Nick Clark · Published July 29, 2026 · PDF

Letta: Agent Memory and Standing Earned From Conduct

An agent that retains everything it has done still has to answer a separate question: on what basis may a stranger act on that record. Letta, as publicly described, is a framework for building stateful agents whose memory persists across sessions, so an agent's history is durable rather than reconstructed each time. A filed provisional addresses an adjacent question: what happens between two agents when one executes under another's recorded determination. In the filed architecture no rating, score, standing quantity, or credential passes between them, and where the followed agent elects to counter-sign, that append decrements its own authorization budget, the same budget gating its own dispatch. This article compares the two structurally.

An agent with durable recall, and a stranger who must decide

A procurement agent has run a long series of vendor evaluations and kept every one. A second agent, built by a different team inside a different company, wants to act on one of them without re-deriving it. It is looking at records keyed to another party's internal identifiers, held in another party's storage, and it must decide what verification it can perform and what, if anything, the first agent bears when they are acted upon.

The familiar repair is a number. Give each agent a score, publish it, let others read it. The filed architecture takes a different position, for a structural reason. In conventional delegation and capability-attenuation designs, cost falls on the relying party while the party relied upon is protected, since an attenuation on delegation narrows only the capability passed downward. Adoption, on that arrangement, is costless to the adopted.

The provisional inverts both halves. Nothing quantitative moves between the agents: no rating, score, standing quantity, or credential. The metered consequence falls on the party being relied upon, in the units of that party's own capacity to dispatch. What the relying agent obtains is not standing but exposure, written only on a counter-signature that has been paid for.

What public materials describe Letta as building

Letta, as publicly described, is a framework for building stateful agents: agents whose memory persists across sessions rather than being rebuilt from a transcript each time the process starts. Public materials present the design as a memory hierarchy, with some information in the agent's immediate context and the rest in external storage the agent can search and retrieve from. What those materials treat as distinctive is that the agent is given tools to manage that boundary itself: it can revise its own memory, influence what is brought into context, and write back what should survive the session. The agent is presented as a durable entity rather than a stateless function called with a prompt.

That is a hard problem, and it sits upstream of everything discussed below. An agent cannot present a coherent account of its own conduct if it does not retain one. Teams reach for a framework of this kind because they want an agent that is the same agent tomorrow.

The two designs are addressed to different layers. Public materials position Letta at agent state and memory persistence inside a single agent. The filed architecture is addressed to what two independently operated agents can do with each other's records across an organizational boundary, and what an adoption costs the agent adopted. The second layer assumes durable state of roughly the first kind, since the verification it specifies runs over records the followed agent must hold and commit to.

The filed mechanism, in its own terms

In the filed architecture a first semantic agent (100), the deferring agent, holds a request to dispatch an action and finds its own dispatch-authority predicate satisfied on recomputation from its memory field (102). It is authorized on its own record to determine the matter for itself. It nonetheless elects to execute under a determination already recorded in the lineage field (104) of a second agent, the followed agent, rather than under one produced by its own admission evaluator (120). An agent whose predicate fails, whose authorization gate (300) stands in the withheld state (310), or which has entered the non-executing cognitive mode (302) does not defer at all. Its holding is a withholding.

Before a deference record (900) is given effect, the deferring agent applies the origin-disjointness test: whether it and the followed agent are assignable to a common origin-equivalence class (200). Where they are, the record is appended as an attempted deference record and given no effect. It decrements no budget, enrolls no class, and appends nothing adverse to the followed agent.

Three ordered conjuncts follow, in a retrievability-form verification over the followed determination record. The presented authority credential and continuity hash field are verified against the counterparty identity record (114). The tested-entry enumeration is verified against a lineage commitment carrying a lineage-extent value and an accumulator root, with positions pairwise distinct and strictly increasing. And the determination must be a member of the closed set, with a matching action class. That

verification is of form and of retrievability and not of merits: the deferring agent does not retrieve the enumerated entries, does not re-derive the determination, and produces no determination concerning the followed agent's conduct.

The consequential step is conditional on the followed agent. Having verified the record, and electing to counter-sign, it appends a deference counter-signature (902) as the second governed observation (610) of a matched pair (600). By that append it decrements its own authorization budget (404) under a deference decrement schedule declared in its signed policy object (112). The schedule declares an amount not less than that declared for issuance of a conduct evaluation artifact (116) by that agent, and is a monotonically non-decreasing function of a count of counter-signatures appended within a declared window, so that the marginal cost of being followed rises with the rate of adoption. The disclosed forms are a step function and a form linear in the count; the parameters are policy-declared, and the filing recites no fixed values for them beyond one illustrative linear embodiment.

