

Atlas

The world's markets, and the cross-border book, modeled.

Cross-border portfolio intelligence for retail investors and advisory firms — self-service review, continuous multi-client monitoring, and institutional research — built compliance-first, with anti-hallucination guarantees at every AI touchpoint.

THE CHALLENGE

Three users, three unmet needs

Retail investor with cross-border holdings

Holdings split across Indian equities/MFs and global brokerages have no single, honest view. Existing tools either ignore the non-India leg entirely or bolt on a spreadsheet — no plain-English read on concentration, currency exposure, or real diversification benefit.

Independent advisor managing many client books

Reviewing each client's portfolio against a stated policy (target allocation, concentration caps, F&O exposure limits) is manual and doesn't scale past a handful of clients. Drift and single-name concentration go unnoticed between reviews. Firms with multiple advisors have no shared client roster.

Research-minded investor wanting institutional-grade signal

13F crowding, insider open-market buying, and macro regime are public data — but assembling them into a usable, point-in-time-correct signal (no look-ahead bias) is out of reach without a quant research team.

THE APPROACH

Atlas — three capabilities, one shared foundation

Atlas Copilot

Self-service, single-portfolio analyser. Upload holdings or a CAS PDF → scored analysis + plain-English narration. Progressive gates degrade silently if data is unavailable — never fabricates.

Portfolio Review

Advisor-grade continuous monitoring for many named clients: drift, concentration, F&O notional, DTAA flags. Firm tenancy — multiple advisors share one client roster and billing plan.

Global Analytics

Institutional/insider research toolkit: 13F crowding, Form-4 insider buying, a 5-indicator macro regime vote, a point-in-time factor model, and Black-Litterman construction.

Architecture & the anti-hallucination guarantee


Three independent mechanisms enforce that no number or claim a user sees is ever model-generated:

MECHANISM	WHAT IT DOES
Structured-in	The model only ever sees a FACTS block of already-computed numbers + rule-engine flag text — never raw holdings, never a client name.
Structured-out	The model must answer through a fixed tool-call schema. Prose outside that call is discarded.
Verify → retry → fallback	A GroundingVerifier rejects any draft whose numerals don't trace to FACTS; an AdviceLanguageScanner rejects buy/sell/guarantee phrases. Two failures fall back to a deterministic, LLM-free template — guaranteed to pass both checks by construction.


Production stance today: the LLM narration path is OFF by default — the deterministic template is what ships. Zero AI-hallucination risk on day one.

WORKFLOW

Atlas Copilot — single analysis



Portfolio Review — continuous monitoring



Tech stack & AI discipline

Frontend	Next.js 14 + React + TypeScript — frontend-global, independently deployed
Backend services	Python 3.11+, FastAPI — Atlas Copilot (hexagonal), Portfolio Review, Global Analytics
Shared math	portfolio_engine/ — pure NumPy, zero I/O (Ledoit-Wolf, allocators, risk metrics)
Storage (metadata)	PostgreSQL (Supabase) — Row-Level Security on every user-data table
Storage (market data)	Hive-partitioned Parquet + DuckDB — data/global_daily/
AI	Anthropic Claude API — structured-in/structured-out only, off by default
External data	SEC EDGAR, FRED (keyless) · Alpaca (IEX daily) · NSE public JSON
Compute (production)	GCP Cloud Run, asia-south1 — India data residency for SEBI/DPDP
Security	AES-256-GCM PII · SHA-256 hashing · Argon2id · RS256 JWT (15-min TTL)

AI PIPELINE & TOKEN UTILIZATION

No multi-agent system. Unlike Learning's LangGraph strategy-builder, Atlas Copilot's narration is a single structured call — deliberately simpler, because the job here is narration-only, not multi-turn negotiation.

Token utilization today: effectively zero

`WEALTH_COPILOT_LLM_NARRATION=false` in production — the deterministic template ships. No Anthropic API call is made for Atlas Copilot narration today.

- **Small, fixed input** — only the FACTS block reaches the model, never raw holdings or full transaction history.
- **One call per analysis** — no agentic tool-calling loop; at most one retry before falling back to the zero-token template.
- **Content-addressed caching** — identical inputs never re-trigger computation or re-narration.
- **No PII to the model** — a SHA-256 hashed `user_id` is the only context identifier.

OUTCOMES

Structural ROI — qualitative, not fabricated

No revenue, pricing, or user-projection figures are invented here — these are the real, structural drivers already true of the codebase today.

- **Shared math, not 4× reimplementation** — `portfolio_engine` is used by Atlas Copilot, Portfolio Review, Global Analytics, and Lens's scanner. One fix benefits all four.
- **Compliance-by-construction** — advice-language scanning, source-ref-checked briefings, and the audit hash-chain are structural, not app-layer checks a direct API call could bypass.
- **AI cost is opt-in, not baseline** — the deterministic template ships today at zero marginal AI cost; LLM narration is a toggleable, bounded-cost enhancement.
- **Content-addressed compute** — identical-input analyses dedupe at the DB layer.
- **One platform, four monetizable surfaces** — all four capabilities share one frontend, one auth system, one plan model.

STATUS

Built & live-verified

- All three capabilities, one shared frontend
- Firm tenancy — Portfolio Review
- Content-addressed, reproducible snapshots
- 3-mechanism anti-hallucination pipeline
- Scanner folded natively into Atlas frontend
- Production on GCP Cloud Run, `asia-south1`

Open items

- Portfolio Review hub embed renders blank — open bug
- LLM narration: prod grounding test run (gate)
- Smart\$ composite — 30% wired (insider cluster only)
- Stock/fund look-through — no reliable free source yet
- No CI/CD for GCP Cloud Run — manual deploys today