

Vyoman


An end-to-end aerial photography business — portfolio, print store, operations console, and AI mission intelligence — designed, built, and operated solo.

vyomanaerials.com · vantage.vyomanaerials.com · shop.vyomanaerials.com

THE PROBLEM

Running a drone photography business end-to-end usually means stitching together a dozen disconnected SaaS tools — content in one place, commerce in another, fulfilment and operations somewhere else. There's no integrated system where a single change propagates everywhere, and AI handles the mechanical 70% of production so the human keeps the creative 30%.

ARCHITECTURE & DATA FLOW



WHAT I BUILT

- › Portfolio site (Next.js 15) with a print store: Stripe checkout, Gelato print-on-demand fulfilment, DHL shipment tracking, and 4-tier digital licensing — entirely automated, no manual intervention on standard orders.
- › A canonical data layer in Neon Postgres linking every portfolio item to its print product and sales record. Update one item in the console → it appears on the public site within seconds via ISR. No code changes, no deploy.
- › Django ops console on Railway: order lifecycle management, automated lifecycle emails, LLM intelligence dashboard, and vision model inference that classifies raw footage against the product registry.
- › Vantage — a 12-phase AI mission intelligence system built in 3 days: streaming LLM briefing generator, airspace WMS compliance checks, weather scoring, shot references via FLUX image generation, and portfolio ROI attribution linking flights to orders.
- › Automated broadcast engine: content created once in the console, published to Instagram and YouTube without manual steps.

WHAT THIS DEMONSTRATES

Full-scope solo product delivery — commerce, operations, AI intelligence, and creative automation integrated into one system with a canonical data layer. End-to-end order pipeline built and operational: Stripe checkout through Gelato fulfilment to DHL delivery, fully integrated. Every architectural decision made and shipped alone.

KEY NUMBERS

29live
products**464**canonical
variants**6**services
deployed**12**AI
mission
phases

TECH STACK

Next.js 15

Django 5.2

Neon Postgres

Stripe

Gelato

Railway

Vercel

Cloudflare R2

Claude

FLUX

Celery

Resend

Sameer Bhalerao

Senior Analytics & AI
Leader · Luxembourg
11+ years · 9 years at
Amazon (L6)
Built and operates this
solo

sameerbhalerao.com
linkedin.com/in/sameervb

Disha Analytics


Compresses two analyst-days of restaurant research into a five-minute, AI-generated intelligence report with a ranked action plan.

disha-analytics.com · console.disha-analytics.com

THE PROBLEM

Independent restaurant owners make operational decisions without structured data. They know something is wrong — mixed reviews, thin margins — but cannot afford consultants and have no tool that turns scattered signals into a clear, prioritised action plan. A senior analyst would take a week. The data exists; the bottleneck is the pipeline to process it.

ARCHITECTURE & DATA FLOW



WHAT I BUILT

- › Input: restaurant name and city. Output: a ranked action plan in ~5 minutes. Eight specialist AI agents run in parallel — competitor, financial, review, menu, hypothesis verifier, validation, synthesis, and owner narrative.
- › Scrapes 8 public sources (TripAdvisor, Yelp, Google Places, and others) and benchmarks every finding against local competitors automatically.
- › Belief graph: every conclusion maps to a specific raw observation via Claims, Signals, and weighted Edges. No black-box outputs — every recommendation is citable and traceable.
- › Labs module auto-generates statistically valid experiments from the report findings: switchback design, duration, sample size, and success criteria written from a single plain-English hypothesis.
- › Three delivery modes: Public Scan (public data only), Validated Report (owner financials included), Owner Walkthrough (decision deck). PDF reports via WeasyPrint.

WHAT THIS DEMONSTRATES

The consulting instinct from Amazon, productised as a B2B SaaS — taking a complex data problem, defining a structured output format, and delivering it under time pressure at production scale. Validated with a restaurant in Luxembourg as the first pilot report.

KEY NUMBERS

8

specialist
AI
agents

8

public
data
sources

~5m

per
report

3

delivery
modes

TECH STACK

Django 5

React 18

PostgreSQL

Celery

Redis

WeasyPrint

Hetzner

Claude API

Docker

Sameer Bhalerao

Senior Analytics & AI
Leader · Luxembourg
11+ years · 9 years at
Amazon (L6)
Built end-to-end, solo

sameerbhalerao.com
linkedin.com/in/sameervb


Trust infrastructure for the informal local economy inside residential societies — workers build a portable reputation, residents find reliable help.

web-production-0d809.up.railway.app

THE PROBLEM

The informal labour market inside residential societies — plumbers, electricians, cleaners — runs entirely on word-of-mouth: no trust signals, no transparent pricing, no accountability structure. Residents cannot find reliable workers; skilled workers cannot build a reputation that travels beyond a single building. The market exists; the infrastructure to make it trustworthy does not.

