Anubhav Sood

Founder-operator who built and runs a company on AI agents: 90 agents, two live payment platforms, evaluation and governance enforced as code, a marketing organisation and two video studios, across a 145,000-file, ~1.9M-line estate.

Bengaluru, India | anubhav.sood.1701@gmail.com | linkedin.com/in/anubhavsood17 | efficientearners.com | mylittlelore.com

Prepared for: OpenAI India, GTM Enablement Lead (Mumbai) · Applied AI Architect (Delhi/Mumbai/Bangalore) · Strategy & Operations Lead, India (Bangalore). All counts in this document were measured or last verified on 2026-09-07 unless a different date is given next to the number. Every figure carries its source; nothing below is invented.

How to read this document. It is long by design. The operator's brief was "do not hold anything back." Section 1 is the two-page version; everything after it is evidence. Source files are abbreviated: INV = *02_Evidence/00_MASTER_INVENTORY.md*, EV05 = *05_evidence_AI_engineering.md*, EV06 = *06_evidence_GTM.md*, EV07 = *07_evidence_ventures_and_links.md*, CV = *CV_OpenAI_Founder.md*, FIT = *FIT_ANALYSIS.md*. Full paths are in Section 14. The underlying engineering evidence was gathered for a frontier-lab application in parallel. The AI-agent system it documents is provider-agnostic: every fact below is reported using neutral vocabulary (agentic coding tools, frontier-tier / fast-tier models, agent instruction files) rather than any one vendor's product names.

1. Executive summary

I am a computer science engineer by training, a developer before my MBA, and the solo founder who has spent fifteen months building and running a real company with AI, on agentic coding tools since March 2026, (Efficient Earners, a trading-education and market-structure business) with agentic coding tools as the entire engineering, marketing and operations department. I designed and now operate 90 AI agents and 25 skills under a written model-routing law, one owner per file, release gates closed by two read-only veto agents, and a session-cost auditor, across a 145,841-file estate of roughly 1.9 million lines of code (EV05 §A, §C). On top of that estate sit two live payment platforms, an NSE market-data engine emitting 15,000+ Pine Script files, a 12-agent marketing organisation, two multi-agent video studios, and a written course-production pedagogy that shipped a 72-chapter video course. I designed a regulatory vocabulary and data-recency firewall (the SEBI compliance perimeter) as build-gate assertions, not a style guide, and it caught 7 of 7 non-compliant captions in one real pass before they shipped (EV06 §A, §B).

Before this I carried an enterprise quota for five years at Ruckus Networks (CommScope) and HFCL, selling AI-driven analytics and cloud subscriptions to engineering teams and closing at the CIO/CFO level, and before that led small teams as a developer. That combination (quota-carrying seller, hands-on architect, and the buyer of my own AI tooling for fifteen months) is the case this document makes. The honest counterpart to it: real revenue is ₹500 to date, the constraint I diagnosed is audience and distribution, not product or engineering, and I say so in Section 13 rather than only in an appendix.

By the numbers

Metric	Value	Source
Files under the working estate (excl. node_modules/.git)	145,841 (up from 111,358 one month earlier, ~30% growth)	EV05 §A, §C
Lines of code: Python / JavaScript	512,421 across 2,046 files / 1,419,620 across 9,990 files (≈1.93M combined)	EV05 §C
Pine Script v6 files / HTML files / Markdown docs	15,117 / 3,416 / 6,472	EV05 §C
Agents: total unique / by tier (project-local agent-instruction frontmatter)	90 unique (86 project-local: 53 ee–/14 es–/17 `tl–`/2 vp–, + 17 global – 8 dual-scope) · 49 frontier-tier / 33 fast-tier / 4 inherit / 0 smallest-tier	INV §4 rows 47 to 49; EV05 §B.1, §C
Skills: project-scoped / global	21 / 4 (25 total)	EV05 §B.2, §C
Git commits: Desk platform / public website	28 (2026-07-30→08-29) / 200 (2026-05-30→09-06)	EV05 §C
Browser self-tests: calculator / workstation	314 / 435 (749 total)	EV05 §A.1, §C
Pre-registered backtest scale: setups / entries / stocks / cells passed	3.6M / 579k / 75 / 0 of 8 (negative result recorded, not suppressed)	EV05 §A.6, §D
Backtest replay coverage	780 stock-weeks (298 benched, 128 gate-clean)	EV05 §C
Founder hours reconciled in the costing system, Jun 2025 to Jul 2026	2,581 hours (₹4.14L cash, ₹29.95L at an internal ₹1,000/hr rate)	EV05 §A.7
AI-tooling spend trajectory (tracked as a first-class cost line)	My primary AI-agent subscription escalated across three tiers: ₹2,399/mo (Oct 2025) → ₹8,900/mo (Mar 2026) → ₹17,800/mo (unconfirmed month, mid-2026)	EV05 §A.7
Products: locked catalogue / wider portfolio audited	4 live SKUs / 12 products at various build stages	INV §2; EV06 §A
Courses: built & QA-passed / in production	"The Read" (72 chapters, 3h16m) built and QA-passed; Blocks 10K (240 sections) built, unsigned; two more scripted/blueprinted	EV06 §E; EV07 §2a

Metric	Value	Source
Books: title / editions	"Trade Less": 150pp Global + 152pp India Edition + audiobook script (3h25m)	EV06 §E
Decisions: dated founder rulings	39 per the founder's own decision log; 35 distinct items independently traceable across the three evidence files (see Section 9)	CV; INV §10, §13 row 12
Verified real revenue, all products, all time	₹500, one UPI payment, captured 28 Aug 2026	EV06 §C

2. Capability matrix: three OpenAI seats

Rows are the requirements as published for OpenAI India's GTM Enablement Lead (Mumbai), Applied AI Architect (Delhi/Mumbai/Bangalore) and Strategy & Operations Lead, India (Bangalore). Strength is judged conservatively: a gap is named as a gap, not talked around.

