



One account. Four services. An AI layer that is yours.

Vespera is a working software platform — portfolios, shops, an AI assistant with an outreach engine, and a study space — built on one identity, one design language and one bring-your-own-AI layer. It is offered for acquisition, licensing or partnership.



Mirror

Portfolios & CVs
mirror.vespera.me



Souk

Shops & landing pages
souk.vespera.me



Anchor

Assistant & outreach
anchor.vespera.me



Atlas

Study & thinking space
atlas.vespera.me

4

services in production

381

HTTP endpoints

69

database tables, one schema

2,379

checks in 71 test scripts

Production-ready, running on a single server, in private beta by invitation, validated with a small circle of users.

CONTACT Ghalmi Ahmed Nour Eddine ·

ghalmiahmednouredine@gmail.com

ONLINE info.vespera.me · Open to offers — let's talk.

01 • SUMMARY

Software that reflects the person using it.

Most website and shop builders hand everyone the same handful of templates. Vespera generates a complete design specification for each person — palette, typography, layout, texture, motion — and renders the page from it. Two people never get the same site. Around that idea sits a platform: one account across every service, one AI layer each person controls, and an assistant that can act on all of it.

A specification, not a template

Mirror and Souk share a design vocabulary of 42 palettes, 41 font pairings, 12 textures and 11 backgrounds; the generator pins what makes a look recognisable and varies the rest.

One identity

One users table, one sign-in, Google sign-in, and a single-use code that carries a session from one service to the next without ever putting a token in a URL.

Bring your own AI

Each person connects their own provider — 9 presets plus any OpenAI-compatible endpoint — or pastes one Groq key and lets Easy mode configure every job. Keys are encrypted; nothing is locked to one vendor.

An assistant with guardrails

Anchor's assistant edits portfolios, shops and study material on request — at most 25 changes and 3 deletions per instruction, every change logged, reversible ones undoable.

Why it matters to a business: the hard, unglamorous parts — identity across origins, per-user AI credentials, provider fallback under rate limits, guarded write access for an AI agent, Gmail sending that respects quotas and consent — are already built, tested and running. A new product inherits them on its first day.

At a glance

SERVICES	Mirror · Souk · Anchor · Atlas — all four in production
LANGUAGES	English · French · Arabic (right-to-left)
BACKENDS	Python · Flask · SQLAlchemy · Alembic · PostgreSQL
FRONTENDS	Next.js 15 · React 19 · TypeScript
AI	Bring your own provider · Groq Easy mode · up to 5 prioritised routes · shared cooldowns
SIZE	381 endpoints · 69 tables · 46,203 lines of Python · 83,435 lines of TS/CSS
TESTS	71 test scripts · 2,379 checks
HOSTING	One server in Paris (EU) · nginx · systemd · Let's Encrypt
STATUS	Production-ready, running on a single server, in private beta by invitation, validated with a small circle of users.

02 • THE SERVICES (1/2)



Mirror

Portfolios & CVs · mirror.vespera.me

97 endpoints 13 test scripts 374 checks

A portfolio and CV builder that generates a complete design specification for each person, then keeps the content live for any site that wants to read it.

WHO IT IS FOR Students, researchers, engineers and creatives who need a distinctive public presence and up-to-date CVs without designing them.

GENERATION

- AI-generated design specification per person — palette, typography, layout, texture, motion — rendered by a spec-driven renderer, not a theme picker
- Content analysis and a writing assistant for summaries and project text
- Deterministic engine seeded from the person's data when no AI is connected — still unique per person
- Style gallery with named looks to start from

CONTENT

- Portfolios with ordered sections and projects
- Rich project pages built from content blocks
- Automatic translation of a whole portfolio into French or Arabic (structure-preserving JSON translation)
- Academic and professional CV generator from the same data
- Media uploads with per-plan limits

PUBLISHING

- Public portfolio API: published content as JSON
- Linked bespoke sites — a separate site on its own domain reads Mirror live and is edited from the Mirror dashboard
- Draft / published states and view counts

ACCOUNT & BUSINESS

- Google sign-in; accounts can exist without a password; e-mail verification and e-mail change flow
- Hosts the shared sign-in and the cross-service handoff for every Vespera service
- Plans, subscriptions and entitlements — per service or as a bundle — with payment arranged off-platform
- Admin console: people, plans, subscriptions, media limits, billing mode, mail settings, audit log, platform-wide overview, impersonation for support



Souk

Shops & landing pages · souk.vespera.me

90 endpoints 9 test scripts 273 checks

A shop builder with generated storefronts, a page builder, and the complete order lifecycle — designed for sellers who are paid directly, outside the platform.

WHO IT IS FOR Independent sellers, makers, small services and food or rental businesses that want a storefront that looks like nobody else's and a clear order process.

STOREFRONTS

- 40 named storefront styles — commerce foundations and art movements from Bauhaus to risograph
- Generated structure: 6 headers, 6 hero layouts, 10 product-card designs, 4 catalogue layouts, 8 price styles, 9 button styles
- Style gallery and one-click restyle; the seeded generator keeps two shops in the same style different
- Shop kinds: products, services, digital, food, rental, mixed

PAGE BUILDER

- Landing pages from 23 block types — hero, features, products, gallery, testimonials, pricing, team, video, timeline, flip book, PDF, downloads, countdown, map, comparison, social links and more
- 60 distinct layouts across 11 of those blocks
- Closed vocabulary: anything the renderer cannot draw is refused at the door; links are restricted to safe schemes and video embeds to an allowlist
- Pages reuse the shop's own design system

COMMERCE

- Products, categories, stock, featured and hidden items
- Cart and checkout
- