

ADR-022: Architectural Decision

ADR-022: ANVIL Filesystem Canonical Map and Cleanup

Status: Accepted (2026-05-07)

Deciders: Alem Basic (CEO) via Phase 2 gap report

Consulted: Petter Graff (architect), FlowForge (devops), Proveo (validator)

Date: 2026-05-07

Context

ALAI's ANVIL (Mac Studio orchestration host) filesystem had accumulated significant structural drift since late 2025.

Phase 1 Findings (Read-Only Audit)

Scope: 5 read-only audit tasks (MCs #99644, #99639, #99642, #99646, #99648)

Inventory:

- 62,838 total paths under `/Users/makinja/`
- 158 LaunchAgent daemons (managed via `~/Library/LaunchAgents/`)
- 91 child directories under `~/system/`
- 31 child directories under `~/ALAI/`

Split-Brain Detection:

- 15 directory names existed under **both** `~/system/` and `~/ALAI/` simultaneously
- Examples: `agents/`, `architecture/`, `clients/`, `docs/`, `services/`, `templates/`, `infrastructure/`, `legal/`, `products/`, `sales/`, `web/`, `org/`, `pipeline/`, `processes/`, `internal/`

Broken References:

- 137 cited-but-missing documentation references (BookStack pages, specs, runbooks referenced in code/configs but not found on disk or in BookStack)
- 30 broken daemon references (LaunchAgents pointing to non-existent scripts or logs)

PHANTOM Daemons:

- 3 daemons marked KeepAlive=true but failing on every restart cycle:
 - `mlx-router` (boot failure — IP binding issue)
 - 2 others flagged in Phase 1

Phase 1.6 Content-Peek

Extended audit (MC #99662) to **read file contents** in suspicious dirs, revealing:

- 4 surprise-canonical paths (`~/aisystem/`, `~/system/security/`, `~/system/schemas/`, `~/system/hooks/`) — live code reads from these, cannot delete
- Merge conflict detection in split-brain pairs

Phase 2 Gap Report

Analysis (MC #99655) synthesized 10 CEO decision items:

1. **BLOCKER:** `com.john.mlx-router.plist` KeepAlive=true causing boot hang every reboot
2. Tree ownership table (which tree is canonical for what purpose)
3. Split-brain resolution strategy (6 pairs required case-by-case decision)
4. Archive strategy (7-day hot backup vs 30-day cold archive)
5. Broken reference cleanup (delete stale daemon plists, update docs)
6. Surprise-canonical upgrade (protect 4 paths from deletion)
7. Storage reclamation estimate (~9GB after cleanup)
8. CLAUDE.md scope-comment enforcement (4-way scope: global, home, system, ALAI)
9. Documentation deliverables (canonical registry, ADR, drift detection design)
10. Drift detection daemon (weekly monitoring to prevent re-introduction)

Root Causes:

- No central canonical path registry → agents recreated dirs arbitrarily
 - No tree ownership model → business docs under `~/system/`, runtime configs under `~/ALAI/`
 - No drift detection → split-brain accumulated over 4+ months undetected
-

Decision

Execute **ANVIL Filesystem Sweep** in 3 phases, with 10 sub-decisions consolidated under this ADR.

Sub-Decision 1: Resolve BLOCKER

Unload `com.john.mlx-router.plist` (KeepAlive=true causing boot hang). Defer router fix to separate MC. Boot must succeed before bulk cleanup.

Implementation: MC #99669 + 3 technical debt MCs (#99665 router fix, #99666 IP binding root cause, #99667 KeepAlive policy review)

Sub-Decision 2: Establish Tree Ownership Model

Define canonical trees with clear ownership (see [Canonical Registry](#) for full table).

Sub-Decision 3: Split-Brain Resolution (6 Pairs)

Both canonical, distinct purpose (rename non-conflicting):

- `architecture` → System: ADRs, technical decisions. ALAI: renamed to `product-architecture/` (product specs)
- `services` → System: daemon definitions. ALAI: renamed to `service-catalog/` (client-facing offerings)
- `templates` → System: code scaffolding. ALAI: renamed to `doc-templates/` (business docs)

One canonical, migrate other (9 ALAI-wins + 3 system-wins):

- ALAI canonical: `clients/`, `infrastructure/`, `internal/`, `legal/`, `org/`, `pipeline/`, `processes/`, `products/`, `sales/`, `web/`
- System canonical: `agents/`, `docs/`

Implementation: MCs #99672 (bulk cleanup), #99695 + #99699 (split-brain merges/renames)

Sub-Decision 4: Archive Strategy

- **Hot backup:** 7-day retention at `~/backups/anvil-fs-sweep-2026-05-07/`
- **Cold archive:** 30-day tar.gz offsite (B2 bucket via existing daemon)
- **What to archive:** All deleted paths, all split-brain losing sides, all deprecated tools

Sub-Decision 5: Broken Reference Cleanup

- Delete 30 broken daemon plists (scripts no longer exist)
- Mark 137 missing doc refs as TODO in code comments (BookStack publish or delete comment)
- Update `specialist-mapping.json` if any agent dirs deleted

Implementation: Part of MC #99672 bulk cleanup

Sub-Decision 6: Upgrade 4 Surprise-Canonical Paths

Discovered via content-peek (MC #99662). See [Canonical Registry](#) for details on:

- `~/aisystem/` — Mehanik infra gate, CF deployments
- `~/system/security/` — `password-share.js`, `client-vault.js`
- `~/system/schemas/` — `mehanik-commit.js` reads JSON schemas
- `~/system/hooks/` — Git pre-commit/pre-publish (NOT Claude Code hooks)

Status: Protected from deletion. Added to canonical registry.