GTM Enablement Lead (Mumbai)

Requirement	Evidence	Artefact	Strength	What closes the gap
Build enablement programs from scratch for a rapidly scaling GTM team	Designed and operates a 12-agent marketing organisation from zero, plus an 8-document B2B enablement pack, plus a course-production pipeline with human gates	Sections 3, 7, 8	Strong	None
Build, manage and automate GTM collateral	A 12-platform, three-tier distribution map with automated posting rails and a no-silent-skip rule; a single-file operations dashboard reassembled from modular source	Section 8	Strong	None
Manage a regional training calendar; run trainings for product launches and process rollouts	Partner-enablement track record at Ruckus (certification tiering, a Rapid Demo programme, Lunch & Learn series); no formal India-regional training-calendar ownership yet	CV	Partial	Direct precedent exists; scale and cadence would be new
Technical depth + strategic thinking; understanding of the AI landscape	Personally architects a 90-agent, evaluation-gated production system; writes a model-routing law and a token-cost law	Sections 3, 5	Strong	None
Thrive in ambiguity; ruthless prioritisation; high-horsepower context-switching	39 dated founder rulings in a half-year; a defined slippage law and a 48-hour re-plan protocol; retired a headline revenue target in writing rather than let it drift	Section 9	Strong	None

Applied AI Architect (Delhi / Mumbai / Bangalore)

Requirement	Evidence	Artefact	Strength	What closes the gap
Senior technical owner for a customer portfolio; "CTO of your book of business"	Has personally been the technical owner (architect, operator and buyer) of my own AI-agent estate for fifteen months, the same posture the seat asks a candidate to bring to a customer's portfolio	Sections 3, 4	Strong	None
Own the technical account plan; discovery with executives and technical teams	Ran first-level technical discovery and prepared bills of materials for 5 years at Ruckus/HFCL before making the commercial case to CIO/CFO	CV	Strong	None
Develop Applied AI Architectures spanning models, data, integration, security, governance, evaluation, deployment	Cloudflare Pages Functions/KV/Durable Objects, Google OAuth, Razorpay, Supabase, an assert/probe/E2E gate chain, a segregated verified-outputs store	Sections 4, 5	Strong	None
Guide customers through technical evaluations, workshops, POCs	Pre-registered backtest cells as a POC-evaluation method; a 14-agent, 25-aspect audit methodology proven on a real commerce launch	Sections 5, 10	Strong	None

Requirement	Evidence	Artefact	Strength	What closes the gap
Familiarity with common LLM frameworks and tools	512,421 LOC of Python in the estate, directed by me; no LangChain-style framework use; my agentic coding harness itself is the framework	Section 11	Partial	An MCP server (in progress) would demonstrate framework-level tool-building directly

Strategy & Operations Lead, India (Bangalore)

Requirement	Evidence	Artefact	Strength	What closes the gap
Establish KPIs, reporting frameworks and operating cadences	A Friday scorecard, a gate calendar (G1 to G6), a biweekly Platform Pulse verdict pass, and a costing OS reconciling founder hours against spend	Sections 8, 9	Strong	None
Build scalable processes across business lines; identify operational gaps	39 folders each self-resumable from one progress file; a severity-laddered gap register (P0 to P3) synthesised across six independent mapper passes	Sections 3, 5	Strong	None
Support market entry strategies and operational readiness	Opened HFCL's South India enterprise business from zero (no accounts, no partners, no references), plus a from-scratch payment-rail rebuild for a new market segment	CV; Section 10	Strong	None
Lead and mentor a regional operations team	Sole operator of the estate; no direct-report people-management history in this evidence base	None	Gap	Named honestly: the agent "team" I manage is not a people-management analogue
Comfort operating in ambiguity; high integrity, low ego	A written decision ledger of 35+ dated rulings, including one that retired the company's own headline revenue target in writing	Section 9	Strong	None

3. The agent operating system

Efficient Earners is not "a founder who uses AI." It is a company whose org chart is a set of agent instruction files, whose management system is a router document, and whose QA department is two read-only agents that hold no other job. This section is the org chart and the rulebook.

Org chart by lane

Lane	Named agents	Tier mix	Source
Command	`ee-ceo-mentor`, `ee-north-star-keeper`, `wealth-cfo`	not individually stated	EV05 §B.1
Engine	`ee-engines-master`, `ee-scenario-quant`, `ee-scenario-backtester`	not individually stated	EV05 §B.1
Terminal / Blocks	`ee-blocks-daily-ops` (05:30 IST daily fail-closed publish)	not stated	EV07 §2b
Marketing (12 agents)	`ee-cmo`, `ee-brand-strategist`, `ee-marketing-director/-producer`, `ee-distribution-director`, `ee-seo-director`, `ee-intel-director`, `ee-platform-pulse/-harvester`, `ee-youtube-director/-producer`, `ee-course-studio`	frontier-tier judgment / fast-tier mechanics split	EV06 §B; INV Reconciliation #2
Growth-Ops (5 agents)	`ee-b2b-desk`, `ee-cohort-director`, `ee-compliance-counsel`, `ee-member-success`, `ee-revenue-ops`	mixed tier	EV05 §B.1; EV06 §B
Desk execution pack (11 agents)	`ee-desk-director/floor/counter/gatekeeper/money/surfaces`, `ee-intake-director`, `ee-marketplace-engineer`, `ee-fleet-engineer`, `ee-release-engineer`, plus the two veto agents below	not individually stated	EV05 §B.1
Veto pair (read-only)	`ee-red-team` (actively tries to refute the phase's own claims), `ee-compliance-auditor` (rule conformance); both restricted to read-only tools, no write access	None	EV05 §B.3, §B.6
Board B trading research (4+ named)	`ee-trading-mentor`, `ee-trade-lab`, `ee-strategy-quant`, `ee-trigger-researcher` (list marked incomplete in source)	not stated	EV05 §B.1
Media studios (33 agents)	`es-*` cinematic-shorts studio (14 agents), `tl-*` long-form video-film studio (17 agents), `vp-*` course-film line (2 agents)	None	EV05 §B.1

86 distinct project-local agent files (53 ee-, 14 es-, 17 `tl-`, 2 vp-) plus 17 global agents, minus 8 identical dual-scope duplicates = 90 unique agents by the estate's own generated count. Model-tier split across the 86 project agents' own instruction-file frontmatter: 49 frontier-tier, 33 fast-tier, 4 inherit, 0 on the smallest tier. (EV05 §B.1, §C; INV §4)

The 29-row router and Model Law M1 to M6

A single document maps task type → skill → agent → model tier in 29 rows, governed by six rules (EV05 §B.3):

- **M1:** the fast-tier model is the default for all file work.
- **M2:** the frontier-tier model fires only on eight named triggers (entry/session/entitlement/allowlist logic; money movement; race/persistence/idempotency root-cause; progress data-model grammar; CSP/routing architecture; the two read-only veto passes; pricing authorship + regulatory file; architecture memos) and must "analyse → decide → return to the fast tier."
- **M3:** the top orchestration tier is capped to a curated ≤100KB packet, ≤15 targeted verification reads, and zero exploration or writes.

