

# T3 — Anthropic agent layer audit

## Anthropic SDK + Agent Loop Audit (MC #10357 T3)

**Lens:** Anthropic Chief AI Architect (Krieger / Kaplan / Cherny / MCP-team composite). **Layer scope:** Claude Code Task-tool sub-agent dispatch surface, hooks, prompt caching, tier routing, ZAKON #28 max-depth, dispatch cap. **Date:** 2026-04-30 22:01 CEST. **Method:** Read-only filesystem inspection. Tool-first per ZAKON NULA.

> **Important scope correction up front:** the CEO brief listed > `~/system/tools/agent-orchestrator.js` as canonical source. That file does not > exist at that path. The only `agent-orchestrator.js` on disk lives at > `~/system/archive/orphan-orchestrators-2026-03-21/agent-orchestrator.js` — > archived 5+ weeks ago. The actual orchestration surface today is: > 1. Claude Code Task tool (Anthropic-canonical, the dominant path), > 2. `~/system/tools/chain-runner.js` (YAML chains, lightly used), > 3. `~/system/kernel/pi-orchestrator.js` (Ollama DAG, separate process). > This doc audits #1 (the Anthropic layer); #2 and #3 fall under T2/T4.

---

## 1. Agent Loop Coherence

### What's there

Today's dispatch path:

CEO prompt

```
-> John (orchestrator, Sonnet by default)
  -> SlashCommand /mehanik (PreToolUse: pre-dispatch-gate.sh validates 13-field marker)
    -> writes /tmp/mehanik-cleared-<MC_ID>
  -> Task tool dispatch (subagent_type=<specialist>)
    -> PreToolUse hook stack: lock-john-dispatch-cap, claude-hooks pre,
      pre-action-da-gate, pre-dispatch-gate, john-max-depth-gate
    -> Claude Code spawns isolated sub-agent process from
      ~/.claude/agents/<name>.md frontmatter (model + tools)
```

```
-> sub-agent runs its own tool loop; only final assistant message
    returns to John
-> /task-postflight (Proveo verification, postflight marker)
-> mc.js done (state injector + alai-hooks evidence-gate enforce marker)
-> CEO report
```

## Anthropic-canonical or not? — verdict per piece

| Component | Verdict | Note | |---|---|---| | Task tool sub-agent dispatch | **Canonical** | Exactly the SDK pattern: isolated context, frontmatter declares model + tools. Result, not reasoning, returns to parent. | | `~/claude/agents/\*.md` with `name:`, `model:`, `tools:`, `description:` frontmatter | **Canonical** | Matches Claude Code documented sub-agent schema. 40 of 40 files have frontmatter. | | PreToolUse hook on `Task\|Agent` matcher | **Canonical primitive, custom payload** | Hooks are an Anthropic primitive; the policies enforced on top (Mehanik clearance, dispatch cap, max-depth) are ALAI-specific. This is the supported pattern. | | Mehanik gate as pre-dispatch deterministic check | **Custom but Anthropic-blessed shape** | Mehanik itself is a sub-agent (`~/claude/agents/mehanik.md`, model: sonnet, tools: Read+Bash) bootstrap-exempt from cap and depth hooks. Writes a marker file consumed by `pre-dispatch-gate.sh`. Anthropic's own pattern: "use a small judge sub-agent + a deterministic hook to enforce." This is *exactly* what Krieger has demoed publicly — Mehanik is well-shaped. | | 13-field marker schema in `/tmp/mehanik-cleared-\*` | **Custom, sensible** | Deterministic state hand-off. Field validation in pre-dispatch-gate.sh lines 73-79 is the right shape. | | ZAKON #28 trip-wire 1 (depth  $\geq 3$  blocks) | **Custom, defensible** | Anthropic does not ship a depth guard. But for a 1-CEO orchestrator with no parallel team, it matches the Boris-Cherny stated principle: "agents are loops; chains-of-chains are workflows pretending to be agents." | | ZAKON #28 trip-wire 2 (emergent-spawn budget approved+3) | **Custom, sensible** | Direct response to MC #10043 drift. Budget-against-Mehanik clearance is a clean primitive. | | Dispatch cap 3/session | **Anthropic would say: too low** | See Section 4. | | Boot enforcer / validation-state-injector on `UserPromptSubmit` | **Canonical primitive, cost-amplifying** | UserPromptSubmit injection is supported. But `validation-state-injector.sh` runs `mc.js list --owner john --priority H --json` on every cache miss (30s TTL) and prepends warnings to every prompt. See Section 3. |

