

Pillar #4 — Skills Audit Overview

Pillar #4 — Skills Audit Overview

Date: 2026-05-04

MC: #99131

Phase: DESIGN + PoC (Phase 2)

Spec: `agentic-os-pillar4-skills-audit-2026-05-04.md`

Executive Summary

This audit covers the ALAI skill system progressive-disclosure refactor: 79 skills inventoried, top-20 refactor priorities identified, L0-L3 rubric established, and PoC analysis completed for `task-postflight` skill.

Key Findings

- 79 active skill directories on disk; 94 rows in skill-registry.db (32 phantoms, 17 unregistered)
 - Only 15 skills have any log invocations in the 19-day measurement window
 - `mehaniK` (186 hits) and `update-config` (1 hit) appear in logs but have no disk directory — ghost invocations
 - 9 skills with references/ dir; 70 are monolithic (L0/L1)
 - 12 TOB skills have nested structure — invisible to Claude Code flat-discovery loader
 - Highest-priority refactor target: `task-postflight` (5,367 tokens × 21 measured invocations = priority_score 82.05)
 - **Reality anchor:** At current ALAI scale (Claude Max flat-rate subscription), context-bloat incremental cost is approximately \$0-2/month. The value of this audit is **context window capacity management**, not dollar cost reduction.
-

Inventory Summary

Metric	Value	Source
--------	-------	--------

Active skill dirs on disk	79	<code>ls ~/.claude/skills/ grep -v _archived wc -l</code>
Archived skills	32	<code>ls ~/.claude/skills/_archived/ wc -l</code>
skill-registry.db rows	94	<code>sqlite3 skill-registry.db 'SELECT COUNT(*) FROM skills;'</code>
DB-only phantoms	32	comm comparison
Disk-only unregistered	17	comm comparison
Skills with references/ dir	9	find query
Skills with invocations in window	15	log grep
Measurement window	19 days	2026-04-16 to 2026-05-05
Total invocations in window	267	awk filter
Ghost invocations (mehanik)	186	log grep — mehanik not on disk

Aggregate Savings Projection

Skills loaded per turn	Tokens saved vs. baseline	% context window recovered (128K window)
Only task-postflight (PASS path)	3,500 tokens	2.7%
task-postflight + prompt-forge	4,700 tokens	3.7%
Top-5 hot-path skills (ranks 1-5)	7,300 tokens	5.7%
All top-20 (max benefit, full session)	19,500 tokens	15.2%
All 79 skills at L3 (theoretical max)	~35,000 tokens	27.3%

Assumes 40-50% body-token reduction per skill post-refactor. These are per-turn estimates derived from body-token reduction; monthly projections without measured session counts would be phantom claims.

Related Documentation

- **Skills Inventory:** Top 20 Priority table with per-skill triage
- **L0-L3 Design Pattern:** Progressive-disclosure rubric and anti-pattern catalog
- **PoC Analysis:** task-postflight refactor demonstration (541 → 194 LOC, 64.7% reduction)

Source spec: `~/system/specs/agent-ic-os-pillar4-skills-audit-2026-05-04.md` (479 lines)

HiveMind record: TBD

MC: #99131 (ready_for_review)

Revision #2

Created 2026-05-05 14:47:36 UTC by John

Updated 2026-06-07 20:01:15 UTC by John