

# LightRAG Stabilization Runbook

## — 2026-05-08

## Genesis

On 2026-05-08 at 14:00, Kelsey Hightower reported LightRAG returning 502 errors. By 19:05 the service had degraded to complete timeout (000). Root cause: MainThread synchronous list comprehension in `lightrag/lightrag.py:872` (`pipeline_process_enqueue_documents`) iterating over 125,341-record `JsonDocStatusStorage` on the asyncio event loop. A single POST to `/documents/text` triggered full file rewrite + pipeline iteration over 121K pending docs → CPU pegged at 100% → `/health` unreachable. The issue was compounded by running `sbnb/lightrag:latest` amd64-only image under Rosetta on Apple Silicon, incurring 2-3× performance tax.

## Six-Step Fix Applied

### S1: Disable Runaway Ingest Agents

Stopped LaunchAgents: `com.alai.lightrag-outbox-ingest`, `com.alai.lightrag-migrate-pump`, `com.alai.lightrag-watchdog`. Kept: `keepwarm`, `backup`, `monitor`.

### S2: Prune Pending Queue

Stopped container, backed up `doc_status.json`. Filtered to `status=processed` only: 8,357 records retained; 116,986 pending/processing/failed quarantined to backup. Restarted container. CPU dropped to 0.31%.

### S3: Verify Queryability

Tested naive mode + `only_need_context=true` (bypasses LLM, returns ALAI corpus chunks). Graph/label endpoint returned 200+ entities. Service functionally restored.

### S4: Image Swap for Native ARM64

Replaced `sbnb/lightrag:latest` (amd64, v1.3.4) with `ghcr.io/hkuds/lightrag:latest` (native arm64, v1.4.16, official upstream). Verified via `docker manifest inspect`.

## S5: Resource Limits

Added cgroup-enforced limits in compose: `cpus: 2.0`, `memory: 4G`.

## S6: Re-Ingest Worker Design

Designed (not implemented) re-ingest worker with: `batch_size=10`, `cooldown=60s`, health-gate, pre-flight LLM availability check, cursor-based restart safety. Build gated on CEO OCD-3 (aging policy decision).

## Verified Post-State

- Container: Up, healthy
- CPU: 4.29%
- Memory: 1.24/4 GiB
- `/health`: HTTP 200 in 3.7ms
- Naive query: Returns ALAI documentation chunks
- Knowledge graph: Queryable

## Known Follow-Ups

- **Child MC #100027 (M, FlowForge):** `LLM_MODEL=qwen3:8b-q8_0` in `.env` not on Ollama; hybrid/local/global modes return 404. Fix: one-line `.env` change to `llama3.1:8b` OR `ollama pull qwen3:8b-q8_0`.
- **CEO OCD-3 aging policy:** Decision required before re-ingest build. 100% of 117K backlog has `unknown_source`. Option A (drop-all) recommended due to zero provenance utility.
- **Martin's asyncio event-loop freeze:** HKUDS 1.4.16 may have fixed root cause. Verify before relaxing `batch_size` in re-ingest worker.

## Evidence Files (Local, Transient)

- `/tmp/lightrag-stabilization-step1-evidence.txt` — `launchctl` list before/after
- `/tmp/lightrag-stabilization-step2-evidence.txt` — `doc_status` counts, `/health` timing post-restart

- `/tmp/lightrag-stabilization-step3-evidence.txt` — query mode results, graph endpoint outputs
- `/tmp/lightrag-stabilization-step4-evidence.txt` — manifest inspect, docker inspect
- `/tmp/lightrag-stabilization-step5-evidence.txt` — docker stats with limits, compose diff
- `/tmp/lightrag-stabilization-step6-evidence.txt` — re-ingest worker design doc
- `/tmp/lightrag-stabilization-progress.txt` — step-by-step progress log
- `/tmp/cache-proxy-99981/` — pre-existing Phase 1 hash-cache proxy evidence (separate)

## References

- Forge file: `/Users/makinja/system/prompts/forged/99982.md` (panel decisions used to drive stabilization)
  - Mehanik clearance: `/tmp/mehanik-cleared-100009` (2026-05-08 19:56)
  - Genesis SENTINEL v3 audit: `project_sentinel_v3_audit_2026-05-01.md`
  - Image diff: `sbnb/lightrag:latest @ 1.3.4` → `ghcr.io/hkuds/lightrag:latest @ 1.4.16`
- 

Revision #2

Created 2026-05-08 18:25:40 UTC by John

Updated 2026-06-14 20:02:39 UTC by John