## What's diverged from Anthropic-canonical

1. **Sub-agent definitions split across two stores.** `~/claude/agents/\*.md` (40 files) is what Claude Code actually reads. `~/system/agents/definitions/` (129 files) is a parallel store. Memo `feedback\_agent\_definitions\_dual\_store.md` documents this with a sync script. **Anthropic's recommendation: one source of truth.** 2. **Specialist-mapping JSON drift.** 40 FS agents vs 30 mapped. 17 FS-only files (the personas: agentforge, codecraft, finverge, vizu, skybound, mehanik, builder, validator, devils-advocate, alem-clone, dzevad-jahic, sylfest-lomheim, gemini-reviewer, redzo-reviewer, proveo, securion, flowforge) have no mapping entry. 7 mapped names point to FS files that don't exist (dorota-huizinga, hadi-hariri, james-bach, lee-robinson, lisa-crispin, maria-santos, resolver). **The mapping JSON has structural rot.** 3. **No**

**prompt-cache breakpoints declared anywhere.** Sub-agent .md files do not place explicit cache markers. With CLAUDE.md ~44 lines plus orchestration-surface.md (89 lines) plus rule-pointers, the system prompt is short enough that Claude Code's automatic 5-min ephemeral cache should hit on stable content — but with 6 UserPromptSubmit hooks injecting dynamic content (boot-enforcer, alem-instruction-checker, validation-state-injector, alai-hooks auto-verify, feasibility-check-advisory, incident-response-mode), **cache invalidation is happening every prompt.** Section 3 quantifies. 4. **Mehanik is sonnet but does pure deterministic verification.** `~/claude/agents/mehanik.md` line 3: `model: sonnet`. Mehanik's job is to grep blueprint files, count constraints, and write a marker. This is Haiku work, not Sonnet. See Section 5.`

**Verdict — Loop coherence:** *Anthropic-canonical core, with custom enforcement that fits the SDK seams correctly. The shape is right. The drift is in inventory hygiene and tier choices.*

---

## 2. Sub-Agent Inventory Health

### Counts (tool-verified)

```
ls ~/.claude/agents/*.md | wc -l           → 40
jq '.mappings | length' specialist-mapping.json → 30
ls ~/system/agents/definitions/ | wc -l    → 129
```

### Tier breakdown across 40 FS agents

| Tier | Count | Files | |-----|-----:|-----| | opus | 4 | anthropic-chief-architect, dzevad-jahic, openai-chief-architect, sylfest-lomheim | | sonnet (incl. claude-sonnet-4-5) | 30 | builder, mehanik, all PI personas (codecraft/proveo/vizu/...), all expert personas except 4 above | | haiku | 6 | claude-code-guide, devils-advocate, gemini-reviewer, redzo-reviewer, sentinel-validator, validator | | no model field | 0 | (clean) |

**Anthropic-recommended tier split for orchestrator-heavy single-user system:** ~50% Haiku (reads/triage/grep), ~40% Sonnet (default work), ~10% Opus (planning + novel architecture). **Current: 15% / 75% / 10%.** Sonnet is over-loaded.