ARCHITECTURE & DATA FLOW



KEY NUMBERS

4
input
languages

P2P
UPI
payments

NLP
job
parsing

Live
pilot,
Bangalore

TECH STACK

Django 5.1

PostgreSQL

Telegram Bot

Claude API

Railway

Python

Sameer Bhalerao

Senior Analytics & AI
Leader · Luxembourg
11+ years · 9 years at
Amazon (L6)
Built solo · ground ops
with Alok

sameerbhalerao.com
linkedin.com/in/sameervb

WHAT I BUILT

- › Residents describe a job in plain English — by typing or by voice, in English, Hindi, Kannada, or Marathi. An AI parser extracts structured fields (skill, date, duration, headcount) without any form or dropdown.
- › An LLM pricing engine (Claude) sets a specific, justified rate from job category, duration, and local market rates — not a range to negotiate from, a concrete number.
- › Matched workers receive a Telegram notification and apply with one tap — no new app required. The platform meets workers where they already are.
- › Trust score accumulates from job completions and resident reviews — verifiable, portable reputation that builds over time.
- › Payments are peer-to-peer via UPI. The platform never handles money. Full field operations documentation: job lifecycle SOPs, dispute resolution, and onboarding guides built with an on-ground partner.

WHAT THIS DEMONSTRATES

Product thinking for a low-tech, voice-first market: 4 languages, Telegram-native, P2P payments, and operational design (SOPs, disputes, field ops) built alongside the code — not as an afterthought.

ComplyLens


Makes CSSF-inspection-ready compliance documentation auditable — reasoning visible, evidence cited, every determination defensible.

complylens.vercel.app

THE PROBLEM

Luxembourg AML/KYC compliance teams produce manual review documentation with no structured evidence chain. In 2024, CSSF handed down €4.7M in fines — up 52% year-on-year. The problem isn't whether a review happened; it's whether the reasoning was auditable by a regulator. AMLR applies July 2027, creating a first-mover window for tools built on traceable AI reasoning.

ARCHITECTURE & DATA FLOW



WHAT I BUILT

- › RiskRationale: classifies client AML risk and generates structured evidence citations mapped to specific CSSF risk indicators — not a black-box score, a reasoned classification with a regulatory anchor for every finding.
- › ReviewMemo: generates full 11-section periodic review memos with a companion CSSF audit pack (ZIP export). Claude drafts the entire structure; the compliance officer reviews, corrects, and signs.
- › AlertDesk: sanctions and PEP alert dispositions with a four-eyes approval workflow. Claude drafts the disposition; a second reviewer approves before the record is locked and immutable.
- › Full audit log records every AI determination. Analysts correct AI output through a structured domain-correction loop that feeds back into the system. Every citation maps to a CSSF indicator via a codified regulatory_framework module.
- › Password-gated live demo (Meridian Holdings synthetic scenario) accessible to prospects. Rapid MVP deployment within 48 hours to validate architecture.

WHAT THIS DEMONSTRATES

LLM product design where the core constraint is regulatory audibility — not just "does the AI output work" but "can a regulator trace every determination back to its evidence." Built for Luxembourg's specific CSSF framework with a clear market window before AMLR 2027.

KEY NUMBERS

3

product
pillars

48h

to MVP

4-
eye

approval
workflow

2027

AMLR
deadline

TECH STACK

FastAPI

Next.js 15

PostgreSQL

Railway

Vercel

Claude

WeasyPrint

Celery

Redis

Sameer Bhalerao

Senior Analytics & AI
Leader · Luxembourg
11+ years · 9 years at
Amazon (L6)
Co-founder with Sachin
Bhalerao

sameerbhalerao.com
linkedin.com/in/sameervb

Soul Spark

A multi-agent AI advisor that unifies finance, health, career, and six other life domains — all inference local, zero personal data leaves the device.

soulspark.me · app.soulspark.me

THE PROBLEM

Personal data is fragmented across a dozen apps — health, money, career, tasks all live separately — which makes it impossible to see cross-domain patterns or turn the combined picture into clear decisions. Cloud-hosted AI services that could help require routing the most sensitive personal data through an external provider. Most people accept this trade-off because local inference seemed impractical at real scale.

ARCHITECTURE & DATA FLOW



WHAT I BUILT

- › Nine life domains — finance, health, career, social, growth, library, travel, inner self, and mindset — unified into a single on-device DuckDB store. Real data sources include Apple Health and Google Sheets, ingested automatically.
- › A packet-native multi-agent runtime (an internal graph runtime with shared memory and a pending-action loop) reasons across all domains simultaneously — cross-domain patterns that no single-domain app can see.
- › All AI inference runs locally via Ollama. Zero cloud inference. Zero subscriptions. Zero personal data exfiltration. The model and the data stay on the device.
- › A conversational advisor moves from "here is your data" to "here is what you should do next" — the shift from dashboards to decisions. Public access via a Cloudflare tunnel; the data and model stay local.
- › 70+ signals, 300+ cross-domain connections surfaced into a single daily brief and an interactive advisor chat grounded in real personal data.

WHAT THIS DEMONSTRATES

Privacy-first AI architecture at scale — a multi-agent intelligence system running entirely on-device for the most sensitive data category possible: personal life. Proves that local-first LLM systems are production-viable, not just theoretical.

