AILOITTE | ENTERPRISE AI AND PRODUCT ENGINEERING

Enterprise Memory Checklist for AI Agents

The 7 non-negotiables for agent memory that is secure, permission-aware, auditable and compliant.

Version 1.1 | October 2026 | ailoitte.com

Purpose

Use this checklist to evaluate whether an AI agent memory approach is enterprise-ready: secure, permission-aware, auditable and compliant. It helps teams avoid the most common failure: shipping "memory" that works in demos but breaks trust and governance in production.

Company		**Industry**	
Use case(s)		**Stage**	☐ Planning ☐ Pilot ☐ Production
Owner		**Date**	

How to score

0 Not defined
No decision, policy or implementation. High risk to proceed.

1 Partially defined
Some decisions exist, but gaps remain. A pilot may be possible with constraints.

2 Defined and implementable
Clear requirements and a feasible plan. Ready for production hardening.

Score every item 0, 1 or 2. A section is only as strong as its weakest item, so each **section score is its lowest item score.** Add the seven section scores for a **total out of 14**, then use the bands on the last page.

What "good" looks like

A production-ready memory system is (1) permission-aware at recall time, (2) governed by retention and deletion, (3) auditable end to end and (4) cost-controlled. The checklist below turns these into checks you can score.

1 Use-case clarity and memory boundaries

Why it matters. Memory should exist to serve a defined outcome (speed, personalization, continuity), not as a generic feature.

Done	Checklist item	Score	Evidence and notes
☐	Top 1 to 3 agent journeys are defined (who uses it, for what decision or workflow).	0 1 2	
☐	Memory scope is defined: what the agent MUST remember vs MUST NOT store, with examples.	0 1 2	
☐	Success criteria are agreed (adoption, resolution time, accuracy, risk reduction).	0 1 2	

Expert prompts for your review

- What breaks if memory is wrong for 24 hours?
- Which fields may persist (preferences, tasks, summaries) and which are prohibited (PII, PHI, secrets)?
- What is the fallback when memory is unavailable or low-confidence?

Evidence to collect

PRD or BRD section, user journey map, memory schema draft, examples of allowed and prohibited memories.

Section score (lowest item score) _____ / 2

2 Identity, RBAC and ownership model

Why it matters. Enterprise memory must obey identity and role boundaries to prevent unauthorized recall and data leakage.

Done	Checklist item	Score	Evidence and notes
☐	SSO and identity provider approach is decided (Okta, Azure AD, Cognito or similar).	0 1 2	
☐	RBAC roles are defined (user, admin, auditor) with ownership of memory objects.	0 1 2	
☐	Admin controls are defined (view, edit, delete, export), with approvals where needed.	0 1 2	

Expert prompts for your review

- Who owns a memory item: the user, the team or the organization?
- Can admins view user memory? Under what policy and audit trail?
- What happens when a user changes teams or permissions?

Evidence to collect

IAM and RBAC matrix, admin roles document, access policy draft, user lifecycle rules.

Section score (lowest item score) _____ / 2

3 ACL-aware recall (the most missed requirement)

Why it matters. Recall must be permission-aware at retrieval time: filtered by ACL, tenant and purpose, and enforced server-side.

Done	Checklist item	Score	Evidence and notes
☐	Recall queries are filtered by tenant, workspace, team and user boundaries.	0 1 2	
☐	Authorization is enforced server-side, not through prompt instructions.	0 1 2	
☐	The test plan includes leakage tests (cross-user, team and tenant) and negative cases.	0 1 2	

Expert prompts for your review

- Can the agent ever recall something a user cannot access in the source system?
- Are permissions evaluated at retrieval time, not only at ingest time?
- How do you prevent prompt injection from bypassing recall rules?

Evidence to collect

Retrieval design, policy enforcement point, test cases, security review notes.

Section score (lowest item score) _____ **/ 2**

4 Retention, TTL, deletion and DSAR readiness

Why it matters. Without retention and deletion controls, memory becomes a compliance and cost liability.