### Anti-pattern findings

**AP-A: Persona agents acting as builders without strict scope.**

- `builder.md` (321 lines, sonnet) is a generic "build this" catch-all with full Read/Write/Bash/Edit. Risk: John reflexively dispatches "builder" instead of resolving via discover.js routing` to a domain specialist (kelsey-hightower for K8s, bruce-momjian for Postgres, etc.). The builder file should either be deleted (force routing through specialists)`

or hard-scoped to "non-domain trivial scaffolding only".

- `proveo.md`, `securion.md`, `vizu.md`, `codecraft.md`, `flowforge.md`, `finverge.md`, `agentforge.md`, `skybound.md` are **company-level umbrella personas**, not agents. Each company already has named experts (kelsey for FlowForge, brad-frost+lea-verou for Vizy, etc.). The umbrella personas are a routing trap: ambiguous which specialist actually runs.

**AP-B: Mapped-but-missing agents.** 7 specialist-mapping entries (`hadi-hariri`, `lee-robinson`, `james-bach`, `lisa-crispin`, `dorota-huizinga`, `maria-santos`, `resolver`) point to .md files that **do not exist in** `~/claude/agents/`. Any dispatch routed to these via `discover.js` routing will crash at sub-agent spawn time. **This is a live broken edge.** **AP-C: Validator duplication.** Three validators: `validator.md` (haiku), `sentinel-validator.md` (haiku), `proveo.md` (sonnet umbrella). Plus `gemini-reviewer.md` and `redzo-reviewer.md`. Anthropic recommendation: one validator role with domain-routed sub-flavors. Five competing files = unclear authority. **AP-D: Two `\*-chief-architect` personas (anthropic + openai) both opus.** Useful for vendor-diverse architecture review, but each dispatch costs Opus rates. Make sure they are summoned with `[CEO\_APPROVED]` only, not by reflex.

## Inventory verdict

|| Count | |---|---:| | Healthy specialist agents (named expert, mapped, used) | ~17 | | Persona umbrellas (delete or hard-scope) | 8 | | Validator/reviewer overlap (consolidate) | 5 | | Bootstrap gates (mehantik, validator, devils-advocate) | 3 | | Mapped-but-missing (broken on dispatch) | 7 | | Total FS agents | 40 |

**Pruning target: 40 → 22.** Delete or merge 18 files. See Section 5.

---

## 3. Prompt Caching & Token Economy

### What lands in the system prompt every John prompt

Inspected hook config (`~/claude/settings.json`) — UserPromptSubmit fires 6 hooks:

```
UserPromptSubmit:
  -> incident-response-mode.sh           (state-dependent, mostly inert)
  -> boot-enforcer.sh                   (inserts 8-hour banner if no boot flag)
  -> user-message-logger.sh             (logging, no injection)
  -> alai-hooks auto-verify              (Kotlin CLI, ~14MB binary, injects verification)
  -> alem-instruction-checker.sh        (CEO directive parser, injects when triggered)
```

```
-> feasibility-check-advisory.sh      (ZAKON FEASIBILITY, injects if conflicts)
-> validation-state-injector.sh      (mc.js list output cached 30s)
```

**The cache-amplification problem:** every UserPromptSubmit hook that prints to stdout becomes part of the conversation. Even if Claude Code's automatic 5-min ephemeral cache were active on the static prefix (CLAUDE.md + tool schemas + sub-agent registry), **a hook-injected line that varies per-call evicts cache from that breakpoint forward.**

## `validation-state-injector.sh` cost analysis

Lines 50-110 — runs `node ~/system/tools/mc.js list --owner john --priority H --status open --json`, parses, prints WARNING lines per H-task missing postflight. Cached 30s in `~/system/state/validation-state-cache.txt`.

- 30s TTL is sensible.
- BUT: the output text changes whenever tasks change state, postflight markers appear, tasks close. **Each change = cache breakpoint invalidation downstream.**
- Estimated frequency: 5-20 invalidations per hour during active session.