KEY NUMBERS

9

life domains

70+

signals tracked

300+

cross-domain links

0

cloud inference

TECH STACK

Python

Streamlit

Ollama

DuckDB

Multi-agent runtime

Cloudflare

Apple Health

Sameer Bhalerao

Senior Analytics & AI Leader · Luxembourg
11+ years · 9 years at Amazon (L6)
Built and runs this solo, daily

sameerbhalerao.com
linkedin.com/in/sameervb

Narrate


Converts any document into a narrated audiobook — voice personality per book, AI cover art, adaptive streaming, and lock-screen captions.

narrate-psi.vercel.app

THE PROBLEM

PDFs, whitepapers, and long documents sit unopened because sustained reading requires uninterrupted attention that modern life rarely provides — and audiobook platforms only serve commercially published content. Document-first sources like government notices, contracts, and circulars are completely inaccessible to anyone who cannot read due to visual impairment, low literacy, or language barriers.

ARCHITECTURE & DATA FLOW



WHAT I BUILT

- › Narrate-Ready pipeline: any uploaded file (PDF, TXT, Markdown) runs through a 9-rule preprocessing spec — document type classification, language detection, header stripping, bullet-to-ordinal conversion, table-to-prose, abbreviation expansion. Published as a canonical, implementation-independent standard.
- › Voice personality per book: one AI call at upload writes voice instructions specific to this book's genre and tone, injected into every TTS chapter call. A strategy whitepaper and a memoir sound noticeably different.
- › Adaptive buffering: a 200-chapter book starts playing in under 30 seconds. Conversion runs in the background as you listen — the same model as YouTube's progressive video loading. GlobalConversion at layout level survives all navigation.
- › AI cover art (DALL-E 3), narrated brief (full-book audio summary), "Where you are" context card (SO FAR / NEXT UP for returning listeners), and intro/outro bookends that frame every listening session.
- › Lock-screen captions in 3 modes (off / phrase clusters / full sentence) via iOS MediaSession API. No database — Cloudflare R2 is the complete data layer for metadata, audio, cover art, and chapter data.

WHAT THIS DEMONSTRATES

Product craft — architecturally complex (adaptive streaming, preprocessing pipeline, MediaSession integration) and useful enough to become a daily habit. The Narrate-Ready preprocessing spec is a publishable document-to-audio standard.

KEY NUMBERS

8+books
narrated**<30s**start, 200
chapters**3**lock-
screen
modes**0**databases
used

TECH STACK

Next.js 15

TypeScript

OpenAI TTS

GPT-4o-mini

DALL-E 3

Cloudflare R2

Vercel

MediaSession API

Sameer Bhalerao

Senior Analytics & AI
Leader · Luxembourg
11+ years · 9 years at
Amazon (L6)
Built solo · personal daily
driver

sameerbhalerao.com
linkedin.com/in/sameervb

Vyoman Earth Intelligence


Turns Luxembourg's public geospatial data and drone imagery into an evidence-backed, human-reviewed site briefing — in minutes, not a GIS analyst's week.

Local MVP 0 prototype · extends the Vyoman Aerials drone operation

THE PROBLEM

Luxembourg's public geospatial data — 10cm orthophotos, LiDAR DEM via the Geoportail — is excellent and almost entirely unused outside government workflows. Drone operators already fly the sites, but turning imagery into a client-ready report still means a manual write-up with no structured evidence trail behind the conclusions.

ARCHITECTURE & DATA FLOW



WHAT I BUILT

- › Coordinates and datasets in. Evidence-cited AI briefing out — Executive Summary, Findings, Review Checklist, Limitations — in under a minute.
- › Live-fetches Luxembourg's public Geoportail WMS orthophotos on demand — no manual downloads, correctly georeferenced GeoTIFF output.
- › Evidence-packet architecture from three prior products — every AI claim traces to a cited dataset, never invented.
- › 2D map and 3D terrain view, human annotation workflow with draft → reviewed → approved states, and a full DuckDB audit trail.
- › Local-first: Ollama with no cloud dependency — deterministic mock-report fallback so the demo always works.

WHAT THIS DEMONSTRATES

The Amazon BI instinct applied to geospatial data: take raw public data, structure it into evidence, and deliver a decision-ready briefing. Three prior products converged independently on the same evidence-review-audit pattern — built on the Vyoman Aerials drone operation, flying in Luxembourg since April 2026.

KEY NUMBERS

3

live data
sources

7

workflow
tabs

10cm

Geoportail
source
resolution

0

cloud
API keys
required

TECH STACK

Streamlit

Python 3.13

DuckDB

GeoPandas

Rasterio

PyDeck

Ollama

ReportLab

Geoportail WMS

Sameer Bhalerao

Senior Analytics & AI
Leader · Luxembourg
11+ years · 9 years at
Amazon (L6)
Built end-to-end, solo

sameerbhalerao.com
linkedin.com/in/sameervb