Merit plays no part in that decrement. No adjudication of correctness by either agent, either principal, or any other party conditions it. The units decremented are the units that gate the followed agent's own dispatch. Responsive to that budget satisfying a floor declared in its own policy object, its authorization gate (300) is written to the withheld state (310) for an enumerated set of action classes. Replenishment runs by its ordinary procedure and by no other, unaffected by how much the agent has been followed, so being followed spends a replenishable quantity and does not terminate the agent.

What the deferring agent receives is not standing. Responsive to receipt of the counter-signature, and not otherwise, it retrieves a corrector-class enumeration (904): the origin-equivalence classes (200) recorded both as having corrected the followed agent in that action class and as having had that correction thereafter conformed to. Both parts of the conformance filter (906) must hold, and a class satisfying correction without subsequent conformity is excluded. Each class is identified by a class-defining reference rather than by the followed agent's local identifier, meaningless outside that

agent. The deferring agent then appends an exposure edge record (910), writes each class into its own replenishment register (402) as an admissible-but-unexercised entry, and dispatches a presentation of its own execution records in that action class to the address set of each class, in the categories its counterparty scope object admits.

Such an entry has two effects and no third. A conduct evaluation artifact (116) from an asserting party (118) assignable to the class becomes admissible to the deferring agent's admission evaluator (120) notwithstanding no prior encounter history, and is not held in the pre-settlement inert state on that ground. And while the exercise field stands unexercised, the entry is not counted in renewal: the agent's quantities stand as they would have stood had it never been written. What was earned is a hearing.

Where the two designs diverge

Divergence turns on what has to cross the boundary between agents. Memory design concerns an agent's relationship to its own state: what it retains, what it retrieves, what it revises. The filed architecture concerns what a second agent may verify about a first agent's records without trusting that agent's storage, and what the first agent bears when the second acts on them.

Two requirements make the filed design structurally distinct, and neither is a memory question. The first is the merit-independent decrement: adoption is metered against the adopted party's own capacity to dispatch, with no party adjudicating whether the adoption was justified. The second is the conversion bar (502), which forbids any stage from converting a deference abstention entry, or a count of corrections received, into a scalar, a default value, a rating, or an operand of a threshold comparison. The architecture therefore requires that no count of adoptions or corrections ever be consumable as a standing quantity by an admission or dispatch decision, a constraint on the whole pipeline rather than a property of any store.

Positioning between the layers is complementary. An agent that cannot maintain a coherent lineage field cannot present a tested-entry enumeration at all. Durable agent state is a precondition for the protocol, not a competitor to it.

Coexistence, and the limits of the filing

Consider a deployment in which each participating agent is built on a stateful framework and the deference protocol runs between organizations. The framework owns retention, revision, and retrieval within an agent. The protocol owns the seam: the origin-disjointness test, the form verification, the counter-signature and its decrement where one is given, the enrollment of corrector classes, and the exercised transition record (908) binding a correction received back to the deference that earned the exposure through which it arrived.

Several things sit outside the filed architecture. It does not make a determination correct; nothing in the mechanism evaluates the merits of an adopted determination, by design. It does not manage context windows, choose what an agent should remember, or specify a storage engine. It does not establish identity from nothing, operating instead on identity primitives and continuity history stores the agents already hold. It does not give the deferring agent credibility with the classes it reaches: the filing is explicit that nothing transmitted describes the followed agent's standing and that no rating is received from those parties. And where a followed agent declines to counter-sign, expressly or by non-response within the carried window, the deferring agent may still execute, under an ungoverned reliance record together with a narrowing of its own counterparty scope object for that agent, which is not a sanction and is released upon a counter-signature within the reinstatement window.

Which layer to evaluate depends on which question is live: whether the agent can still account for itself next week, or the terms on which a counterparty may act on that account and what a counter-signature costs the agent acted upon.

Disclosure Scope

This article describes subject matter disclosed in U.S. Provisional Application No. 64/117,812, a pending application. Descriptions of the deference record (900), deference counter-signature (902), authorization budget (404), replenishment register (402), conformance filter (906), and exposure edge record (910) reflect that filing. Quantities described as policy-declared are declared in an agent's signed policy object (112); the filing recites no fixed values for them, and none should be inferred here.

References to Letta are to public materials and are used for comparison only; no relationship, endorsement, or infringement is asserted.

Deference and Earned Exposure (/earned-exposure)

All 49 steps → (/inventive-steps)

Standing earned through recorded conduct; deference that survives across agents.

