

# ANUBHAV SOOD

AI skills and systems: a one-page addendum to the CV. Everything below I designed, built and operate (Jun 2025 to Sep 2026); counts measured 2026-09-07

Bengaluru, India | +91 63669 77533 | anubhav.sood.1701@gmail.com | [linkedin.com/in/anubhavsood17](https://www.linkedin.com/in/anubhavsood17) | [efficientearners.com](https://efficientearners.com) | [mylittlelore.com](https://mylittlelore.com)

---

## SYSTEMS BUILT AND RUNNING

- **Agent operating system.** 90 AI agents and 25 skills (49 frontier-tier, 33 fast-tier, none on the smallest) organised in lanes: command, engine, platform, growth, compliance, desk, and three media studios. A 29-row task router with a six-rule model law; one owner per file. Two read-only veto agents (compliance auditor, red team) at every release phase; a daily fleet-learning loop as the only path from a run-time lesson to an agent definition. A 10-researcher  $\square$  3-synthesiser workflow; a pre-tool-use hook. Nine pre-wired local dev servers.
- **Live commerce platforms.** The Desk: Cloudflare Pages Functions, KV and Durable Objects, Google OAuth, Razorpay webhooks (signature  $\square$  idempotency  $\square$  own record  $\square$  provider corroboration  $\square$  grant), 18,969 lines of JS. mylittlelore.com: PayPal checkout Worker, R2 signed 24-hour delivery, KV ledger; 14-agent, 25-aspect audit fixed a stranded-buyer bug on launch day. Public site on Cloudflare Pages behind Wix DNS, 200 commits.
- **Data and tools.** NSE market-data engine over 87 symbol trees: five-phase contract-driven pipeline, sha-pinned computation, launchd daily and weekly runs, 15,117 Pine Script v6 files emitted; "Blocks", a five-horizon market-structure tool over 212 instruments with a 05:30 fail-closed publish gate. Two workstations with 314 and 435 self-tests; a costing OS reconciling 2,581 founder hours with AI spend as a cost line.
- **Media and course production.** Two multi-agent video studios (14 and 17 agents): script gates, scene maps, ElevenLabs voice-over, Hindi subtitle burn-in, 9:16 cuts, film QA before render. A course studio on a written pedagogy produced a 72-chapter, 3h16m course and a 240-section, ~10-hour AI-voiced course through 21 adversarial verifier passes and a banned-pharse scan (240/240 clean).
- **Marketing organisation.** Twelve agents under one front door with a judgment/mechanics split; a 12-platform, three-tier distribution map with UTM attribution and a no-silent-skip rule. A biweekly Platform Pulse in which an agent drives my logged-in Chrome via a browser connector across 13 dashboards read-only and returns a report, dashboard readings and a docket routed to the owning lane; Brevo lifecycle drafted by agents, sent by me.

## PRACTICES I RUN AS LAW

- **Context engineering.** A three-layer agent-instruction authority chain with operator rulings outranking every document; never-invent-a-number and stop-and-list rules. Read budgets (tail -n 120, top 80 lines, grep the row); generated files never hand-edited. Every one of 39 project folders resumable from one progress file.
- **Model routing and token economics.** M1 fast tier by default; M2 frontier tier only on eight named triggers, then back. M3 top tier capped to orchestration and  $\leq 15$  verification reads, zero writes; M4 never the smallest tier. M5 no promotion for comfort; M6 escalations carry a one-page findings map. A session-cost auditor prices every transcript against published per-million rates with cache-read and cache-write multipliers.
- **Evaluation and verification.** Pre-registered test cells (3.6M setups, 75 stocks, negative result recorded, not re-cut); a 9,014-line pre-deploy assertion script. Post-deploy probes with auto-rollback; Playwright E2E matrix E01 to E20. A segregated verified-outputs store; counts valid only if measured in the current session.
- **Safety and governance as code.** SEBI vocabulary, data-recency and disclosure firewalls enforced as build-gate assertions (rejected 7 of 7 draft captions; the 30-day floor later matched the regulator's July 2026 circular); secret names documented, values never written. Agents never deploy, set DNS or rotate secrets; single-writer files. A recorded anti-pattern register including the one production incident caused by two agents on one file.
- **Orchestration patterns.** Director/producer split; fan-out research with JSON schemas. Phase-gated release with a file-ownership map, a phase closing only on pasted command output; a 48-hour re-plan protocol. A parking list so nothing enters the runbook mid-flight.

## TOOLS AND STACK

- Agentic coding tools (agents, skills, hooks, workflows, 22 plugins); MCP connectors (Chrome, Gmail, Drive, Wix, Telegram, Discord, iMessage, Bright Data) | Python (pandas, numpy), JavaScript, Pine Script v6, HTML/CSS | Cloudflare Pages, Workers, KV, Durable Objects, R2 | Razorpay, PayPal | Google OAuth | Supabase | Brevo, ZeptoMail | Wix | GitHub | launchd | Playwright | ElevenLabs | ffmpeg | TradingView

## IN PROGRESS, SEPTEMBER 2026

- An API-based LLM application with tool use and a 40-item eval set built from the compliance firewall; an MCP server exposing the F&O calculator as tools. A public reference architecture of the Desk. Not yet done, and said so.