- **M4:** never the smallest tier, "just a grep" included.
- **M5:** no promotion for comfort: a risky task gets a second fast-tier pass or a read-only frontier-tier veto, never a frontier-tier implementer.
- **M6:** escalations carry a ≤1-page findings map; the top tier never re-reads what the fast tier already read.

An inferred (not measured) token-share target sits alongside the law: normal regime fast-tier≥85%/frontier-tier≤12%/top-tier≤3%; bootstrap 60 to 70%/15 to 25%/10 to 15%; gate-day 70%/30%. (EV05 §B.3)

DISPATCH FLOW: TASK TO MODEL TIER, WITH THE VETO LOOP

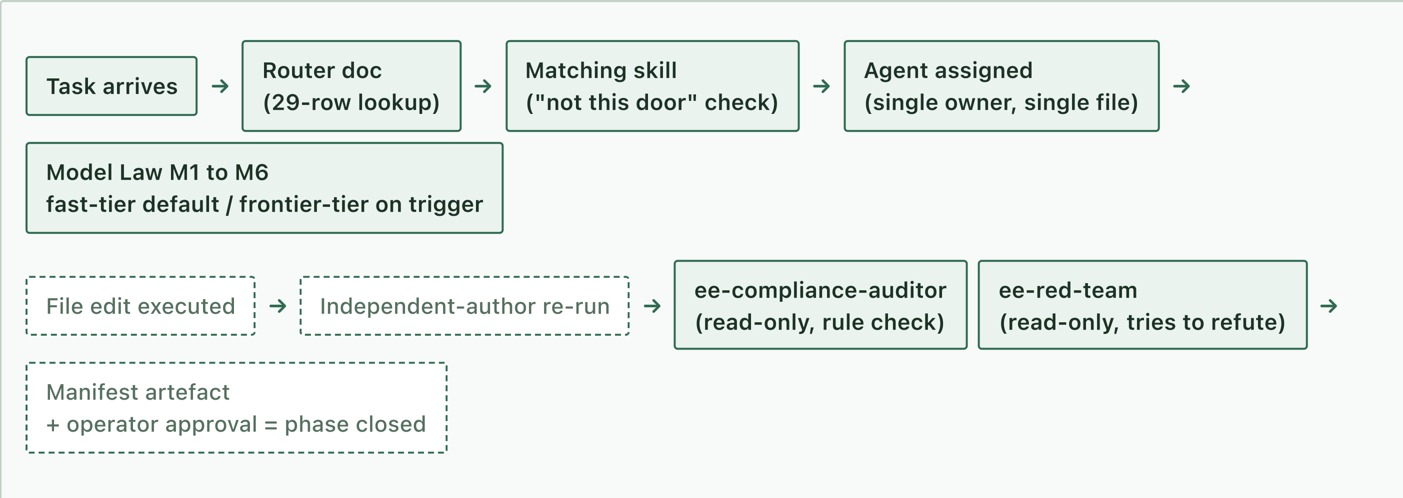

Single-writer and veto pattern

Every gate or money-touching file has exactly one owning agent; a documented ownership map states plainly that "two agents never share a file" (EV05 §B.3). This is not a theoretical rule: an anti-pattern register records a real production incident it exists to prevent: two agents editing one file in one wave caused a port regression on 2026-08-20 (EV05 §B.3). Every gate/money change is closed by an independent-author re-run plus a read-only pass from both veto agents before a release phase counts as done: "a phase with only agent sign-off is not closed" (EV05 §B.3).

Fleet-learning loop

A dedicated skill is, by its own description, "the ONLY sanctioned path from a run-time lesson to a definition or skill edit; agents themselves never hold that pen," run once daily (EV05 §B.3). This is a governed feedback loop from execution transcripts back into agent and skill definitions, distinct from letting any agent freelance-edit its own instructions.

The 10→3 research workflow

A custom workflow script fans out 10 fast-tier researchers (each briefed on a distinct track with a required structured JSON schema of 8 to 15 sourced claims) into 3 frontier-tier synthesisers that write draft operating-skill documents from the aggregate. External research findings are treated explicitly as "inputs, never overrides" against the business's own house law (EV05 §B.3).

Hooks and launch configuration

One project-level hook, triggered before every agent dispatch, prints a colour-glyph banner identifying which of ~50+ named agents is being dispatched, and is explicitly designed to never block the dispatch (EV05 §B.4). A local launch configuration defines nine dev-server configurations (the website, retail and B2B terminals, an agent-guide server, an HTML hub, a weekly-ops server, the F&O calculator, an MF-BuyZone app, and a Command-Center static server), each on a distinct port (EV05 §B.4).

Anti-pattern register

Pattern	What happened	Fix	Source
Two agents, one file, one wave	Port regression, 2026-08-20	Single-writer-per-file rule + veto close before a phase counts as done	EV05 §B.3
Session-scope collision	Two accounts working the same estate produced same-day collisions; router self-reports "violated daily (61 recorded roots)"	File-modification-time comparison before editing a single-writer doc; append to a session log on stop (target, not fully achieved)	EV05 §B.3
OS-rebuilding as productive procrastination	Four command-center generations in five months, diagnosed by the founder himself	Tripwire R1 (2026-07-27): no further OS-build hours until a ₹1,00,000 revenue month is banked, bug fixes excepted	EV06 §D

4. Platform architecture walk-throughs

4.1 The Desk (ee–app–`platform`): the paid member commerce platform

Cloudflare Pages Functions, a KV namespace (GATE_KV) and a built-but-unbound Durable Object credits ledger, Google OAuth, ZeptoMail for transactional mail, and Razorpay for payments and webhooks. 67 JS files / 18,969 lines, 63 HTML files, a 9,014-line pre-deploy assertion script, 28 commits (2026-07-30→08-29) (EV05 §A.2, §C).

GATE ORDER AND WEBHOOK CHAIN

Purchase click → Razorpay checkout → Webhook received → Signature verify → Idempotency check → Write own record → Re-fetch & corroborate against provider → Entitlement granted (read-back confirmed)

The credit-balance model is not a stored number but a fold over an append-only ledger, never stored directly; a single writer owns entitlement state (EV05 §E). Before any deploy: a static pre-deploy assertion gate + a live-HTTP post-deploy probe with auto-rollback + a Playwright end-to-end harness (tests E01 to E20) (EV05 §A.2). KV/DO/OAuth secrets are never printed or written to a tracked file; agents never run a production deploy, set DNS, or rotate secrets. Deployment stays operator-only (EV05 §D).