Sub-Decision 7: Storage Reclamation

Estimate: ~9GB freed after archiving:

- Deprecated tools: ~2GB
- Duplicate split-brain content: ~4GB
- Stale logs/caches: ~3GB

Sub-Decision 8: CLAUDE.md Scope Enforcement

4-way scope model (see [Canonical Registry](#) for full table).

Sub-Decision 9: Documentation Deliverables

Create 3 artifacts (this ADR is one of them):

1. **Canonical Registry** ([see page](#)) — ITIL CMDB / Backstage-style catalog
2. **This ADR** — Architectural decision record

3. **Drift Detection Design** ([see page](#)) — spec for future daemon build

Sub-Decision 10: Drift Detection Daemon (Future Build)

Design deliverable: See [Drift Detection Design](#)

Purpose: Weekly check to prevent split-brain re-introduction

Status: Design done, build deferred to separate MC (not part of this sweep)

Consequences

Positive

1. **BLOCKER resolved** — ANVIL boots without hang (mlx-router unloaded)
2. **~9GB storage reclaimed** — Deprecated/duplicate content archived
3. **Canonical registry established** — Future agents have authoritative "where does X belong" source
4. **Tree ownership clarified** — Orchestration runtime (system) vs company state (ALAI) now explicit
5. **Split-brain eliminated** — 15 dir name conflicts resolved via merge/rename/migrate
6. **Broken references cleaned** — 30 phantom daemons deleted, 137 missing docs flagged
7. **4 surprise-canonical paths protected** — Prevents accidental deletion of live-referenced dirs
8. **CLAUDE.md scope model documented** — 4-way load context now explicit
9. **ADR record created** — Future context for "why is the tree structured this way"
10. **Drift detection designed** — Automated prevention of chaos re-introduction

Negative

1. **Some content judgment deferred** — 3 split-brain pairs (architecture, services, templates) required CEO call for "both canonical" edge case
2. **Organizational audit deferred** — This sweep is **mechanical** (what/where), not **semantic** (should it be there). Example: `~/ALAI/web-worktrees/ucenje-v2` (personal scholarly project under commercial brand tree?) — deferred to separate workstream
3. **3 TD MCs created** — mlx-router fix, IP binding root cause, KeepAlive policy review — technical debt carried forward

4. **137 missing doc refs** — Flagged as TODOs, not resolved (requires BookStack authoring or code comment deletion)
5. **Daemon fleet audit incomplete** — 158 daemons inventoried, 30 broken refs deleted, but full health audit (success rate, error patterns) deferred

Neutral

1. **4 surprise-canonical paths upgraded** — Content-peek revealed live references; upgraded to protected status (good outcome, but unplanned scope expansion)
 2. **5 scope-creep items excluded** — `~/system/drafts/`, `company-prompts/`, `exo-env/`, `hook-native/`, `minions.db` — flagged in gap report but NOT touched in this sweep. Pending separate org audit (semantic-fit review)
-

Implementation

Phase 1: Read-Only Audit (5 Tasks)

Completed MCs:

- #99644: Inventory `~/system/` children (91 dirs)
- #99639: Inventory `~/ALAI/` children (31 dirs)
- #99642: Detect split-brain dir names (15 pairs)
- #99646: Audit daemon fleet (158 LaunchAgents, 30 broken refs)
- #99648: Catalog cited-but-missing docs (137 refs)

Phase 1.6: Content-Peek (Extension)

Completed MC: #99662

Phase 2: Gap Report

Completed MC: #99655 — Synthesize findings into 10 CEO decision items

Phase 3 Wave 1-A: BLOCKER Unload

Completed MC: #99669 — Unload `com.john.mlx-router.plist` `KeepAlive=true`

Technical Debt MCs Created: #99665, #99666, #99667

Phase 3 Wave 1-B: Bulk Cleanup

Completed MC: #99672 — Archive deprecated tools, delete broken daemon plists, reclaim ~9GB

Phase 3 Wave 1-C: Split-Brain Resolution

Completed MCs: #99695, #99699 — Merge/rename/migrate 6 split-brain pairs

Phase 3 Wave 2: Documentation

This Deliverable (MC #99637 child):

- Canonical registry: [see page](#)
- This ADR
- Drift detection design: [see page](#)

Phase 3 Wave 3: Validation & Publication

MC #99703: E2E validation (Proveo) — [see evidence](#)

Current: BookStack publication (MC TBD)

References

Authoritative Documents

- **Canonical Registry:** [Canonical Registry page](#) (also `~/system/specs/canonical-registry.md`)
- **Drift Detection Design:** [Drift Detection Design page](#) (also `~/system/specs/anvil-fs-drift-detection-design.md`)

Mission Control Tasks

Parent: MC #99637 (ANVIL-FS Sweep — PARENT)

Phase 1-3 MCs: See Implementation section above for complete list of 16 MCs.

Related ADRs

- ADR-012: AWS App Runner canonical for Drop (anti-phantom drift reference)
- ADR-021: Bilko blueprint-aligned cleanup (parallel effort, same timeframe)

Context Documents

- `~/ .claude/CLAUDE.md` — John's identity, ZAKONs
 - `~/CLAUDE.md` — Home orchestration guardrails
 - `~/system/CLAUDE.md` — System-specific rules
 - `~/ALAI/CLAUDE.md` — Brand guidelines
-

Revision #2

Created 2026-05-07 15:29:49 UTC by John

Updated 2026-06-07 20:01:35 UTC by John