AILOITTE | COMPLIANCE UNBLOCKER

Is compliance slowing down your agent memory rollout?

A practical Q&A and evidence checklist to help Security and Compliance say "approve with controls".

Version 1.1 | October 2026 | ailoitte.com

Why this exists

Most AI agent rollouts do not stall because the model is weak. They stall because memory becomes a governance and audit problem: permission gaps, unclear retention and deletion, and missing traceability. This guide gives you a fast way to answer the questions reviewers ask.

1 Fast triage (60 seconds)

If you answer "no" to any of these, compliance will rightly slow the rollout.

- [ ] Do memory reads respect RBAC and document-level ACLs, enforced server-side rather than through prompt instructions?
- [ ] Can you show an audit trail for memory writes and recalls: who, what, when and why?
- [ ] Do you have retention and TTL rules plus deletion workflows, including DSAR and right-to-erasure where applicable?
- [ ] Do you prevent "store everything" through data minimization and sensitive data redaction or approval?

2 The 6 compliance objections we hear, and the control that answers each

Objection	The control that answers it	What to show
"Where is memory stored and who can access it?"	Memory is tiered (session, task and long-term). Every read and write goes through a policy layer with tenant and RBAC or ABAC checks.	Data flow diagram and access model
"Can the agent recall something the user should not see?"	ACL-aware recall enforced server-side at retrieval time, using source permissions and tenant filters.	ACL enforcement tests and blocked retrieval examples

Objection	The control that answers it	What to show
"Can we delete it? What about DSAR and right-to-erasure?"	Retention, TTL and deletion workflows across every store: structured, vector and evidence.	Deletion runbook and deletion audit events
"Are you storing sensitive data or secrets?"	Minimization, classification and redaction. Sensitive memories require approval or are blocked.	Memory taxonomy and redaction logs
"How do we explain why the agent said this?"	Audit logs plus provenance and citations for retrieval, tracing which memory items were used.	Audit event schema and a sample trace
"How do we monitor abuse and cost?"	Observability and alerts for anomalous recall and spend spikes, with quotas and budgets per team.	Dashboards and alert thresholds

Tip: if your organization uses SharePoint, Google Drive or Confluence, make sure the retrieval layer enforces source permissions server-side at query time, never through instructions in the prompt.

3 Evidence checklist for "approve with controls"

Preparing for a Security or Compliance review? These artifacts shorten the review and cut the back-and-forth.

- [ ] **Architecture and data flow** Memory architecture diagram covering stores, pipelines and policy gates.
- [ ] **Access model** Tenant isolation strategy, RBAC or ABAC integration and the ACL enforcement approach.
- [ ] **Retention and deletion** TTL policy table, DSAR and right-to-erasure workflow, and proof of deletion across stores.
- [ ] **Auditability** Event schema for memory reads and writes, plus a sample trace for one agent response.
- [ ] **Data minimization** Memory taxonomy (allowed, blocked, sensitive) and redaction or approval rules.
- [ ] **Testing** Prompt injection and unauthorized recall tests, plus red-team findings and fixes.
- [ ] **Operations** Cost and usage dashboards, alerting, and an incident response runbook for policy violations or leakage.

4 Six high-signal questions to send us

Pick the one that is blocking you. We reply with the control and the evidence checklist for your environment.

1. How do we enforce ACL-aware recall against SharePoint, Google Drive or Confluence permissions? **Ask us**
2. What does a DSAR or right-to-erasure workflow look like for memory, across structured and vector stores? **Ask us**
3. What exactly do we log for memory reads and writes to satisfy audit? **Ask us**
4. How do we prevent storing PII, PHI or secrets in memory? **Ask us**
5. How do we prove tenant isolation in retrieval and vector search? **Ask us**
6. What is the minimum set of controls to launch safely in 30 to 60 days? **Ask us**

Want a compliance unblock session?

Book a short call and we will turn your blocker into a concrete control and an evidence plan your reviewers can approve.

Book a call: **meetings-na2.hubspot.com/ailoitte/30min** | Email your question: **info@ailoitte.com**

Related free resources

- Enterprise Memory Reference Architecture: the diagram, both pipelines and a control map.
- AI Agent Memory Checklist: score your design in about 20 minutes.
- Memory Risk Controls Case Snippet: these controls applied in a real rollout.