4.2 mylittlelore.com: a second live storefront, my own build

Static HTML/CSS/JS, no build step, deployed on Cloudflare Pages under my own GitHub organisation (repo private) (EV07 §1b). A Cloudflare Worker drives PayPal checkout, an R2 bucket issues signed 24-hour download links, and a KV namespace is the customer ledger.

PURCHASE → DELIVERY FLOW

22 bundles, $7.99 ea. (3rd free) → PayPal checkout worker → KV ledger → R2 signed link (24h expiry) → Buyer download

Live and taking real payments since 2026-08-14. A 14-agent, 25-aspect same-day audit (links, catalogue, SEO/GEO, performance, accessibility, mobile, security, legal, copy honesty, checkout code, hosting, design system) found and fixed the actual design-breaking bug, diagnosed and patched a stranded-buyer payment bug, and fixed ~60 defects across 12 commits; both site and worker redeployed the same evening (EV07 §1b).

4.3 The market-data engine ("EE-Engines")

Python (an NSE data-ingestion library, pandas/numpy for the volume-profile math), a five-phase contract-driven pipeline (each phase with its own written contract), Pine Script v6 code generation, and scheduled automation (EV05 §A.1).

FIVE-PHASE PIPELINE, SHA-PINNED

Phase 0 fetch → Phase 1 extract → Phase 2 parent-blocks (sha-pinned) → Phase 3 render → Phase 4 pool-combine

87 symbol trees (75 stocks across 15 sector pools + NIFTY/BANKNIFTY + 10 derived-pool mirrors); 15,117 `.pine` files estate-wide; two browser apps carry 314 and 435 self-tests. A hard size/token/duplicate gate runs on every Pine build. Runs unattended daily (16:15 IST Mon to Fri) and weekly (Sat 09:00 IST) via a scheduled-job system. A real macOS constraint (background jobs blocked from the working volume) was engineered around with an internal-disk launcher relaying to the terminal, arrived at after proving the fix with four probe agents (EV05 §A.1, §D).

4.4 Blocks: the market-structure tool

A five-horizon (Yearly/Half-yearly/Quarterly/Monthly/Weekly) market-structure charting surface over 212 NSE instruments across 17 sectors, EOD-only by design (EV07 §2b).

05:30 IST DAILY FAIL-CLOSED PUBLISH

Incremental candle fetch → Build script (must end "Leak scan PASSED") → Gate check → PASS → publish

FAIL → previous tree stays live, nothing partial ships

Operated by a dedicated agent, scheduled daily, with a run log appended every session (EV07 §2b).

4.5 The website (`ee-website-pages`)

Static site, published by Cloudflare Pages, DNS fronted by Wix. This is a deliberate split-hosting decision, not an accident. 200 commits, 2026-05-30 to 2026-09-06 (EV05 §A.3, §C). A hard rule reserves any push/merge to the production branch as operator-only, since that push is a live deploy via Cloudflare Pages' GitHub integration (EV05 §D).

4.6 The costing OS

A single self-contained HTML business-books application (68KB, browser-local storage, no backend, no login) plus a Python reproducibility script (19KB), replacing a spreadsheet with Earnings/Spend/Activity ledgers, P&L by month/quarter/FY, and a trading rollup (EV05 §A.7). It reconciled 2,581 founder hours and ₹4.14L cash (₹29.95L at an internal ₹1,000/hr rate) across Jun 2025 to Jul 2026 from brokerage exports and an actual AI-tooling invoice, and tracks my own AI-agent subscription spend as a first-class, per-project allocated cost line (EV05 §A.7).

4.7 The personal PWA (SDT)

A hand-built installable Progressive Web App ("Survival Dominance Tracker"): a 1,979-line single-file front end, a manifest declaring standalone display and two icon sizes, and a service worker for offline caching. No backend, vanilla HTML/JS/PWA APIs only (EV05 §A.10).

5. Evaluation, verification and cost engineering

Assert / probe / E2E

Every Desk deploy passes through three independent checks before it counts as done: a 9,014-line static pre-deploy assertion gate, a live-HTTP post-deploy probe with auto-rollback, and a Playwright end-to-end harness with tests named E01 to E20 (EV05 §A.2). "Definition of done" is stated explicitly: build script + assertion-gate PASS at or above the session floor + the relevant end-to-end suite + a manifest append (EV05 §A.2).

Self-tests and verified outputs

Two browser trading apps carry 314 and 435 self-tests respectively. A dedicated top-level folder (governed by its own instruction file) segregates machine-verified output (weekly/monthly/quarterly/backtest-practice tiers) from work-in-progress (EV05 §B.6).

Pre-registered backtests

A pre-registered study committed 8 test cells before running 3.6 million setups across 75 stocks and 579,000 entries, and reported the negative result (0 of 8 passed) rather than re-cutting the study to find a positive one (EV05 §A.6, §D). A separate 78-week backtest replay benched 298 of 780 stock-weeks, with 128 flagged gate-clean (EV05 §C).

A self-written session-cost auditor and the costing system

A self-written session-cost auditor walks the main session transcript plus every background-task and sub-agent-workflow transcript, sums each message's usage block, and prices the total against a hard-coded per-million-token rate table across four model tiers, with separate cache-read (0.1×) and cache-write (2.0× for 1-hour, 1.25× for 5-minute) multipliers, explicitly noting the figures are "burn rate against usage limits, not an invoice" on my subscription plan (EV05 §B.4). This sits alongside the costing OS (Section 4.6) that reconciles the same spend against founder hours.

Measured-not-quoted law

"A gate count is valid ONLY if produced in THIS session; documents quote counts with a timestamp and the word 'measured'" (EV05 §B.6). This document follows the same rule: every count above carries the date it was taken.

6. Safety and governance as code

Because I am an *unregistered* market educator in India (no SEBI RA/RIA), the entire product surface is built under a hard, code-and-copy-enforced vocabulary firewall: the single most load-bearing constraint on the business, and the strongest safety-and-governance-adjacent evidence in the estate (EV05 §D; EV06 §A).