**Anthropic recommendation:** UserPromptSubmit hooks that inject dynamic state should **append after** the stable system prompt and **not** be read as part of the cached prefix. The current architecture treats them as additional system context, defeating prefix cache.

## `mc.js list` injection on every prompt — confirmed cost amplifier

Yes, `validation-state-injector.sh` runs `mc.js list --owner john --priority H --status open --json` (or hits cache). Even at 30s cache, on a 4-hour session = ~480 prompt submissions × full output included = constant cache invalidation, plus ~30 mc.js invocations (cache miss path).

**Quantified:** assume CLAUDE.md + tool schemas = ~3500 cached tokens; each invalidation forces re-tokenization. At Sonnet pricing \$3/MTok input / \$0.30/MTok cache-read, the gap = \$2.70 per 1M tokens × N invalidations. At 480 prompts × 3500 tokens of *should-have-been-cached* prefix = 1.68M tokens. If 50% miss cache due to injector churn = ~\$2.27 saved per 4-hour session if injector becomes append-only (or moves to a separate context window).

## YouTube RAG ingestion — verified Ollama, NOT Anthropic

`~/system/kernel/pi-orchestrator.js`:

- Line 4796: `[IDLE] YouTube batch started (PID ' + child.pid + '), model=qwen3:8b-q8\_0 on FORGE`
- 71 references to OLLAMA / 11434 / 10.0.0.2 (FORGE box) vs 2 references to "anthropic"
- The 2 anthropic references are imports/strings, not API calls

**Verdict: pi-orch idle YouTube ingestion is on local Ollama (FORGE 10.0.0.2:11434, qwen3:8b-q8\_0). ZERO Anthropic API spend on that path. Not a runaway cost. CEO concern resolved.**

## Caching verdict

| Component | Cache health | Anthropic verdict | |---|---|---| | CLAUDE.md (44 lines) | Should hit auto-cache | Good — short, stable | | orchestration-surface.md (89 lines) | Loaded on demand | Good — not in every prompt | | sub-agent .md frontmatter | Per-dispatch, isolated | Good — separate context | | UserPromptSubmit injectors | **Cache-defeating** | **Fix** | | YouTube RAG | Local Ollama, no API cost | Good | | mc.js list as ambient state | Injected per prompt | **Move to /status command, not auto-injected** |

---

## 4. Max-Depth ZAKON #28 + Dispatch Cap

### Trip-wire summary (tool-verified from hook source)

| Wire | File | Behavior | Bypass | |---|---|---|---| | Dispatch cap | `lock-john-dispatch-cap.sh` | Block 4th `Task\Agent` per session | `[CEO\_APPROVED]` token in prompt | | Trip-wire 1 (depth) | `john-max-depth-gate.sh` lines 71-118 | Block dispatch where target MC has  $\geq 3$  ancestors | | | Trip-wire 2 (emergent-spawn) | `john-max-depth-gate.sh` lines 122-195 | Block `mc.js add` when emergent count > approved+3 | | | Trip-wire 3 (cross-domain) | `john-max-depth-gate.sh` lines 197-260 | SOFT — write drift-stop memo, allow once | n/a (soft) |

Bootstrap-exempt subagents (counted nowhere): `mehanic`, `validator`, `devils-advocate`, `anthropic-chief-architect`.

## Are these correctly tuned?

**Dispatch cap = 3 / session.** This is the wire that blocked us today and forced `[CEO\_APPROVED]` for a 5-team architecture audit. **Anthropic position on parallel sub-agents (per Boris Cherny**

## / public Claude Code talks):

- The Task tool *intentionally* allows many sub-agents in parallel — that is the Anthropic optimization story (sub-agents have isolated context, so spawning 5 reviewers in parallel does NOT bloat parent context).
- The well-known Anthropic example: research agent dispatches 5–10 sub-search agents in parallel, fans-in summaries.
- Anthropic does NOT publish a "max sub-agents per session" guardrail. The intended limiter is *cost telemetry + budget alerts*, not a hard count.

**3/session is too low for legitimate orchestrator workflows.** For a session that does: 1. Mehanik clearance (exempt) 2. Specialist build dispatch 3. Proveo validation (counted) 4. Documentation via Skillforge (counted) = already 3 dispatches for ONE task. A second task in same session = blocked.

# Recommended retuning

Anthropic-aligned tuning:

| Setting | Today | Recommend | Rationale | |---|---:|---:|---| | Dispatch cap (counted) | 3 | **8** | One task = ~3 (build + proveo + skillforge); 8 = 2-3 tasks/session before review | | Max-depth (gen 4+ blocked) | yes | **keep yes** | Catches the #10043 pathology — directly load-bearing | | Emergent-spawn budget (approved+3) | yes | **keep yes** | Load-bearing per ZAKON #28 Section "If Only One Rule" | | Bootstrap exemption list | mehanik, validator, devils-advocate, anthropic-chief-architect | **add: skillforge, proveo (validator), all reviewers (gemini-reviewer, redzo-reviewer)** | Validators and doc-writers are mechanical, not strategic | | `[CEO\_APPROVED]` log | logged | **add: weekly review trigger** | Per ZAKON #28: if overrides > 2/week, escalate. Currently no automation. |

**Net effect:** legitimate multi-team audits like today's MC #10357 (5 teams × at least 1 dispatch each = 5) would not require `[CEO\_APPROVED]`; only *drift* dispatches would.

# Verdict

- Dispatch cap **3 is too restrictive**, not Anthropic-aligned. Raise to 8.
- ZAKON #28 trip-wires 1 + 2 are well-designed. **Keep as-is.**
- Trip-wire 3 (soft, cross-domain memo) is good — exactly the Anthropic "soft signal + audit trail" pattern.
- Bootstrap exemption list should grow to include validators and skillforge.

---

# 5. Day-to-Day Agent Layer Perfection (Single-User CEO Orchestrator, Zero Revenue)

## What MUST remain

### Identity / orchestration core (no compromise):

- `~/claude/CLAUDE.md` (44 lines) — Anthropic-canonical orchestrator system prompt. Keep small.
- `~/claude/agents/mehanik.md` — pre-dispatch deterministic gate. **Tune model from sonnet → haiku.**
- `~/claude/agents/validator.md` — postflight verification (haiku, correct).
- `~/claude/agents/devils-advocate.md` — pre-action critique (haiku, correct).

### Specialist roster (keep, all properly mapped, all opus/sonnet justified):

- CodeCraft: martin-kleppmann, bruce-momjian, petter-graff, sentinel-developer, sindresorhus (5)
- Vizu: brad-frost, lea-verou (2)
- Securion: parisa-tabriz, sentinel-architect (2)
- Skybound: paul-hudson (1)
- AgentForge: chip-huyen, georgi-gerganov, anthropic-chief-architect (opus, persona) (3)
- Finverge: markos-zachariadis (1)
- FlowForge: kelsey-hightower (1)
- Lexicon: lexicon (1)
- SkillForge: skillforge (1)
- Proveo: angie-jones, sentinel-tester (2)

### Total essential specialist roster: 19. Plus 3 bootstrap = 22 files. Hooks that earn their cost:

- `pre-dispatch-gate.sh` — Mehanik marker enforcement. **Keep.**
- `john-max-depth-gate.sh` — ZAKON #28 trip-wires. **Keep, retune as Section 4.**
- `lock-john-dispatch-cap.sh` — **Keep, raise cap to 8.**
- `bash-danger-gate.sh` — destructive-command block. **Keep.**
- `alai-hooks` (Kotlin) bash + evidence-gate — **Keep, has been validated 2026-04-29.**
- `postflight-gate.sh` — Hard Constraint #3 enforcement. **Keep.**
- `boot-enforcer.sh` (UserPromptSubmit) — **Keep**, 8-hour TTL is fine.

