

JobSpark Systems — Operational Systems Proof Brief

Operational systems architecture · deterministic, event-driven automation · governed AI

Prepared by Kyle Spivey · Ponte Vedra Beach, FL · spiveykyle111@gmail.com · 904-315-6646

JobSpark runs the live state of a real operation — calls, jobs, crews, trucks, contracts, and payroll — as a single, always-current picture. I designed and built it solo: 70,400 lines of production code where deterministic, event-driven automation owns the business process and AI is confined to just three points. This brief shows how a job moves through the system, the numbers behind it, and the architecture underneath.

How a job actually moves through JobSpark

Almost everything below is deterministic, event-driven automation. Humans only answer the phone, watch the live board, and approve reviews. The three amber tags mark the only places an AI model runs.




Figure 1. One job's lifecycle — top row left→right, bottom row right→left. Blue is automated, teal is human, amber marks the 3 LLM call sites.

The numbers

Every count is code-verified against branch main (snapshot Jul 28, 2026) and labeled production vs. partial vs. demo — not marketing math.

AREA	AUDIT-VERIFIED SNAPSHOT
Codebase	70,400 production lines across 245 files — designed and built solo.
Frontend / backend	158 files (37,700 lines) · 87 files (32,700 lines) · 71 API routes.
Workflows	56 normalized production workflow groups across an 18-stage pipeline.
State model	21 state machines · 79 finite state values · guarded transitions.
Event history	41 canonical action / event / history types · append-style history + replay.
Exceptions	34 categories · 16 failure-routing points · 7 repair / retry / reconcile mechanisms.
Human gates	11 human-in-the-loop approval boundaries.
AI surface	3 direct LLM call sites — intake parse · query agent · closeout contract parse.
Security & tenancy	5 workspaces · 5 roles · route middleware + Firestore rules on financial / terminal writes.
Data model	33 Firestore collection types · 24 indexes.
Communications	Providers hooked up — status / inbound webhooks · crew SMS ingest · reply interpretation.
Testing	42 test files · 167 test-case signals (38 failure-path) · emulator-backed financial-integrity paths.

How JobSpark is different

- **Deterministic core, bounded AI.** Only 3 LLM call sites in 70,400 lines. Payroll, closeout, assignment, warnings, and repair are exact — server-guarded and transaction-protected. AI only reads and interprets: intake, history/ops questions, and the closeout contract the crew sends in.
- **Humans touch four things.** Answer the phone, watch the live board, and approve the closeout and maybe-job reviews. Everything else runs as deterministic, event-driven automation.
- **Evidence over adjectives.** Every number traces to a file and line against branch main, labeled production vs. partial vs. demo. The claims are meant to survive an engineer reading the code.
- **Built from a real operation backward.** Exercised in parallel beside a real moving operation averaging about 30 jobs per week, and designed around what people do under pressure — not a clean demo. It is not yet that company's system of record.

The architecture, in one view

The same system, seen whole: a thin, agent-assisted surface on top and a deterministic factory underneath — 56 coordinated workflows, 21 state machines, 34 exception categories, and full event history, with the AI surface kept deliberately tiny. Each panel cites the code that backs it.

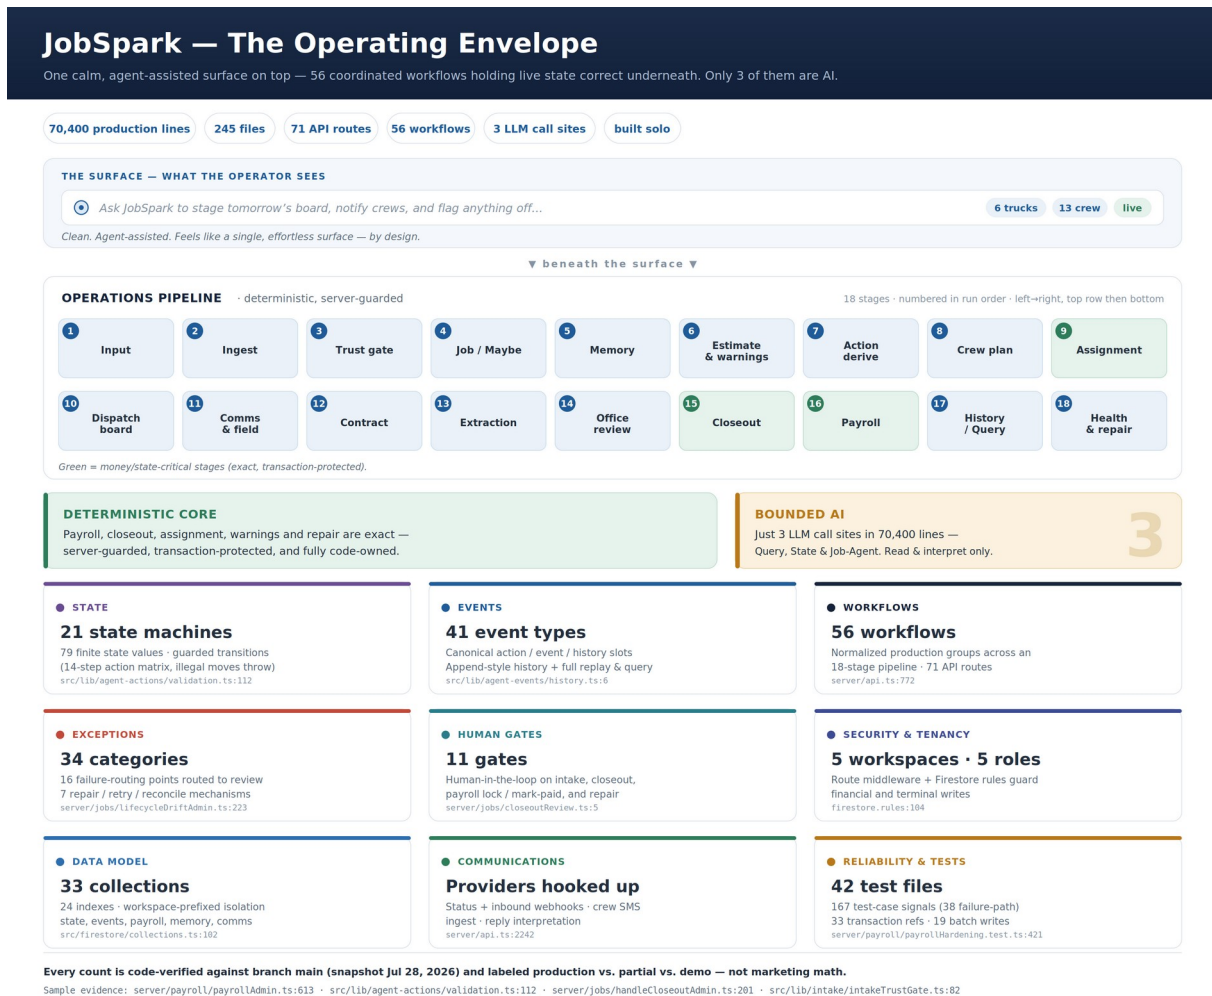


Figure 2. The operating envelope — deterministic core, bounded AI, and the engineering scale beneath the surface.

Live platform: jobsparksystems.app · Website: jobsparksystems.com · LinkedIn: linkedin.com/in/kylespivey1

The point isn't more AI. It's the operational engine that makes a small, bounded amount of AI safe, measurable, and useful.