The firewall, as three enforced layers

Layer	What it bans / requires	Enforcement	Source
Vocabulary firewall	Bans buy/sell/target/stop-loss/signal/recommendation/advisory/"indicator"-as-product-name/returns/win-rate/accuracy/any P&L figure, on every surface, including internal reports	Build-gate assertion, not a style guide	EV05 §D; EV06 §A
Data-recency envelope	No named-security price data inside the trailing period	Tightened 1 day → 3 months → a uniform 30-day floor, effective 1 Jul 2026, matching SEBI's own circular almost word for word	EV06 §A
Disclosure firewall	Outputs (zones, tiers) may be shown; the construction math is never explained; SEBI-adjacent lineage is "speakable, construction never"	Enforced in copy review and agent prompts	EV06 §A

The 7-of-7 rejection

At the 28 Jul 2026 "CAS launch," a compliance-firewall pass rejected 7 of 7 draft Instagram captions in one campaign: a fabricated broker-email claim, an invented statistic, an order-priority rule that read as advice, and a regulator hashtag on three captions (EV06 §A, §B). This is presented in full as Case Study 2 (Section 10).

The 30-day floor and the regulatory precedent it anticipates

The internal 30-day recency floor converged, almost word for word, with SEBI's own July 2026 circular. The internal audit calls it "the regulatory floor with zero headroom," not a safety margin any more (EV06 §A). My own research also surfaced the tail-risk case the firewall exists to prevent: a real regulatory action against a finfluencer for stock-specific recommendations inside "education"-branded content, with a nine-figure disgorgement ordered (EV06 §A).

Secret handling and operator-only deploys

Secret *names* are documented; secret *values* are never printed or written to a tracked file. Two files that are almost certainly credentials were located and never opened during either evidence pass. Agents never deploy, set DNS, bind data stores, set secrets, or push a production release. That line is drawn by me, not a tooling limitation (EV05 §D).

PII rules

A member-email storage rule was deliberately reversed, from "never store member email" to encrypted storage, as part of a dated 2026-08-24→27 ruling batch, each captured as a verbatim-quoted operator decision with its own stated rationale (EV06 §D).

The client-work IP firewall

Signed 2026-07-27, this bright-line rule structurally separates any freelance/agency client build from Efficient Earners' own IP: zero reuse of EE source, logic, parameters or derived output; no EE strategy or performance claims in a client pitch. Never build or operate a distributed signal bot, auto-execution system, or advisory output for a third party; every client engagement carries a written scope, an IP clause, and a GST invoice. An explicit "no similar-but-different" decline rule closes the gap a looser rule would leave open (EV06 §F).

7. Media and course production

Two video studios

Studio	Agents	Stage chain	Source
Cinematic shorts (`es–*`)	14	Cinematic-video studio for shorts; script gates, scene maps, voice-over generation, film QA before render	EV05 §B.1; CV
Long-form video-film (`tl–*`)	17	Long-form video-film studio: Hindi/Devanagari subtitle burn-in, 9:16 reel cuts, film QA before render	EV05 §B.1; INV Reconciliation #3

A separate two-agent line (vp–*) builds course-film content (EV05 §B.1).

Course studio pedagogy

A written instructional-design method, locked as law rather than guidance (EV06 §E):

- **Backward design:** every module starts from one observable end-capability statement, decomposed into one performance objective per video, each mapped to a checkable self-test item.
- **Prerequisite DAG** with zero forward references: videos are topologically sorted; a video may use only terms taught in strictly earlier videos.
- **Spiral schedule:** every core concept resurfaces in at least two later modules.
- **Chain plan:** 5 to 9 videos per module, a hard 2 to 3 minute cap (4-minute ceiling), chained into sequences.
- **Two human gates:** Gate L (learning map) and Gate M (beat-level script), before any video production begins.

The 72-chapter course

"The Read": 72 chapters, 5 modules, ≈3h16m runtime, built and QA-passed, not yet deployed pending a hosting decision (EV06 §E; INV §2 row 4).

Blocks 10K: 240 sections, 21 verifier passes, 240/240 scan

A ₹10,000 course teaching hand-built block reading. The density-scanning/scoring engine math is never taught (an absolute, repeated IP firewall). 4 modules, ≈240 sections, ≈84,000 spoken words, ≈10 hours of AI-voiced English content plus 2 founder-recorded Hindi/Hinglish workshop hours. Production: 4 frontier-tier module leads producing blueprints, batched fast-tier drafting fleets, **21 adversarial fast-tier verifier passes** (38 fixes applied), full frontier-tier reads of all 121 firewall-sensitive sections (40 more fixes), a deterministic banned-substring scan (**240 of 240 sections clean**) and a personal founder read of the two highest-risk sections. Estimated build cost ≈$40 to 55. Gate A signed 2026-07-12; Gate C (full draft) delivered 2026-07-18; unsigned/not launched as of 2026-08-15 (EV07 §2a).

The book and editions

"Trade Less": a 150-page trading-psychology book published across 5 external channels (Kindle, paperback, Apple Books, Gumroad, Wix Store), plus a fully-localized 152-page India Edition (NIFTY/MCX-specific), plus a written audiobook script (~3h25m narration + companion PDF), ready for narration (EV06 §E).

Long-form flagship engine

Governs 7 to 15 minute explainer scripts under a founder-set editorial motto written into the production law: "Simple, less information, but precise information." One idea per scene, every claim carries one specific number. Two operator gates: Gate R (draft research report) and Gate S (script + scene map). Every script routes through a mandatory SEBI compliance scan before production (EV06 §E).

Hindi subtitle pipeline

Devanagari subtitle burn-in and 9:16 reel cuts exist as a documented capability inside a dedicated video-production skill (INV Reconciliation #3).

8. GTM and marketing system

Twelve named agents (reconciled from an earlier "fifteen" figure; see Section 14 methodology note) run under one front door, a frontier-tier routing agent, which routes, sequences and reports but never writes content, ships, or approves a gate itself (EV06 §B; INV Reconciliation #2).

The 12-platform, three-tier distribution map

Tier	Mechanism	Platforms	Source
Tier 1	Bot API, fully automated	Telegram	EV06 §B
Tier 2	Free-tier social API, deliberately LinkedIn-only to protect a 20-post/mo cap	LinkedIn	EV06 §B
Tier 3	Paste-packet, keyless, always-ready; plus browser-posting via a Chrome connector	Substack, Medium, WhatsApp, Quora, X, Threads, Instagram, Pinterest	EV06 §B

A hard rule: a missing key **never silently skips a platform**. It falls back to a manual paste packet (EV06 §B). A house UTM convention stamps every outbound link with source/medium/campaign/content, and caught one real pre-convention link shipping bare and unstamped during a dry-run rehearsal (EV06 §B).