## What can be deleted / consolidated

**1. Persona umbrella agents (8 files, ~32KB):** `agentforge.md`, `codecraft.md`, `finverge.md`, `flowforge.md`, `proveo.md`, `securion.md`, `skybound.md`, `vizu.md`. *Rationale:* every

company has named specialists; umbrella agents create routing ambiguity. **Delete.** Force `discover.js routing` to resolve to named expert. **2. Builder catch-all:** `~/claude/agents/builder.md`. *Rationale:* 321 lines of generic "build this" tools is the AP-A trap. **Delete or hard-scope** to "<150 LOC trivial scaffolding only, no production paths". Specialists already cover the real work. **3. Reviewer overlap (3 files):** `gemini-reviewer.md`, `redzo-reviewer.md`, `sentinel-validator.md`. *Rationale:* `validator.md` + named domain experts already cover review. **Consolidate to 1.** Pick the strongest of the three. **4. Persona ghosts:** `alem-clone.md`, `dzevad-jahic.md` (opus!), `sylfest-lomheim.md` (opus!). *Rationale:* unmapped, opus-tier, unclear use case. **Verify with CEO; archive if not actively used.** Each unscheduled Opus dispatch is a money fire. **5. Mapping JSON cleanup:** delete 7 mapped-but-missing entries (`dorota-huizinga`, `hadi-hariri`, `james-bach`, `lee-robinson`, `lisa-crispin`, `maria-santos`, `resolver`) OR populate the .md files. **Live broken edge** today. **6. Mehanik tier:** retune `~/claude/agents/mehanik.md` line 3 from `model: sonnet` to `model: haiku`. Mehanik does grep + count + marker write — pure I/O. Sonnet is overkill; saves ~5x on every pre-dispatch check. **7. Validation injector demotion:** move `validation-state-injector.sh` from `UserPromptSubmit` matcher to a `/status` slash command. Stop polluting prompt cache. CEO can ask "what's pending?" when it matters.

## Concrete pruning + retention list

**RETAIN (22 files in `~/claude/agents/`):**

```
mehanik.md (retune to haiku)
validator.md
devils-advocate.md
anthropic-chief-architect.md
openai-chief-architect.md
martin-kleppmann.md
bruce-momjian.md
petter-graff.md
sentinel-developer.md
sentinel-architect.md
sentinel-tester.md
sindre-sorhus.md
brad-frost.md
lea-verou.md
parisa-tabriz.md
paul-hudson.md
chip-huyen.md
georgi-gerganov.md
markos-zachariadis.md
kelsey-hightower.md
```

```
angie-jones.md
lexicon.md
skillforge.md
claude-code-guide.md (haiku, useful for hook design Q&A)
```

### **DELETE (8 umbrella personas):**

```
agentforge.md, codecraft.md, finverge.md, flowforge.md,
proveo.md, securion.md, skybound.md, vizu.md
```

### **REVIEW WITH CEO (4):**

```
builder.md (delete or hard-scope)
alem-clone.md (purpose unclear)
dzevad-jahic.md (opus – justify or delete)
sylfest-lomheim.md (opus – justify or delete)
```

### **CONSOLIDATE TO 1 (3 → 1):**

```
gemi-reviewer.md, redzo-reviewer.md, sentinel-validator.md
→ keep one, delete two
```

### **FIX MAPPING:**

```
~/system/agents/specialist-mapping.json:
- delete or stub-create 7 mapped-but-missing entries
- add the 8 umbrella deletes (or remove mapping entirely once .md files go)
```

## Net inventory after pruning

```
40 .md files → 22 .md files
30 mapping entries → 19 entries (1:1 with valid FS)
6 UserPromptSubmit hooks → 4 (move validation-state-injector to /status)
Mehanik: sonnet → haiku (~5x cheaper per gate run)
Dispatch cap: 3 → 8 (Anthropic-aligned)
```

