

Positioning

Questions this doc answers

- What is the core hook for enterprise buyers?
- What language should we avoid (e.g., "AI memory tool")?
- What competitive story should we tell to differentiate from app-layer players?

Messaging spine

1. **Pain:** "Your organizational AI context shouldn't be living in Slack threads." This resonates with CTOs who cannot afford repeated synchronous explanations.
2. **Consequence:** Organizations lose decisions, customer context, and compliance knowledge every time the tool resets.
3. **Promise:** Reflect is a **persistent organizational memory layer** deployed inside your network—the AI tools you already use (ChatGPT, Claude, Cursor, etc.) read/write the same memory automatically.
4. **Moat:** Private deploy + explicit visibility controls + graph/temporal tooling. Competitors either lock you inside one tool (Claude/ChatGPT) or leave you with an app-layer wrapper (Mem0, Notion).

Session briefings & async diligence

Reflect's core architectural contribution is the **session-initialization briefing**: a structured snapshot of the user's memory state delivered to the model on connect, before any prompt is processed. The briefing contains identity, named topic clusters (via Louvain detection), open threads, tagging conventions, and behavioral guidance—so the AI knows what exists before it writes.

The async diligence bundle ensures your docs meet the same expectation: markdown + PDF downloads, question boxes, and copy-paste prompts plus the question-bank entries derived from transcripts.

What not to say

- Avoid calling Reflect "an AI memory tool" (Mem0/Supermemory semantics).
- Don't lead with integrations; lead with the pain of repeated re-explanations and the cost of synchronous meetings.
- Mention competitive comparisons in the `/diligence/competitive-differentiation` doc—not in the hero.