Platform Pulse: the biweekly browser-driven circuit

Every alternate Friday, a fast-tier "harvester" agent drives my real logged-in browser through a fixed **13-rail circuit** (video-platform studio dashboards, a social-business dashboard, podcast dashboards, a self-publishing dashboard, an e-book platform, audio editions, two commerce platforms, a storefront, both websites, a newsletter platform, mail), collecting numbers into a schema'd raw-data file, strictly read-only on every platform (EV06 §B). A frontier-tier agent then runs one capped verdict pass: a ≤2-page report, a dashboard-readings merge, and a modification docket routed by finding-type to the owning lane. The pulse recommends; it never edits another lane's files. Token discipline is designed in: text-first reads, ~6 tool calls per platform, 2 retries then log-and-move-on (EV06 §B).

Lifecycle email and the dashboard

Capture → welcome → weekly → renewal machinery on an email-marketing platform, with renewal-nudge drafts scheduled at T-7/-3/0/+3, drafted by the agent, always sent personally by me, no auto-renewal, ever (EV06 §B). A single-file HTML dashboard is "the single surface I look at" across 4 brands and ~13 distribution rails, designed so a missing input lowers a confidence score rather than flattering a vanity number (EV06 §B).

Brand-voice lock and cadence

A locked brand-voice document (locked 2026-05-01) defines two separate voice references (long-form vs. marketing surfaces) as 8 attributes, each with a worked wrong/right example, stated as absolute: "If a long-form sentence fails any one of these, it gets rewritten. No exceptions" (EV06 §B). Standing rhythm: daily AM/PM production; Monday week plan; Friday scorecard; Saturday 20:00 IST member delivery ("a streak that outranks every other output, forever"); Sunday structural read; alternate-Friday pulse. Content mix follows a fixed three-stream floor (MARKET / BLOCKS / ROAD) per cycle (EV06 §B).

9. Commercial and decision system

Pricing iterations and the falsifying test

Book pricing: $15 to 20 target → $19 locked → $9.99 locked/$15 ceiling (D23, 2026-07-22) → shipped at $12.99 Kindle/$19.99 paperback, a documented deliberate deviation to capture the 70% royalty band (EV06 §C). Course pricing used an explicit falsifying-test design, not intuition: sell 10 units at ₹2,999 to the warm list, with a pre-declared decision rule (≥10 = validated; 3 to 9 = price is fine, reach is the problem, do not touch it; 0 to 2 with price cited = retest once at ₹1,999, never lower) (EV06 §C). Desk credits went through three pricing-model iterations in five weeks: credits → a signed 1-credit=₹1 anchor (R-09) → retired and converted to a 10%-flat, no-expiry cashback wallet modelled on a well-known reward-points pattern (R-10), specifically to reduce a compliance exposure a prior decision had flagged (EV06 §C).

Rails tried, in order

Wix Pricing Plans/Payhip (ruled out: Payhip's recurring billing requires Stripe, and Stripe India is invite-only) → Razorpay Payment Links selling closed-loop credits → direct rupee purchases via Razorpay, entitlement opened only after independent corroboration against Razorpay's own record (EV06 §C). This rebuild is Case Study 1 (Section 10).

Ledgers and unit economics

A revenue ledger and a member roster are header-only, structured schemas (GST split, cashback accrual, rail, recurring-vs-one-off, status) governed by a stated law: "numbers are recorded, never invented; null is null" (EV06 §C). A refund-float register books dated workshop seats as a liability, not revenue, until delivered (D-038). Razorpay's own published rate (2% + 18% GST = 2.36% effective) nets ₹2,500 → ₹2,441.00, a 97.64% gross margin per SKU, verified directly at razorpay.com the same session. Client-build unit economics: ₹55,000 gross → ₹46,610 GST-net, 50% advance / 50% on delivery (EV06 §C).

The ladder, the odds, and the gates

A signed decision on 2026-07-27 explicitly retired the ₹1 Cr Jul to Dec cumulative target in writing and replaced it with a monthly run-rate ladder: ₹10,00,000/month from December 2026, with August and September individually targeted at ₹5,00,000 GST-net each. This was recorded alongside its own honest gap (modelled August base case: ₹1.20L accrual / ₹1.01L cash) and calibrated odds signed at the same sitting: Aug 5%, Sep 15%, Dec 10%, scheduled to be scored, not silently revised, at an Oct-1 quarterly review (EV06 §C, §D). This is Case Study 4.

Gate (Aug 31 2026)	Criterion
G1	Rails live
G2	≥1 client build closed with cash
G3	Opt-in list ≥150
G4	≥3 paid SKUs and ≥$100 own sales
G5	August cash ≥₹2.24L survival floor
G6	Mandatory (not pass/fail): honest variance write-up of actual vs. ₹5L target: "no silent drift, no retro-lowering"

Rule: missing ≥2 of G1 to G5 auto-simplifies September to two lanes only. (EV06 §D)

The decision ledger

Item	What it decided	Outcome
D23 (2026-07-22)	Book price locked $9.99, $15 ceiling	Shipped instead at $12.99/$19.99, documented deviation
D-027 (2026-07-04)	Terminal-only distribution; trading contribution reported, never a target	Signed
D-028	No NSE-derived data reaches a paying user until a scoped legal opinion lands	Killed 4 products outright
D-030 (2026-07-04)	Proprietary code never leaves the house: no Pine-script delivery, ever	Signed; obsoleted the prior pool-factory build
D-035 (2026-07-27)	Retires the ₹1 Cr target in writing	Replaced with the monthly ladder
D-037 (2026-07-27)	Five lanes defunded for Aug to Sep with named unfreeze triggers	Deferred, not killed
D-038	Delivered-only commerce law: no presales anywhere	Live
D-039 (2026-07-27)	Client-work IP firewall (Section 6)	Signed, live
R1 tripwire (2026-07-27)	No further OS-rebuild hours until a ₹1L revenue month is banked	Signed
R-01 to R-15	15 dated rulings across the Aug 2026 Desk rebuild: architecture, DNS, deploy authorization, credits reclassification, member-email storage	Signed
G-01	Desk canonical-hostname DNS provisioning	Open as of 2026-09-06
G1 gate (Aug 10 2026)	"Rails live": written criterion overtaken by the pivot to an owned Razorpay Desk	Contested; logged as an open, unreconciled contradiction rather than silently resolved
G1 to G6 (Aug 31 2026)	Re-based kill-gate, see table above	Scored per rule
2026-08-06 ruling	"Blocks is a TOOL": not a course, not an education product, not a service	Signed
2026-08-06 opportunity call	Split verdict on the global-market extension question	Case Study 5
2026-08-28 self-graded verdict	"Right enough to sell, wrong shape to scale: ship it, do not add to it," six named defects, the day payments went live	Recorded

