

Glossary

Purpose: Align vocabulary for internal teams, prospects, and AI prompts.

Term	Definition
MCP (Model Context Protocol)	Streamable HTTP interface that lets AI tools read/write memories via <code>read_memories</code> , <code>write_memory</code> , <code>browse_memories</code> , and the other verbs.
<code>allowed_vendors</code>	JSON array field that scopes visibility. <code>["*"]</code> = all tools, otherwise the tool must be explicitly listed to read the memory.
shared memory	Team-level pool returned by <code>read_team_memories</code> , shared via <code>share_memory</code> only when the author opts in.
supersession	A memory that has been superseded by a newer version. The graph helpers expose these relationships so agents know what is stale.
session briefing	Structured snapshot of the user's memory state delivered to the model on connect—contains identity, topic clusters, open threads, and behavioral guidance.
deployment modes	<code>hosted</code> , <code>isolated-hosted</code> , <code>self-host</code> with tailored network, egress, and telemetry controls.
question bank	<code>content/diligence/_source/question-bank.yaml</code> – the transcript-derived dataset that feeds the FAQ, docs, and prompt copy boxes.
graph tools	Future MCP helpers (<code>get_open_tickets</code> , <code>get_unresolved_threads</code> , <code>get_recent_decisions</code> , <code>get_superseded_decisions</code> , <code>get_current_state(topic)</code>) that synthesize operational state.
audit trail	Logging of every read/write/auth event stored in <code>usage_events</code> , powered by the same SQLite/Postgres backend.
async diligence	The end-to-end experience (markdown + prompts + PDFs) that allows prospects to deploy their own AI to evaluate Reflect.