop short of governance: no admissibility gate at retrieval or generation, no carried lineage of why a memory surfaced, no structural control over who sees which slice.

A Category That Got the Problem Right

A distinct product category has formed around a correct diagnosis: stateless inference is a dead end for agents. Memo, Zep, and Letta, the last formerly MemGPT, each give a language-model agent persistent memory so that it can carry knowledge across turns and sessions instead of starting blank on every call. As described in their public materials, the shared shape is consistent. Conversation history and facts extracted from it are written to a vector store or a structured memory, and on each new turn the system retrieves the slices judged relevant and re-injects them into the model's context window. The agent thereby appears to remember, because the orchestration layer keeps reminding it. The diagnosis is right, the engineering is real, and the category exists because the need is genuine.

The discovery substrate shares the diagnosis exactly. An agent without persistent, structured state cannot reason coherently over time, and memory is not an add-on to cognition but a precondition for it. Where the substrate diverges is in what memory is permitted to be: not a store that an orchestration layer reads from and re-injects, but a typed object the agent carries, governed at the moment it is read and at the moment it is used.

What's Missing Is Governance

Re-injection has no admissibility gate. When a memory framework retrieves a slice and places it in the prompt, nothing evaluates whether that retrieval is permitted, whether the memory is current and unrevoked, or whether using it would violate a policy that applies to this query. The decision is a similarity match, and the constraint, if any, is whatever the prompt manages to encode. There is likewise no carried record of why a given memory was injected that travels with the agent's output, so the provenance of a remembered fact is opaque after the fact. And in any setting where memory crosses users or organizations, there is no structural control over who may see which slice; sharing is governed, if at all, by application code around the store rather than by the memory itself. The frameworks remember well; they do not govern what is remembered, why it surfaces, or who may see it.

The Discovery Object Is the Structural Form They Approximate

In the discovery substrate, the unit of memory is the discovery object: a persistent, memory-resident semantic entity that carries the full context of a traversal as typed fields, not a query string or an embedding. Its intent field holds what the traversal seeks; its memory field holds accumulated semantic commitments as structured records rather than as re-injected text; its policy reference field holds the governance constraints that apply; its lineage field holds the record of every admitted and rejected step. Retrieval is traversal, and every traversal step is admissibility-gated, so a memory

is surfaced only if reaching it is permitted under the current policy, lineage continuity, and temporal validity, and the reason it was surfaced is recorded in lineage. Addressing is navigational rather than table-based: a structured alias resolves by walking the live index under the same governance, so possessing a reference to a memory does not confer access to it. And when multiple agents' traversals intersect, memory is shared through a policy-governed collaborative merge that permits the exchange only if both parties' policies allow it and records the exchanged elements in both lineages. The features these products build around the store, persistence, retrieval relevance, and cross-agent sharing, are present in the substrate as properties of a governed object rather than as orchestration around an ungoverned one.

Complementary Positioning

This is a difference of structural foundation, not of ambition, and the two are composable. A memory framework that has solved persistence, extraction, and retrieval relevance has built much of what an agent needs; what it lacks is the governance layer that makes remembering safe in regulated, multi-tenant, and cross-organization settings. The substrate supplies that layer by making memory a typed, admissibility-gated, lineage-bearing object rather than a store to be re-injected. An agent platform can adopt the discovery object as its memory form and gain per-retrieval admissibility, carried provenance, and governed cross-agent sharing without abandoning the persistence and extraction work it has already done.

Disclosure Scope

The discovery object as a traversal-native semantic entity carrying intent, context, memory, policy, and lineage as typed fields is disclosed in the cognition filing (U.S. Application No. 19/647,395 and its international counterpart) at Section 10.2; navigational, governance-integrated alias resolution at Section 10.7; and policy-governed collaborative multi-discovery-object traversal at Section 10.19. This article

compares those disclosed mechanisms with the publicly described persistent-memory frameworks Memo, Zep, and Letta, and positions the discovery object as the governed structural form those frameworks approximate. References to those products are to their public materials and are used for comparison only; no relationship, endorsement, or infringement is asserted.

Semantic Discovery (</semantic-discovery>)

[All 36 steps → \(/inventive-steps\)](/inventive-steps)

Search, inference, and execution as one governed step.

