

Runbook: LightRAG ingest LaunchAgent fix (MC #10286)

Overview

This runbook documents the investigation and fix applied to three LightRAG-related LaunchAgents on the ALAI Mac Studio host in MC #10286. The fix was validated by Proveo (Angie Jones) with a PARTIAL verdict: 3 PASS, 1 PARTIAL (AC3), 1 FAIL (AC4 — same-day unverifiable). CF Access root cause is tracked separately in MC #10298.

1. Symptom — How to Detect This Failure

These signals indicate the `com.alai.lightrag-outbox-ingest` LaunchAgent is failing silently:

- **Outbox file grows, doc count does not:** `wc -l ~/system/logs/mc-task-outcomes.jsonl` increases after each `mc.js done`, but `curl http://localhost:9621/documents | jq .total` stays flat over days.
 - **SQLite checkpoint stops advancing:** `sqlite3 ~/system/state/outbox-ingest.sqlite "SELECT MAX(ingested_at) FROM processed"` returns a timestamp from days ago.
 - **Watchdog calendar_err alert:** Daemon-fleet-watchdog fires a `calendar_err_N` alert for `com.alai.lightrag-outbox-ingest` or `com.john.lightrag-monitor`.
 - **HTTP 302 in error log:** `tail ~/system/logs/lightrag-outbox-ingest.err` shows 302 or redirect errors when posting to `https://lightrag.alai.no/documents/text`.
 - **PID column is "-" with non-zero LastExitStatus:** `launchctl list | grep lightrag` shows `PID="-"` with non-zero `LastExitStatus` for a timer-scheduled daemon (`StartInterval`) is abnormal; for calendar daemons it is normal between scheduled windows.
-

2. Root Cause

The primary failure was in `com.alai.lightrag-outbox-ingest`:

- The plist `LIGHTRAG_URL` environment variable was set to `https://lightrag.alai.no` (the public Cloudflare-proxied URL).
- CF Access service token was returning HTTP 302 on `POST /documents/text` requests from the local host, causing all upload attempts to time out or silently fail.
- LightRAG itself was healthy at `http://localhost:9621` — this is the correct direct URL for host-local callers.

Workaround applied: Changed `LIGHTRAG_URL` to `http://localhost:9621` in the plist. The CF Access token 302 root cause (why the local host receives a redirect instead of being authorized) is tracked in **MC #10298** (priority: M).

The other two daemons were not functionally broken:

- `com.alai.lightrag-backup`: Calendar Sunday-only schedule. `PID="-"` between fires is `launchd-normal`. `LastExitStatus=0`. No defect.
- `com.john.lightrag-monitor`: `exit 256 = bash exit 1 = warnings-only state` (Ollama route 302, SSH not configured). These are pre-existing infrastructure gaps, not failures. The script exits 1 to flag warnings; this is by design.

3. Fix Procedure

Preconditions: You have shell access to the Mac Studio host. LightRAG is running locally on port 9621.

Step 1: Verify current plist URL

```
grep -A1 "LIGHTRAG_URL" ~/Library/LaunchAgents/com.alai.lightrag-outbox-ingest.plist
```

If the value is `https://lightrag.alai.no`, proceed. If already `http://localhost:9621`, skip to Step 4.

Step 2: Edit the plist

```
nano ~/Library/LaunchAgents/com.alai.lightrag-outbox-ingest.plist
```

Change the `LIGHTRAG_URL` string value from `https://lightrag.alai.no` to `http://localhost:9621`.
The correct plist line:

```
<key>LIGHTRAG_URL</key><string>http://localhost:9621</string>
```

Step 3: Unload all 3 lightrag plists

```
launchctl unload ~/Library/LaunchAgents/com.alai.lightrag-outbox-ingest.plist
launchctl unload ~/Library/LaunchAgents/com.alai.lightrag-backup.plist
launchctl unload ~/Library/LaunchAgents/com.john.lightrag-monitor.plist
```

Step 4: Reload all 3 lightrag plists

```
launchctl load -w ~/Library/LaunchAgents/com.alai.lightrag-outbox-ingest.plist
launchctl load -w ~/Library/LaunchAgents/com.alai.lightrag-backup.plist
launchctl load -w ~/Library/LaunchAgents/com.john.lightrag-monitor.plist
```

Step 5: Drain the outbox manually (if backlog exists)

```
node ~/system/tools/lightrag-outbox-ingest.js
```

The script is idempotent — it uses `outbox-ingest.sqlite` with `correlation_id` as PRIMARY KEY dedup gate. Running it multiple times is safe. Expected output when backlog is cleared: `processed: 0, skipped: N, failed: 0`.

Step 6: Kickstart the ingest daemon to verify immediate fire

```
launchctl kickstart -k gui/$(id -u)/com.alai.lightrag-outbox-ingest
```

Check the log immediately after:

```
tail -20 ~/system/logs/lightrag-outbox-ingest.log
```

Expected: A `[ingest] DONE` line with exit success.

