

Audit Summary

Pillar #4 Skills Audit — Summary

Source: `~/system/specs/agent-ic-os-pillar4-skills-audit-2026-05-04.md` (§1-§3 + Reality Anchor)

MC: [#99131](#) | [#99176](#)

Date: 2026-05-05

Phase: DESIGN + PoC (Phase 2)

Executive Summary

This audit covers D2 (Top-20 Refactor Priority Table), D3 (Progressive Disclosure Design Pattern), and D4 PoC analysis for the task-postflight skill refactor.

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.
-

Environment Constants

```
WORKING_DIR=/Users/makinja
SKILLS_ROOT=/Users/makinja/.claude/skills
TELEMETRY_LOG=/Users/makinja/system/logs/skill-use.log
```

```
REGISTRY_DB=/Users/makinja/system/databases/skill-registry.db
TOKENIZATION_FORMULA=bytes_div_3.7
PRICE_USD_PER_MTOK_INPUT=3.00
SESSIONS_PER_MONTH_BASELINE=600
TELEMETRY_WINDOW_DAYS=19
OUTPUT_DIR=/tmp/pillar4-99131-out/
```

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, no disk)	186	log grep — mehanik not on disk

Appendix A — Inventory CSV (83 rows)

Source: `~/system/specs/agent-ic-os-pillar4-skills-inventory.csv`

83 total lines: 3 comments + 1 header + 79 data rows, 19 columns, RFC-4180 compliant

Methodology:

- Tokenization formula:** `skill_md_tokens_est = skill_md_bytes / 3.7` (GPT-4o empirical average for English markdown, $\pm 15\%$ error)
- Telemetry window:** 19 days (2026-04-16 to 2026-05-05) — `invocations_30d` is lower bound, not exact 30-day count
- Sessions per month baseline:** 600

- **Price per Mtok (input):** \$3.00

Key findings from CSV:

- **32 phantom rows:** Exist in skill-registry.db but no corresponding directory on disk (e.g., `algorithmic-art`, `brand-guidelines`, `sentry-*`, `tob-*` variants)
- **17 unregistered:** Disk skills NOT in skill-registry.db (`ask-board`, `prompt-forge`, `task-postflight`, `lightrag-*`, etc.)
- **Ghost invocations:** `mehanik` (186 hits), `update-config` (1 hit) — no disk directory exists
- **TOB nested structure:** All 12 `tob-*` skills have `README.md + skills/` subdir at root; no `SKILL.md` — invisible to flat-loader

[→ See Top-20 Refactor Priority Table](#)

Appendix B — Inventory README

Source: `~/system/specs/agent-ic-os-pillar4-skills-inventory.README.md`

Frequency Sources

Two independent sources per spec §CONSTRAINTS #4:

1. **skill-use.log (PRIMARY)** — `/Users/makinja/system/logs/skill-use.log`
 - 300 total entries
 - 267 entries within 30d window (2026-04-04 to 2026-05-04)
 - Hook fires on `SKILL=<name>` log line
2. **skill-registry.db (SECONDARY)** — `/Users/makinja/system/databases/skill-registry.db`
 - Query: `SELECT name, use_count FROM skills ORDER BY use_count DESC;`
 - 94 rows in registry

Hook Coverage Gap

- Hook fires only when `tool_name="Skill"` (i.e., the `/command` invocation)
- **Skills invoked by sub-agents as inline tool calls (not via Skill tool) are NOT counted**
- **Skills triggered via `references/` file Read operations are NOT counted**
- `invocations_30d` column = **minimum lower bound**, not exact count

Monthly Cost Assumptions

```
SESSIONS_PER_MONTH_BASELINE = 600
```

```
PRICE_USD_PER_MTOK_INPUT = 3.00
```

```
est_monthly_context_tokens = (frontmatter_description_bytes / 3.7) * 600
```

```
est_monthly_invocation_tokens = skill_md_tokens_est * invocations_30d
```

```
est_monthly_cost_usd = (est_monthly_context_tokens + est_monthly_invocation_tokens) * 3.00 /  
1,000,000
```

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.

Cross-links:

- [Top-20 Refactor Priority Table](#)
 - [Progressive Disclosure Design Pattern](#)
 - [PoC: task-postflight Refactor](#)
-

Revision #2

Created 2026-05-05 12:59:38 UTC by John

Updated 2026-06-07 20:01:12 UTC by John