---

# ANTHROPIC RECOMMENDATIONS — ranked

1. **(this session)** Raise dispatch cap 3 → 8 in `~/claude/hooks/lock-john-dispatch-cap.sh` line 77. Single-line change. Restores Anthropic-canonical multi-sub-agent ergonomics. 2. **(this week)** Retune `~/claude/agents/mehanik.md` model from sonnet → haiku. Mehanik is mechanical I/O. Saves ~80% per dispatch check. 3. **(this week)** Move `validation-state-injector.sh` out of UserPromptSubmit; expose as `/status` slash command. Restores cache hits. 4. **(this week)** Fix specialist-mapping.json — 7 mapped-but-missing entries are a live broken-dispatch trap. 5. **(this month)** Delete 8 umbrella persona files. Force routing through named specialists. 6. **(this month)** Consolidate 3 reviewer files into 1. 7. **(this month)** Add weekly cost telemetry on `[CEO\_APPROVED]` overrides — ZAKON #28 already specifies "if > 2/week, escalate" but no automation exists. Hook into `~/claude/hooks/john-max-depth-gate.log` parser → cron → CEO Slack DM. 8. **(this month)** Audit 4 opus-tier persona files (anthropic + openai chief-architects justify; dzevad-jahic + sylfest-lomheim must justify or downgrade).

---

## CONFIDENCE: HIGH

Tool-verified file counts, frontmatter, hook configurations, pi-orch model provider, cache-injector behavior. All claims grounded in `~/claude/agents/`, `~/claude/hooks/settings.json`, `~/system/kernel/pi-orchestrator.js`, `~/system/agents/specialist-mapping.json`.

## RISKS

- Deleting umbrella persona files may break implicit references in older skills/commands. **Mitigation:** grep for "agentforge.md", "codecraft.md", etc. before delete; replace with named specialist routing.
- Raising dispatch cap to 8 without monitoring could re-enable the #10043 drift pattern. **Mitigation:** ZAKON #28 trip-wires 1+2 remain — they catch *recursive* drift, which is the actual failure mode. Cap raises only the *parallel-fan-out* limit.
- Moving validation-state out of auto-injection means John might forget to check pending postflights. **Mitigation:** add a single-line check in PostToolUse on `mc.js done` — block if marker missing (already done by `alai-hooks evidence-gate`).

---

## File Path Manifest (audit references)

- `~/Users/makinja/.claude/CLAUDE.md`
- `~/Users/makinja/.claude/agents/` (40 .md files inspected)
- `~/Users/makinja/.claude/agents/mehanik.md`
- `~/Users/makinja/.claude/agents/anthropic-chief-architect.md`
- `~/Users/makinja/.claude/hooks/lock-john-dispatch-cap.sh`
- `~/Users/makinja/.claude/hooks/john-max-depth-gate.sh`

- ``/Users/makinja/.claude/hooks/pre-dispatch-gate.sh``
  - ``/Users/makinja/.claude/hooks/boot-enforcer.sh``
  - ``/Users/makinja/.claude/hooks/validation-state-injector.sh``
  - ``/Users/makinja/.claude/settings.json``
  - ``/Users/makinja/system/agents/specialist-mapping.json``
  - ``/Users/makinja/system/rules/orchestration-surface.md``
  - ``/Users/makinja/system/rules/zakon-28-max-depth-boundary.md``
  - ``/Users/makinja/system/kernel/pi-orchestrator.js``
  - ``/Users/makinja/system/tools/chain-runner.js``
  - ``/Users/makinja/system/archive/orphan-orchestrators-2026-03-21/agent-orchestrator.js``  
*(archived; brief mistakenly listed as canonical)*
- 

Revision #2

Created 2026-05-01 08:08:30 UTC by John

Updated 2026-06-07 20:00:45 UTC by John