PRIMARY TECHNICAL DISCLOSURE

- [Governed Semantic Discovery: Search, Inference, and Execution Through Adaptive Traversal \(/articles/governed-semantic-discovery-search-inference-and-execution-through-adaptive-traversal\)](/articles/governed-semantic-discovery-search-inference-and-execution-through-adaptive-traversal)

SECONDARY TECHNICAL

- [The Adaptive Index as Unified Search-Inference-Execution Substrate \(/articles/semantic-discovery/unified-substrate\)](/articles/semantic-discovery/unified-substrate)
- [Three-in-One Traversal: Search, Inference, and Execution in a Single Step \(/articles/semantic-discovery/three-in-one-traversal\)](/articles/semantic-discovery/three-in-one-traversal)
- [The Discovery Object: A Traversal-Native Semantic Agent \(/articles/semantic-discovery/discovery-object\)](/articles/semantic-discovery/discovery-object)
- [Post-PageRank Semantic Ranking: Relevance Through Governed Traversal \(/articles/semantic-discovery/post-pagerank\)](/articles/semantic-discovery/post-pagerank)
- [Persistent Semantic State: Eliminating Prompt Reconstruction \(/articles/semantic-discovery/persistent-state\)](/articles/semantic-discovery/persistent-state)
- [Traversal Lineage as Index Evolution Signal \(/articles/semantic-discovery/traversal-lineage\)](/articles/semantic-discovery/traversal-lineage)
- [Anchor Semantic Neighborhood Publication \(/articles/semantic-discovery/semantic-neighborhoods\)](/articles/semantic-discovery/semantic-neighborhoods)
- [Inference-Time Execution Control as Traversal Primitive \(/articles/semantic-discovery/inference-governance\)](/articles/semantic-discovery/inference-governance)
- [Anchor Self-Organization Under Entropy and Load Pressure \(/articles/semantic-discovery/anchor-self-organization\)](/articles/semantic-discovery/anchor-self-organization)

- [Alias Resolution as Navigational Traversal \(/articles/semantic-discovery/alias-resolution\)](/articles/semantic-discovery/alias-resolution)
- [Three Discovery Operating Modes: Human Search, Agent Reasoning, Answer Synthesis \(/articles/semantic-discovery/operating-modes\)](/articles/semantic-discovery/operating-modes)
- [Model-Agnostic Semantic Discovery \(/articles/semantic-discovery/model-agnostic\)](/articles/semantic-discovery/model-agnostic)
- [Affect-Modulated Discovery Traversal \(/articles/semantic-discovery/affect-modulated-traversal\)](/articles/semantic-discovery/affect-modulated-traversal)
- [Confidence-Gated Discovery Traversal \(/articles/semantic-discovery/confidence-gated-traversal\)](/articles/semantic-discovery/confidence-gated-traversal)
- [Integrity-Tracked Traversal Drift Detection \(/articles/semantic-discovery/integrity-tracked-drift\)](/articles/semantic-discovery/integrity-tracked-drift)
- [Biological Identity-Scoped Access During Discovery \(/articles/semantic-discovery/biological-access\)](/articles/semantic-discovery/biological-access)
- [Rights-Grade Anchor Governance for Content Discovery \(/articles/semantic-discovery/rights-grade-anchors\)](/articles/semantic-discovery/rights-grade-anchors)
- [Forecasting-Shaped Discovery Traversal \(/articles/semantic-discovery/forecasting-shaped\)](/articles/semantic-discovery/forecasting-shaped)
- [Capability-Constrained Anchor Accessibility \(/articles/semantic-discovery/capability-constrained\)](/articles/semantic-discovery/capability-constrained)
- [Collaborative Multi-Object Discovery Traversal \(/articles/semantic-discovery/collaborative-traversal\)](/articles/semantic-discovery/collaborative-traversal)
- [Discovery-Driven Sensor Invocation Closed Loop \(/articles/semantic-discovery/sensor-invocation-loop\)](/articles/semantic-discovery/sensor-invocation-loop)
- [Cross-Platform Credentialed Reader Activation \(/articles/semantic-discovery/credentialed-reader-activation\)](/articles/semantic-discovery/credentialed-reader-activation)
- [LLM-as-Bootstrap: Why Anchor Inference Engines Shrink as the Lineage Matures \(/articles/semantic-discovery/maturation-engine-substitution\)](/articles/semantic-discovery/maturation-engine-substitution)
- [Personal Cognitive Asset: How Per-User Lineage Re-Weights the Same Substrate \(/articles/semantic-discovery/personal-lineage-layer\)](/articles/semantic-discovery/personal-lineage-layer)
- [Loki, the Dog, and the Symbol Grounding Problem \(/articles/semantic-discovery/hybrid-symbol-grounding\)](/articles/semantic-discovery/hybrid-symbol-grounding)

