

Overview & Outcomes

ANVIL Filesystem Sweep — Overview & Outcomes

What and Why

The ANVIL Filesystem Sweep (MC #99637) was a comprehensive cleanup and canonicalization effort to resolve structural drift in ALAI's Mac Studio orchestration host filesystem. After 4+ months of agent activity without central governance, the system accumulated 15 split-brain directory names (same name under both `~/system/` and `~/ALAI/`), 158 daemon references (30 broken), and no canonical path registry. This sweep established tree ownership rules, resolved all split-brain conflicts, reclaimed storage, and documented the canonical map for future agent compliance.

Headline Numbers

- **62,838 paths inventoried** across `/Users/makinja/`
- **158 daemons audited** via LaunchAgent plist examination
- **3 PHANTOM daemons unloaded** (KeepAlive=true boot failures, BLOCKER resolved)
- **~9GB storage reclaimed** after archiving deprecated tools and split-brain duplicates
- **30 valid tar archives** created in `~/backups/anvil-fs-sweep-2026-05-07/` (36 on-disk including W1C additions)
- **6 split-brain pairs resolved** via merge/rename/migrate strategy
- **9 ALAI-canonical dirs established** (clients, infrastructure, legal, products, sales, etc.)
- **4 surprise-canonical paths protected** (live-referenced by tools: aisystem, system/security, system/schemas, system/hooks)
- **5 CEO-excluded items deferred** (drafts, company-prompts, exo-env, hook-native, minions.db)

Phase Chain

Phase	MC ID	Description	Status
-------	-------	-------------	--------

Phase 1.1	#99644	Inventory + heatmap	ready_for_review
Phase 1.2	#99639	Daemon-path graph	ready_for_review
Phase 1.3	#99642	Doc-disk reference audit	ready_for_review
Phase 1.4	#99646	Canonical map synthesis	ready_for_review
Phase 1.5	#99648	Verification	ready_for_review
Phase 1.6	#99662	Content-aware re-classification	ready_for_review
Phase 2	#99655	Gap report (10 CEO decisions)	ready_for_review
Phase 3 W1-A	#99669	BLOCKER: unload 3 PHANTOM daemons	ready_for_review
Phase 3 W1-B	#99672	Bulk safe cleanup	ready_for_review
Phase 3 W1-C	#99695	Split-brain merges (6 pairs)	ready_for_review
Phase 3 W1-C+	#99699	Final 3 split-brain pairs	ready_for_review
Phase 3 W2	#99701	Documentation (registry, ADR, drift spec)	ready_for_review
Phase 3 W3-A	#99703	E2E validation (Proveo)	ready_for_review
Phase 3 W3-B	TBD	BookStack publish (this page)	in_progress

Status

COMPLETE — Proveo E2E validation returned **PARTIAL (9/10 PASS)**. All critical probes passed:

- Inventory diff: PASS (all intended deletions done, ALAI/clients=13 subdirs, new ALAI subdirs present)
- Daemon health: PASS (53 daemons, 3 phantoms absent, running daemons healthy)
- Boot health: PASS (exit 0, no FATAL errors)
- Discover.js verify: PASS (all 9 checks OK)
- Canonical paths preservation: PASS (4 surprise-canonical + 5 CEO-excluded all intact)
- Active reference integrity: PASS (security refs, mehanik-marker, hooks, DEPLOY-MAP valid)
- ALAI/CLAUDE.md: PASS (0 broken refs, pointer present)
- Wave 2 documentation: PASS (3 files with exact line counts)
- Archive integrity: **PARTIAL** (36 tars vs 30 expected, explainable via W1C additions; no corruption)

- MC chain integrity: PASS (all 16 MCs verified)

Open Items

- **Organizational audit** (semantic-fit review): Deferred to separate workstream. Examples: `~/ALAI/web-worktrees/ucenje-v2` (personal scholarly project under commercial brand tree?), `~/projects/` vs `~/companies/` placement criteria.
- **5 CEO-excluded paths**: Pending separate org decision (drafts, company-prompts, exo-env, hook-native, minions.db archived but not deleted).
- **3 technical debt MCs**: #99665 (mlx-router fix), #99666 (db-ttl-sweep), #99667 (distillation-scorer) — open, not yet actioned.

Related Pages

- [Canonical Registry](#) — Authoritative path ownership table
- [ADR-022](#) — Architectural decision record
- [Drift Detection Design](#) — Future build spec
- [Validation Evidence](#) — Proveo 10-probe E2E report

Revision #2

Created 2026-05-07 15:25:57 UTC by John

Updated 2026-06-07 20:01:33 UTC by John