The founder's own decision log names 39 dated rulings in total; this table lists the 35 distinct items independently traceable across the three evidence files read for this document. Several of the 39 are grouped here as single multi-item rows (e.g. R-01 to R-15). (INV §10, §13 row 12)

The 48-hour re-plan protocol

Invoked at least once on record: 2026-07-04, the founder proposed re-sequencing H1; after a 48-hour re-plan plus a 5-question opportunity protocol and a fix-first pass by three agents, it was signed "confirmed, log and release" (EV06 §D).

The constraint diagnosis

Stated once and never re-litigated: "the binding constraint is AUDIENCE, not product and not engineering; every rail reduces upstream to the same empty list" (EV06 §D, §B). A separate, independently-audited investor report reached the identical conclusion from a completely different evidence trail five weeks later, without citing the first finding (EV06 §D).

10. Five case studies

1. The Razorpay rail step-ladder rebuild

Context. The Desk's payment rail had gone through two prior implementations (Wix Pricing Plans/Payhip, then Razorpay-credits) and needed a from-scratch rebuild in August 2026 after the founder ordered "delete those codes... no patches; we go step by step from scratch" (EV06 §C).
What was built. A six-step verified ladder: account-prerequisite probe → credentials probe (self-disabling diagnostic) → ₹1 test link → checkout record + redirect road → webhook chain (signature verify → idempotency check → own record → re-fetch-and-corroborate → confirm amount → grant with read-back) → cashback/mail.
Mechanism. Each step was verified live by me before the next step could begin: no step N+1 without a real console read-back on step N.
Outcome. Payment acceptance opened 2026-08-28; the first real transaction (₹500, UPI) was captured the same day and preserved through a later account reset. One account-migration dead-end cost a full night and is recorded as a named trap for future rebuilds.
What it demonstrates. Careful, verified, incremental delivery on a real-money integration: the exact posture a customer needs to see before trusting an agent-driven build process with payment logic.
Source: EV06 §C; skill file `ee-razorpay-rail/SKILL.md`

2. The SEBI firewall catching 7 of 7

Context. At the 28 Jul 2026 "CAS launch," a batch of Instagram captions was drafted for a real campaign.
What was built/decided. The compliance vocabulary firewall, a build-gate assertion and not a style-guide suggestion, ran against the draft batch.
Mechanism. The firewall checks for buy/sell/target/stop-loss/signal/recommendation/advisory vocabulary, unverifiable statistics, and disclosure-tag misuse, on every surface including internal drafts.
Outcome. 7 of 7 captions were rejected: a fabricated broker-email claim, an invented statistic, an order-priority rule that read as advice, and a regulator hashtag on three captions.
What it demonstrates. A regulatory constraint enforced as a machine-checked invariant, not a policy document. It caught real non-compliant content before it reached a single follower, in a campaign that was already in flight.
Source: EV06 §A, §B

3. The My Little Lore 14-agent launch-day audit

Context. mylittlelore.com went live with real PayPal payments on 2026-08-14.
What was built. A same-day 360° site audit dispatching 14 parallel audit agents across 25+ aspects: links, catalogue, SEO/GEO, performance, accessibility, mobile, security, legal, copy honesty, checkout code, hosting, design system.
Mechanism. Each agent worked its aspect independently; findings were consolidated and fixed same-evening across 12 commits to both the site and the checkout worker.
Outcome. Found and fixed the actual design-breaking bug (an image-dimension CSS regression), diagnosed and patched a stranded-buyer payment bug where a buyer could pay and never reach their download, and fixed ~60 other defects, including a reviews section showing literal placeholder text and a FAQ describing a cart mechanic that did not exist.
What it demonstrates. Parallel multi-agent QA applied to a real commerce launch, on launch day, with real money already flowing.
Source: EV07 §1b

4. Retiring the ₹1 Cr target in writing

Context. The original headline plan target was ₹1 Cr cumulative revenue, Jul to Dec 2026.
What was decided. On 2026-07-27, a signed decision explicitly retired this number in writing. It was not silently dropped.
Mechanism. Replaced with a run-rate ladder (₹10,00,000/month from December, ₹5,00,000 GST-net individually for August and September), recorded alongside its own honest gap against the modelled base case, and calibrated odds (Aug 5%, Sep 15%, Dec 10%) signed at the same sitting, explicitly scheduled to be scored, not revised, at an Oct-1 quarterly review.
Outcome. A defensible, falsifiable target replacing an unmoored one, with a pre-committed scoring date.
What it demonstrates. Willingness to kill a headline number in writing rather than let it quietly slide. It is the same rigor the pre-registered backtest cells (Section 5) apply to trading hypotheses, applied here to the business plan itself.

Source: EV06 §C, §D

5. The 2026-08-06 "grab the architecture, kill the lane" opportunity call

Context. A live, real-time question: should the Blocks terminal extend to global markets?
What was decided. A formal 5-question opportunity protocol was run, and it returned a split verdict rather than a single yes/no.
Mechanism. Grab the market-agnostic architecture change (capped at ≤8 founder-hours: "only architecture decays") while killing-with-a-tripwire the broader global-market extension as a funded lane for the quarter, because it would consume exactly the hours needed to hit gates G2/G3.
Outcome. Stated explicitly: "it does not risk the Aug-31 gate, it fails it arithmetically." The opportunity was not rejected on vibes; it was rejected on an hours-budget calculation against a signed gate.
What it demonstrates. Structured, protocol-driven decision-making under a real opportunity-cost trade-off, with the reasoning written down rather than just the verdict.