Provisional application

PRIMARY TECHNICAL DISCLOSURE

- Earned Exposure and Deference Records: Standing Built From Recorded Conduct (/articles/earned-exposure/earned-exposure-and-deference-records)

SECONDARY TECHNICAL

- The Valuation-Free Ordinal Standing Quantity (/articles/earned-exposure/valuation-free-ordinal-standing)
- Determination-Invariance Under an Agent's Own Policy Succession (/articles/earned-exposure/determination-invariance-policy-succession)
- Elective Exposure Enrollment Without Deference (/articles/earned-exposure/elective-exposure-without-deference)
- Interventional Attribution by Receipt-Path Control (/articles/earned-exposure/interventional-attribution-receipt-path)

- The Portable Exoneration Record (/articles/earned-exposure/portable-exoneration-record)
- Attribution, Exoneration, and the Domain Baseline: An Ordered Branch Over the Latency Differential (/articles/earned-exposure/attribution-versus-exoneration-branch-domain-baseline)
- Bounded Count of Dispatches Per Deference Record: Capping How Many Executions One Counter-Signed Adopted Determination Authorizes (/articles/earned-exposure/bounded-count-dispatches-deference-record)
- Coarse Lineage-Segment Granularity: Verifying a Tested-Entry Enumeration Without Disclosing Which Lineage Entries Were Tested (/articles/earned-exposure/coarse-lineage-segment-granularity-tested-entry)
- Correction-Exposure Decay: An Authorization Quantity That Depletes With Agent Operation and Renews Only on Recorded External Evaluation (/articles/earned-exposure/correction-exposure-decay-authorization-quantity)
- Cause-Keyed Gate Restoration: A Distinct Return Path for Decay-Induced Withholding of Agent Authorization (/articles/earned-exposure/distinct-return-path-decay-induced-withholding)
- Edge Severance Events and the Three-Valued Survival State: Resolving Whether a Counterparty Relationship Outlived a Payment-Class or Obligation-Class Termination (/articles/earned-exposure/edge-severance-events-three-valued-survival)
- Emitting the Deference Correction Linkage Record to the Followed Agent: Reporting a Downstream Correction Without Producing a Rating (/articles/earned-exposure/emission-deference-correction-linkage-record-followed)
- The No-Address Omission Branch: What a Deferring Agent Records When an Exposure Dispatch Reaches Nobody (/articles/earned-exposure/exposure-dispatch-no-address-omission-branch)
- Channel-Baseline Normalization: Taking Channel Overhead Out of the Latency Differential (/articles/earned-exposure/latency-differential-channel-baseline-normalization)
- Reason-Type Resolution From Severance-Survival States: Typing an Agent's Relationship Edges Without Asking the Counterparty (/articles/earned-exposure/reason-type-resolution-survival-states)
- Reason-Type Retyping by a Superseding-Type Entry: Recomputing an Edge's Reason-Type on Later Severance Evidence Without Disturbing Modifications Already Applied (/articles/earned-exposure/reason-type-retyping-superseding-type-entry)
- Renewal-Side Replenishment of the Authorization Budget: Restoring an Agent's Metered Spending Capacity Only on Exposure to an Uncounted Origin-Equivalence Class (/articles/earned-exposure/renewal-side-replenishment-authorization-budget)
- Following Does Not Count as Being Tested: Retention of the Short Decay Constant When an Agent Defers, and the Action Class That Stays Untested (/articles/earned-exposure/retention-short-decay-constant-deference-class)
- Standing-Determination Designation: Substituting a Verified Counterparty Determination for One Conjunct of the Dispatch-Authority Predicate (/articles/earned-exposure/standing-determination-designation-substitution-inside-dispatch)

- Uncounted-Class Renewal: Gating Agent Authorization on a Windowed Register of Origin-Equivalence Classes (/articles/earned-exposure/uncounted-class-renewal-gated-windowed-register)

APPLICATIONS · GENERAL

- Vendor Qualification Without Ratings: Standing Earned From Recorded Conduct (/articles/earned-exposure/vendor-qualification-without-ratings)
- Credentialing Autonomous Practitioners: Deference That Survives Across Agents (/articles/earned-exposure/professional-credentialing-agents)

APPLICATIONS · SPECIFIC

- Marketplace Seller Reputation Systems: Presented Scores Versus Earned Admissibility (/articles/earned-exposure/marketplace-seller-reputation)
- **Letta: Agent Memory and Standing Earned From Conduct (/articles/earned-exposure/letta-agent-memory)**
- Trustpilot Signals and Conduct-Derived Admissibility (/articles/earned-exposure/trustpilot-signals)
- Copilot Coding Agents: Review Gates and Earned Exposure (/articles/earned-exposure/github-copilot-agent)
- Supplier Risk Scores and Recorded Counterparty Conduct (/articles/earned-exposure/supplier-risk-scoring)
- Verified Profiles and Standing Built From Recorded Conduct (/articles/earned-exposure/linkedin-verification)

TERMINOLOGY

- counterparty identity record (/glossary#counterparty-identity-record)
- deference record (/glossary#deference-record)
- edge (/glossary#edge)
- reason-type (/glossary#reason-type)
- severance event (/glossary#severance-event)
- standing quantity (/glossary#standing-quantity)
- untested class (/glossary#untested-class)

Deference and Earned Exposure overview → (/earned-exposure)