AILOITTE | CASE SNIPPET

Reducing Enterprise Risk in Agent Memory Rollout

How we made AI agent memory permission-aware, auditable and compliant, without slowing delivery.

Version 1.1 | October 2026 | ailoitte.com

Context

An enterprise team was moving its AI agent from pilot to production. It needed memory for continuity and workflow state, and it had to meet security and compliance expectations before rollout.

Risks we saw

- **ACL gaps:** users could recall information they should not have access to.
- **Missing audit trail:** no clear answer to "why did the agent say this?"
- **Undefined retention and deletion:** exposure to DSAR and right-to-erasure requests.

What we implemented: the control map

Control	What it prevents
1 **ACL-aware recall**	Unauthorized recall and data leakage
2 **Tenant isolation**	Cross-tenant contamination
3 **Retention, TTL and delete workflows**	Compliance gaps and storage sprawl
4 **Audit logs (writes and recalls)**	Gaps in non-repudiation and investigations
5 **Sensitive data guardrails**	PII, PHI and secrets persisting in memory
6 **Observability and alerts**	Silent failures, drift and abuse

Outcome

SECURITY AND COMPLIANCE

Review shifted from "block" to "approve with controls".

ENGINEERING

Rollout unblocked with a clear governance model.

OWNERSHIP

Memory became an owned system (policies, logs and lifecycle), not a hidden feature.

Want a fast risk assessment of your agent memory?

Book a short consult with our AI engineers. We will walk through your design and highlight the top risks and controls to close them.

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

Related free resources

- AI Agent Memory Checklist: score your design against the 7 requirements.
- The Enterprise Memory Layer eBook: the full architecture and 90-day roadmap.

Client details are withheld under NDA.