

ALAI Opinion MC task plan — 2026-08-23

ALAI opinion findings ? Mission Control execution plan

Date: 2026-08-23

Opinion source: /Users/makinja/system/evidence/alai-rapport-opinion-2026-08-22/opinion-raw.md

Mission Control DB: /Users/makinja/system/databases/mission-control.db

Role split

The implementation split is explicit: **GPT builder** work uses `openai-codex/gpt-5.6-sol`; **Fable/Opus architect** work uses `claude-fusion/claude-opus-5`. Architects define the bounded operating contract and measurable governance. Builders implement and prove the Pi and Claude Code critical paths. This is a four-task execution batch, not a new broad ALAI rewrite.

Findings converted into MC tasks

#900140 — ALAI Lite 30-day critical-path operating contract (Architect)

- Role: architect
- Model: `claude-fusion/claude-opus-5`
- Opinion finding: ordinary delivery must become one short requirement-to-health-check path, with orchestration outside the critical path.
- Verbatim basis: “Requirement → coding agent edits correct repository → local build/tests → pull request → CI build/unit/lint/security → one human production approval → automated deployment → health check and rollback.”
- Deliverable: a 30-day architecture contract naming the only blocking controls, the observational controls, freeze exceptions, production approval, health probe and rollback

boundary.

#900141 — ALAI Lite Pi extension blocking-cap implementation (GPT Builder)

- Role: builder
- Model: `openai-codex/gpt-5.6-sol`
- Opinion finding: Pi should retain irreversible safety and runtime support, not task/RAG/evidence/checkpoint enforcement on normal coding.
- Verbatim basis: “The Builder recommended retaining only irreversible safety, context tracking, provider fallback, Git/CI protection, one production approval, and post-deployment health/rollback.”
- Deliverable: deterministic inventory of active versus blocking Pi extensions, a maximum-five blocking cap, ordinary-edit canary, safety-negative canaries and rollback proof. Fusion remains optional and non-blocking.

#900142 — ALAI Lite Claude Code five-gate convergence canary (GPT Builder)

- Role: builder
- Model: `openai-codex/gpt-5.6-sol`
- Opinion finding: Claude Code local development must not require task tokens, RAG, delegation, evidence documents or checkpoints.
- Verbatim basis: “For Claude Code, the opinion recommended eliminating task tokens, RAG requirements, delegation requirements, evidence-document requirements, and checkpoints from local development.”
- Deliverable: prove fresh CC sessions load exactly five irreversible-harm PreToolUse gates, safe Bash/Edit/Write works, legacy processes converge, and destructive/secret/deploy/project-path negative controls still block.

#900143 — ALAI Lite delivery SLO and weekly governance plan (Architect)

- Role: architect
- Model: `claude-fusion/claude-opus-5`
- Opinion finding: simplification succeeds only if delivery speed, false blocks, gate count and system-maintenance share are measured.
- Verbatim basis: “The success criteria were: first edit within five minutes, local feedback within ten minutes, pull request to production within thirty minutes, zero false-positive agent blocks, no more than five active critical-path gates, and less than ten percent of

engineering time spent maintaining the AI system.”

- Deliverable: a weekly scorecard and go/no-go rule that rejects new control-plane work unless it demonstrably shortens product delivery.

Execution plan / redoslified

1. **#900140 Architect first:** freeze the 30-day critical-path contract and interfaces. This prevents builders from inventing another orchestration layer.
2. **#900141 GPT builder:** implement and prove the Pi blocking-cap profile against the architect contract.
3. **#900142 GPT builder:** implement and prove CC five-gate convergence, including legacy-session handling.
4. **#900143 Architect closeout:** consume measured builder evidence and publish the weekly delivery/governance scorecard.

Tasks #900141 and #900142 may execute in parallel only after #900140 publishes the contract. #900143 is last because it evaluates actual evidence rather than planned behavior.

Started work

MC #900140 is the first started task. John begins by reading the opinion, checking current Pi/CC evidence and writing the first architecture-step record at

`/Users/makinja/system/evidence/900140/started-architecture-step.md`. No builder code begins before this contract step.

Machine-checkable manifest

```
{
  "opinion_source": "/Users/makinja/system/evidence/alai-rapport-opinion-2026-08-22/opinion-raw.md",
  "roles": {
    "builder": {
      "model": "openai-codex/gpt-5.6-sol"
    },
    "architect": {
      "model": "claude-fusion/claude-opus-5"
    }
  },
  "tasks": [
    {
```

```
"mc_id": 900140,
"title": "ALAI Lite 30-day critical-path operating contract (Architect)",
"role": "architect",
"assigned_model": "claude-fusion/claude-opus-5",
"opinion_basis": "The opinion requires one bounded requirement-to-deployment
architecture and removes orchestration from ordinary delivery.",
"opinion_evidence_quote": "Requirement → coding agent edits correct repository → local
build/tests → pull request → CI build/unit/lint/security → one human production approval →
automated deployment → health check and rollback."
},
{
"mc_id": 900141,
"title": "ALAI Lite Pi extension blocking-cap implementation (GPT Builder)",
"role": "builder",
"assigned_model": "openai-codex/gpt-5.6-sol",
"opinion_basis": "The opinion narrows Pi to irreversible safety and runtime support
rather than task-process enforcement on every action.",
"opinion_evidence_quote": "The Builder recommended retaining only irreversible safety,
context tracking, provider fallback, Git/CI protection, one production approval, and post-
deployment health/rollback."
},
{
"mc_id": 900142,
"title": "ALAI Lite Claude Code five-gate convergence canary (GPT Builder)",
"role": "builder",
"assigned_model": "openai-codex/gpt-5.6-sol",
"opinion_basis": "The opinion explicitly removes task, RAG, delegation, evidence-
document and checkpoint prerequisites from Claude Code development.",
"opinion_evidence_quote": "For Claude Code, the opinion recommended eliminating task
tokens, RAG requirements, delegation requirements, evidence-document requirements, and
checkpoints from local development."
},
{
"mc_id": 900143,
"title": "ALAI Lite delivery SLO and weekly governance plan (Architect)",
"role": "architect",
"assigned_model": "claude-fusion/claude-opus-5",
"opinion_basis": "The opinion defines measurable delivery and maintenance targets that
require an architect-owned weekly governance contract.",
"opinion_evidence_quote": "The success criteria were: first edit within five minutes,
local feedback within ten minutes, pull request to production within thirty minutes, zero
```

false-positive agent blocks, no more than five active critical-path gates, and less than ten percent of engineering time spent maintaining the AI system."

```
    }
  ],
  "execution_order": [
    900140,
    900141,
    900142,
    900143
  ],
  "started": [
    {
      "mc_id": 900140,
      "actor": "john",
      "evidence_path": "/Users/makinja/system/evidence/900140/started-architecture-step.md"
    }
  ]
}
```

Revision #1

Created 2026-08-22 23:15:56 UTC by John

Updated 2026-08-22 23:15:56 UTC by John