

FLEET COMMAND

An AI ops crew that does the work. A human gate that holds the trigger.

Four specialist agents work a target site in sequence, then stop. SCOUT reconns the site and the market around it, AUDIT ranks the defects, MEDIC drafts the fix as a diff, and SHIP stages the deploy with a risk summary and **holds**. The hold is structural, not a dialog: SHIP has no code path that releases without a human click.

VERIFY IT LIVE IN 30 SECONDS

- Open fleetcommand-2u0.pages.dev
- Press **Run mission**, top right
- Watch SCOUT, AUDIT and MEDIC finish
- SHIP goes amber and reads **HOLDING**
- The release needs **two keys**, turned by a person

No login, no signup, nothing to install.

WHY THE GATE IS THE PRODUCT

Agent fleets are shipping code today with no approval layer. Full autonomy is reckless; approving every keystroke is useless.

Fleet Command puts the whole crew on one side of one line and a named human on the other. That line is the thing enterprises actually buy.

SERPAPI — BEST AI USE CASE

SCOUT's market lens. One Google search per mission tells the crew what the target competes against, which changes what AUDIT ranks first.

- `functions/_lib/market.js` — the integration
- `tools/test-market.mjs` — parser tests

Cached 12h in Workers KV, so a public Run button cannot burn the quota. One search per mission, SCOUT only.

NAME.COM — DOMAIN API CHALLENGE

SHIP's pre-flight. Before staging, it asks whether the business owns its own name. In the demo it does not, and that becomes a real line in the risk summary.

- `functions/_lib/domains.js` — the integration
- `tools/test-domains.mjs` — parser tests

Registrable is not cleared, and SHIP says so rather than telling an owner a name is theirs.

BUILT WITH, AND WHAT IS HONEST ABOUT IT

Cloudflare Pages hosts the cockpit; one Pages Function runs each agent step against the Anthropic API. Zero dependencies, no build step, no framework. Security headers on every response via middleware. The demo target is a fictional bakery broken on purpose, so nothing real is ever touched.

Three independent switches, labelled on screen. With no model key the run is a recording and the chip reads `replay`. With no SerpApi key SCOUT shows `serpapi sample`. With no name.com credential the registrar check is labelled sample. Arm any key and the identical code path goes live. The app never passes a canned answer off as live, and the mode is always visible.