APPLICATIONS · GENERAL

- [Enterprise Knowledge Management Through Governed Traversal \(/articles/semantic-discovery/enterprise-knowledge-management\)](/articles/semantic-discovery/enterprise-knowledge-management)
- [AI-Native Search That Replaces PageRank With Contextual Relevance \(/articles/semantic-discovery/ai-native-search\)](/articles/semantic-discovery/ai-native-search)
- [Semantic Discovery for Scientific Research \(/articles/semantic-discovery/scientific-research-discovery\)](/articles/semantic-discovery/scientific-research-discovery)
- [Semantic Discovery for Legal Case Research \(/articles/semantic-discovery/legal-case-research\)](/articles/semantic-discovery/legal-case-research)
- [Semantic Discovery for Patent Landscape Analysis \(/articles/semantic-discovery/patent-landscape-analysis\)](/articles/semantic-discovery/patent-landscape-analysis)

- [Semantic Discovery for Medical Literature Search \(/articles/semantic-discovery/medical-literature-search\)](/articles/semantic-discovery/medical-literature-search).
- [Semantic Discovery for Competitive Intelligence \(/articles/semantic-discovery/competitive-intelligence\)](/articles/semantic-discovery/competitive-intelligence).
- [Semantic Discovery for Regulatory Compliance Search \(/articles/semantic-discovery/regulatory-compliance-search\)](/articles/semantic-discovery/regulatory-compliance-search).
- [Discovery-Coordinated Multi-Sensor Perception \(/articles/semantic-discovery/coordinated-perception\)](/articles/semantic-discovery/coordinated-perception).
- [Post-AirTag Cross-Platform Object Tracking \(/articles/semantic-discovery/post-airtag-tracking\)](/articles/semantic-discovery/post-airtag-tracking).
- [Use the World as Memory: The Brain Strategy for AI \(/articles/semantic-discovery/world-as-memory\)](/articles/semantic-discovery/world-as-memory).

APPLICATIONS · SPECIFIC

- [Google Search Retrieves Results, Not Understanding \(/articles/semantic-discovery/google-search\)](/articles/semantic-discovery/google-search).
- [Perplexity Answers Questions Without Discovery State \(/articles/semantic-discovery/perplexity\)](/articles/semantic-discovery/perplexity).
- [Elasticsearch Indexes Documents, Not Discovery \(/articles/semantic-discovery/elasticsearch\)](/articles/semantic-discovery/elasticsearch).
- [Algolia Optimizes Relevance Without Discovery State \(/articles/semantic-discovery/algolia\)](/articles/semantic-discovery/algolia).
- [Pinecone Finds Vectors, Not Understanding \(/articles/semantic-discovery/pinecone\)](/articles/semantic-discovery/pinecone).
- [Weaviate Stores Semantics Without Discovery Governance \(/articles/semantic-discovery/weaviate\)](/articles/semantic-discovery/weaviate).
- [You.com Answers Questions but Does Not Govern Discovery \(/articles/semantic-discovery/you-com\)](/articles/semantic-discovery/you-com).
- [Brave Search Built an Independent Index Without Governed Traversal \(/articles/semantic-discovery/brave-search\)](/articles/semantic-discovery/brave-search).
- [Kagi Charges for Better Results, Not Governed Discovery \(/articles/semantic-discovery/kagi\)](/articles/semantic-discovery/kagi).
- [Metaphor Systems Predicts Links but Does Not Govern Traversal \(/articles/semantic-discovery/metaphor-systems\)](/articles/semantic-discovery/metaphor-systems).
- [Glean Indexes Enterprise Knowledge Without Governing Its Discovery \(/articles/semantic-discovery/glean\)](/articles/semantic-discovery/glean).
- [Coveo Personalizes Retrieval, Not Discovery Governance \(/articles/semantic-discovery/coveo\)](/articles/semantic-discovery/coveo).
- [Apple Find My Lacks Cross-Authority Reader Activation \(/articles/semantic-discovery/apple-find-my\)](/articles/semantic-discovery/apple-find-my).
- [Google Find My Network Needs Credentialed Cross-Activation \(/articles/semantic-discovery/google-find-my\)](/articles/semantic-discovery/google-find-my).
- [IETF DULT Specifies Behavior, Not Architecture \(/articles/semantic-discovery/ietf-dult\)](/articles/semantic-discovery/ietf-dult).

- [Glean Enterprise Search and Work AI \(/articles/semantic-discovery/glean-enterprise-search\)](/articles/semantic-discovery/glean-enterprise-search)
- [GraphRAG, but with Governance: Where Microsoft's Architecture Stops Short \(/articles/semantic-discovery/microsoft-graphrag\)](/articles/semantic-discovery/microsoft-graphrag)
- **[Memory Layers for Agents: Why Mem0, Zep, and Letta Get Close \(/articles/semantic-discovery/memory-for-agents\)](/articles/semantic-discovery/memory-for-agents)**

[Semantic Discovery overview → \(/semantic-discovery\)](/semantic-discovery)