Pillar #9 Interplay & OAuth

```
## §8 — Pillar #9 Interplay + OAuth (D6 / AC#6)

### Topic 1 — Memory-layer location (laptop vs VM vs hybrid)

**Decision:** Mem0 = laptop-only (ANVIL) for now. Qdrant port 6333 and Ollama port 11434
are both ANVIL-local. vm-alai-support (Pillar #9) does not have direct access to ANVIL
ports.

**Topology gap:** Pillar #9 VM (vm-alai-support, 4.223.110.181) cannot reach ANVIL localhost:900
directly. Mem0 server is bound to 127.0.0.1. Resolution options: (a) CF tunnel rule exposing
Mem0 port via CF Access (preferred — no public binding, CF handles auth); (b) rsync Qdrant
snapshot to VM on a schedule (read-only replica); (c) move Mem0 to vm-alai-support (requires
Qdrant + Ollama on VM — adds ~$10/mo GPU-less Ollama inference cost). Chip-huyen EC-4: Mem0
bound to 127.0.0.1:9000 today (ANVIL-only). CF tunnel option is the lowest-risk path.
This is a Phase 3 decision — surfaces to §11 item #3.

### Topic 2 — OAuth-CLI-on-VM read/write authority boundary

LLM-client construction paths for each framework:

| Framework | LLM client construction | OAuth-compatible |
|-----------|-------------------------|-----------------|
| Mem0 self-hosted | /Users/makinja/system/mem0/config.py lines 67-77: `{"provider":"ollama","co
| claude-mem | /opt/homebrew/lib/node_modules/claude-mem/package.json — no @anthropic-ai/sdk dep
| mem-search | NOT VIABLE — no code path exists | N/A |
| Memipalace | NOT VIABLE — no code path exists | N/A |
| LightRAG-resurrect | /health response: `llm_binding_host:https://ollama.basicconsulting.no` —

EVIDENCE: config.py lines 67-77 (file confirmed on disk); claude-mem package.json; LightRAG /hea

All three viable frameworks are COMPATIBLE WITH PILLAR #9 OAuth model (no Anthropic API key requ

### Topic 3 — State-sync timing (rsync windows)

Qdrant data dir: /Users/makinja/.qdrant/storage (ANVIL local, not yet confirmed path).
If Mem0 is moved to VM: rsync window recommendation = every 4h during active sessions
(per Pillar #9 spec §3.3 state-sync design). For the current laptop-only topology, no
rsync needed — Mem0 is single-source-of-truth on ANVIL.

### Topic 4 — Multi-client SVE namespace isolation

Current state: `user_id='john'` hardcoded in discover.js line 677.
Qdrant payload_schema shows user_id as keyword field — Qdrant already supports per-user
filtering natively.

Two designs:
- **Design A (recommended): metadata filter** — single mem0_john collection, query with
  `payload filter user_id=`. Cost: zero additional infra. Risk: one corrupt
```

```
  write with wrong user_id bleeds facts. Mitigation: server.py write endpoint validates
  user_id against allow-list.
- **Design B: per-client collection** — `mem0_john`, `mem0_snowit`, `mem0_adnancesko`, etc.
  Clean isolation, harder to cross-search. Config change per client in config.py.

Recommendation: Design A for Phase 3 (lower ops overhead). Design B if client-count
exceeds 10 or audit trail is required. Surfaces to §11 item #2.

### Topic 5 — DR access path

If ANVIL (MacBook) goes offline:
- Mem0 data: no off-laptop copy today. Qdrant snapshots must be added to the rsync-to-VM
  step (Step 1 of migration plan above).
- LightRAG backups at /Users/makinja/system/backups/lightrag/20260503-040002/ — 4 tarballs
  with MANIFEST.sha256.
- Pillar #9 VM already has CF tunnel access; CEO Telegram bridge handles text dispatch.
- RTO for memory-only recovery: 1h if Qdrant snapshot is available on VM; 4h cold (restore
  from backup).

---
```

Revision #2
Created 2026-05-07 10:24:45 UTC by John
Updated 2026-06-07 20:01:28 UTC by John