lzt-harness

Source: Private · **Status:** Active · **Period:** 2025-02 (active)

Private operator harness for Claude Code: multi-agent routing, destructive-command hooks, and an intent-filtered MCP broker.

The challenge

A single-agent Claude Code session loses context faster than I can ship useful work, and the failure modes of unattended agentic execution (silent git operations, accidental file deletion, runaway tool calls) need belt-and-suspenders enforcement rather than a hopeful prompt.

My role

I built a multi-agent harness on top of Claude Code with planner, implementer, and reviewer routing, hooks that block destructive operations by default, and an MCP broker that surfaces intent in response metadata.

What changed

Effective session length

Baseline: About 30 minutes per single-agent session

4–6 hours across multi-agent sessions

Silent work-loss incidents

Baseline: About 2 per month

0 in 2026-Q3 (remaining cases caught by the teleport safety net as advisories)

Audit chain readability

Baseline: Opaque — single agent, no per-tool trace

Per-tool intent log, browsable after the session

Stack

Go | Rust | Python | Shell | MCP