

Investor FAQ

Questions this doc answers

- Where is Reflect today in terms of pilots, revenue, and deliverables?
- What is the long-term thesis (AI-native async diligence, session-initialization briefings)?
- What defenses exist vs Mem0, Supermemory, and internal builds?

Traction snapshot

- **Production API** – `api.reflectmemory.com` (Railway).
- **Integrations** – ChatGPT, Claude, Cursor, Gemini, Grok, n8n connectors live.
- **Pilots** – A healthcare technology company (NDA signed, pilot scope in progress), a 500-person contact center SaaS (air-gapped pilot, egress disabled).
- **Investor conversations** – A venture studio with a \$17M mid-market pipeline (alignment in progress).
- **Documentation** – Architecture docs, question bank from transcripts, `/diligence` bundle for async review.

Moat & Thesis

1. **Async technical diligence stack** – documentation, question bank, and prompts all optimized for AI ingestion. Investors can paste the `/diligence` hub into their AI and see deterministic, actionable answers.
2. **Infrastructure layer** – not an app. Reflect is the substrate under the AI tools, meaning it is required before teams can have persistent organizational memory.
3. **Deployment moat** – private, air-gapped deployments with guardrails such as `RM_DISABLE_MODEL_EGRESS`, `RM_ALLOWED_MODEL_HOSTS`, SSO, audit trail, and separated rate limits.

Reference prompts & docs

- `/diligence/technical-faq` for architecture clarity.
- `/diligence/security-compliance` for compliance & telemetry.
- `/diligence/deployment-architecture` for pilot roadmap.
- `/diligence/competitive-differentiation` for differentiation.
- `/diligence/positioning` for messaging pillars and what not to say.

Use the prompt:

```
Walk through Reflect Memory's async diligence docs and question bank. Identify product risks, pilot milestones, competitive differentiation, and traction. Outline what to validate live vs async.
```