Step 7: Confirm watchdog detects healthy state

```
bash ~/bin/daemon-fleet-watchdog.sh 2>&1 | grep lightrag
```

Expected: All 3 labels in `calendar_ok` state. No `calendar_err_*` or `not_loaded` transitions.

4. Verification Commands

```
# 1. All 3 plists loaded with LastExitStatus=0
launchctl list | grep lightrag

# 2. Checkpoint DB row count (should match mc-task-outcomes.jsonl line count)
sqlite3 ~/system/state/outbox-ingest.sqlite "SELECT count(*) FROM processed"

# 3. Most recent ingest timestamp
sqlite3 ~/system/state/outbox-ingest.sqlite "SELECT MAX(ingested_at) FROM processed"

# 4. LightRAG pipeline health
curl http://localhost:9621/documents/pipeline_status

# 5. LightRAG document total count
curl http://localhost:9621/documents | jq .total

# 6. Outbox log last run summary
grep "DONE" ~/system/logs/lightrag-outbox-ingest.log | tail -5

# 7. Watchdog recent transitions for lightrag
grep lightrag ~/system/logs/daemon-fleet-watchdog.log | tail -20
```

5. Known Limitations

- **AC4 cannot be verified same-day:** `com.alai.lightrag-outbox-ingest` fires on `StartInterval=21600` (6 hours). Verifying that `launchd` autonomously fires the next scheduled cycle requires waiting at least 6 hours after the kickstart. Same-day verification only demonstrates manual-kickstart success. Rely on the `daemon-fleet-watchdog` for ongoing health monitoring.
- **Log timestamps absent:** `lightrag-outbox-ingest.js` does not emit timestamps to its log file. This makes it impossible to distinguish manually-triggered runs from `launchd`-autonomous fires in the log tail. Consider adding a timestamp at script start as a follow-up TD.
- **CF Access 302 root cause unresolved:** The public URL `https://lightrag.alai.no` still returns HTTP 302 for host-local service token requests. The localhost bypass is a workaround. If the CF tunnel configuration changes or `localhost:9621` changes port, the plist must be updated again. See MC #10298 for the proper fix.

- **com.john.lightrag-monitor DRAFT comment:** The plist still contains a stale "DRAFT — pending Alem approval" comment referencing MC #8545. The daemon IS installed and running. This comment is cosmetic noise but should be cleaned up.
- **AC3 drain was incremental, not single-session:** The 312-entry outbox was drained incrementally across multiple sessions starting 2026-04-17. Any future outbox drain may similarly require multiple passes if entries arrive between runs.

6. Watchdog Coverage

The daemon-fleet-watchdog at `~/bin/daemon-fleet-watchdog.sh` covers all 3 LightRAG plists via its glob at line 39:

```
for plist in "$HOME"/Library/LaunchAgents/com.{alai,john}.*.plist
```

This glob automatically includes any new LightRAG LaunchAgents matching the pattern without code changes. The watchdog runs every 15 minutes via `com.alai.daemon-fleet-watchdog`.

Alert states to watch for:

- `calendar_err_256` — daemon exits with code 1 (warnings/errors)
- `calendar_err_512` — daemon exits with code 2 (script error)
- `not_loaded` — plist unloaded from launchd (critical)

Healthy state: `calendar_ok` (LastExitStatus=0, plist loaded)

7. Related MCs

MC	Title	Status	Notes
#10286	Fix LightRAG ingest LaunchAgents — drain 312 outbox + add watchdog	DONE (PARTIAL verify)	This fix. Delivered by Kelsey Hightower. Proveo: 3 PASS, 1 PARTIAL, 1 FAIL.
#10298	CF Access service token 302 root cause investigation	OPEN (priority: M)	Why does https://lightrag.alai.no return 302 for local host? Resolves the need for the localhost bypass.

8. Evidence Links

- Proveo full report: `/tmp/postflight-10286/proveo-report.md`
 - Proveo JSON: `/tmp/proveo-10286-1777555315.json`
 - Watchdog glob source: `~/bin/daemon-fleet-watchdog.sh:39`
 - Plist (fixed): `~/Library/LaunchAgents/com.alai.lightrag-outbox-ingest.plist` —
LIGHTRAG_URL=http://localhost:9621
 - Checkpoint DB: `~/system/state/outbox-ingest.sqlite` — 312 rows as of 2026-04-30
 - Ingest log: `~/system/logs/lightrag-outbox-ingest.log` — 6286 lines, multi-session history since 2026-04-17
 - Watchdog log transitions: `~/system/logs/daemon-fleet-watchdog.log` — 12:33:44Z
calendar_ok to not_loaded, 12:44:21Z not_loaded to calendar_ok
-

Revision #2

Created 2026-04-30 13:32:50 UTC by John

Updated 2026-05-31 20:07:03 UTC by John