Done	Checklist item	Score	Evidence and notes
☐	Retention and TTL rules are defined per memory type (preferences, work history, summaries).	0 1 2	
☐	User and admin "forget" workflows are defined, including exceptions.	0 1 2	
☐	A DSAR and right-to-erasure process is defined (SLA, audit, verification).	0 1 2	

Expert prompts for your review

- What is the default retention window for each memory type?
- Can you prove deletion across every store (database, vector store, logs, backups)?
- How do you handle legal hold or regulated retention requirements?

Evidence to collect

Retention policy, deletion workflow, DSAR SOP, data map across stores.

Section score (lowest item score) _____ **/ 2**

5 Auditability and observability

Why it matters. You need to answer who stored what, who recalled it and why it influenced an answer, and detect misuse early.

Done	Checklist item	Score	Evidence and notes
☐	Memory writes are logged (actor, timestamp, type, source, purpose).	0 1 2	
☐	Memory recalls are logged (query context, result IDs, decision path).	0 1 2	
☐	Monitoring exists for anomalies (recall spikes, sensitive patterns).	0 1 2	

Expert prompts for your review

- Can you reconstruct why the agent gave a specific answer last week?
- Do you capture provenance (source links or IDs) for each memory item?
- Do you have alerts for abnormal recall behavior or exfiltration attempts?

Evidence to collect

Logging schema, audit dashboards, alert rules, runbook for investigations.

Section score (lowest item score) _____ **/ 2**

6 Safety guardrails and sensitive data controls

Why it matters. Memory can amplify risk: if sensitive information is stored or recalled, it can be repeated at scale.

Done	Checklist item	Score	Evidence and notes
☐	Sensitive data detection is defined (PII, PHI, secrets) with block and redact rules.	0 1 2	
☐	A human override and moderation workflow is defined for risky memories.	0 1 2	
☐	Safe recall and safe output rules are defined (what can be surfaced, and to whom).	0 1 2	

Expert prompts for your review

- What happens when a user pastes a secret into chat? Do you store it?
- Do you maintain allowlists and denylists for memory categories?
- How do you prevent summarization from retaining restricted data?

Evidence to collect

DLP rules, moderation SOP, guardrail policies, examples of blocked and redacted cases.

Section score (lowest item score) _____ **/ 2**

7 Architecture, encryption and cost controls

Why it matters. Production memory needs predictable performance, encryption, isolation and spend control across tenants and workloads.

Done	Checklist item	Score	Evidence and notes
☐	Memory store strategy is defined (structured store, retrieval strategy, vector usage).	0 1 2	
☐	Encryption and key ownership are defined (KMS or CMK, rotation, secrets).	0 1 2	
☐	Quotas, rate limits and storage caps are defined per user, team and tenant, with monitoring.	0 1 2	

Expert prompts for your review

- Where does each memory type live (database, vector store, cache), and why?
- What are the cost drivers (storage growth, embedding churn, recall frequency)?
- How will you enforce quotas and prevent runaway memory growth?

Evidence to collect

Architecture diagram, security controls checklist, quota policy, cost model assumptions.

Section score (lowest item score) _____ **/ 2**

Overall scoring and next step

Copy each section score here and add them up.

1. Use-case clarity and memory boundaries	_____ / 2
2. Identity, RBAC and ownership model	_____ / 2
3. ACL-aware recall (the most missed requirement)	_____ / 2
4. Retention, TTL, deletion and DSAR readiness	_____ / 2
5. Auditability and observability	_____ / 2
6. Safety guardrails and sensitive data controls	_____ / 2
7. Architecture, encryption and cost controls	_____ / 2
Total score	_____ / 14

0 to 6: Red

Do not proceed to production. Define governance, access controls, retention and deletion, and auditability first.

7 to 10: Amber

Pilot-ready with constraints. Run a controlled pilot while closing compliance and audit gaps.

11 to 14: Green

Strong readiness. Proceed to production design and hardening with Security and Compliance sign-off.

Recommended next step

_____

_____

_____

Optional: free expert review (15 to 20 minutes)

Share your completed checklist and our AI engineers will give you (1) your top 3 risks, (2) the recommended controls and (3) a phased plan to reach production safely.

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