Source: EV06 §D, 2026-08-06 entry

11. Tooling stack

Group	Tools
Agentic coding tools	An agentic coding harness (agents, skills, hooks, workflows), 22 installed plugins, a custom pre-dispatch hook, a local launch configuration (9 dev-server entries)
MCP / connectors	4 user-level MCP servers (Telegram, Discord, iMessage, a chat-test harness); app connectors in daily use: Chrome, Gmail, Google Drive, Wix, computer-use, Desktop Commander, Bright Data. Heavy MCP *operator*; no MCP server authored yet (Section 13)
Languages	Python (pandas, numpy), JavaScript, Pine Script v6, HTML/CSS
Cloud	Cloudflare Pages, Workers, KV, Durable Objects, R2; GitHub; scheduled-job automation
Payments	Razorpay (webhooks, Payment Links, Route payouts discussed for a future phase), PayPal
Auth / data	Google OAuth, Supabase (Postgres schema for the Terminal's auth layer)
Comms	Brevo (lifecycle email), ZeptoMail (transactional), Wix (DNS + Velo bridge)
Media	Voice-over generation (ElevenLabs), ffmpeg, TradingView (Pine Script v6 target)
Automation	Scheduled-job automation, a social-posting API integration, a Telegram bot API integration
QA	Playwright end-to-end tests (E01 to E20), a static assertion gate, a live-probe gate, 749 browser self-tests, the two read-only veto agents

12. How this transfers to OpenAI

GTM Enablement Lead (Mumbai)

1. **Stand up an India GTM-collateral system in the first 30 days** the way the 12-platform, three-tier distribution map was built: a working automation layer plus a documented manual fallback for every channel, so nothing silently breaks when a key expires.
2. **Build the regional training calendar around the same C0 to C7 gated production pipeline** used for course content: a commission stage, a design gate, a build stage, and a founder/manager approval gate before anything ships to reps.
3. **Run a biweekly enablement-effectiveness pulse**, modelled directly on the Platform Pulse: a fixed circuit of what to check, a capped verdict pass, and a docket routed to the team that owns the fix, not a vague "let's improve enablement" conversation.

Applied AI Architect (Delhi / Mumbai / Bangalore)

1. **Bring the assert/probe/end-to-end gate pattern into the first customer's Applied AI Architecture.** A static pre-deploy gate, a live post-deploy probe with auto-rollback, and a segregated verified-outputs store are a concrete, reusable reference, not a slide.
2. **Use the same evaluation rigor that ran the pre-registered backtest study** (a pre-committed test plan, a negative-result-recorded-not-suppressed norm) as the customer's own evaluation-framework template.
3. **Teach in two registers from day one.** The same rigor that produced a beginner-register 72-chapter course and a prop-desk-register B2B positioning doc from one underlying IP applies directly to briefing a customer's engineering team and their CIO in the same week.

Strategy & Operations Lead, India (Bangalore)

1. **Port the gate-and-scorecard rigor directly.** A Friday scorecard ordered lead-measures-first-lag-last, a signed gate calendar with pre-committed pass/fail criteria, and a mandatory honest-variance write-up even when a gate is not met, are already a working operating cadence, ready to be adapted to a regional team rather than a one-person company.
2. **Bring the severity-laddered gap-register method** (P0 critical → P1 journey-breaking → P2 hardening → P3 polish, synthesised across independent passes) to the first India operational review, as a concrete format rather than a free-form audit.
3. **Apply the same 48-hour re-plan protocol** to the first real regional-priorities conflict: a structured re-sequencing exercise with a signed outcome, rather than an ad hoc reshuffle.

13. Honest gaps and what is in progress (September 2026)

Gap	Honest status	Source
API-based application with an eval set	No shipped product calls a frontier lab's API at runtime; all AI work is mediated through my agentic coding harness. In progress: a small application with tool use and a 40-item eval set built from the compliance-firewall vocabulary rules. Not yet done, and said so.	EV05 §B.4; AI_Skills_OpenAI.md
MCP server	Heavy MCP *operator* (4 user-level servers + app connectors + 22 plugins), but no MCP server has been authored. In progress: an MCP server exposing the F&O calculator as tools.	INV Reconciliation #1; AI_Skills_OpenAI.md
Public reference architecture	Internal architecture documents (the router, the ownership map, the razorpay-rail skill) exist but are not published or partner-facing. In progress: a public reference architecture of the Desk.	AI_Skills_OpenAI.md
No hosted CI	No hosted CI system was found; the estate substitutes a local pre-deploy/post-deploy gate pattern instead (Section 5). Whether that was a deliberate solo-founder choice or worth changing is an open question I have not yet answered.	EV05 §D
Traction	₹500 verified revenue, all products, all time, as of 28 Aug 2026, one UPI payment. Zero book sales across 5 channels. The honest story is the production system and the constraint diagnosis (audience, not product), not invented traction.	EV06 §C
Not a career SA	Significant customer-facing technical-role experience is the bar several of these seats ask for; this candidate prepared BOMs and ran first-level technical discovery for 5 years, then became a hands-on architect as a founder, a different path to the same competence, named as such.	FIT
9 years vs. the 10+ bar	5 years quota-carrying sales + earlier development years + 15 months as founder totals approximately 9 years against a stated 10+ year minimum on the most senior of these seats.	FIT

14. Source index

Every file path cited in this document, once, with its abbreviation.

Abbrev.	Path
INV	04_AI/02_Evidence/00_MASTER_INVENTORY.md
EV05	04_AI/02_Evidence/05_evidence_AI_engineering.md
EV06	04_AI/02_Evidence/06_evidence_GTM.md
EV07	04_AI/02_Evidence/07_evidence_ventures_and_links.md
CV	04_AI/03_CV/CV_OpenAI_Founder.md
AI_Skills_OpenAI.md	04_AI/03_CV/AI_Skills_OpenAI.md
FIT	04_AI/01_Roles/FIT_ANALYSIS.md
OPENAI_India_roles	04_AI/01_Roles/OPENAI_India_roles.md

Methodology note on the marketing-agent count. An earlier pass in the source evidence cited "fifteen named agents"; a same-session reconciliation against the live agent-definitions directory found twelve. This document uses twelve throughout, per `INV` RECONCILIATIONS #2 (2026-09-07).

What is deliberately excluded. No personal trading P&L, return, or win-rate figures appear anywhere in this document. Flippy Tales (flippytales.in) is a family member's brand, not a venture of mine, and is not otherwise described here. No secrets, account IDs, or customer names appear above. This document is written using vendor-neutral vocabulary throughout: "agentic coding tools," "AI agents," "frontier-tier / fast-tier models," and "agent instruction files" describe the same real, operating system documented in the underlying evidence files.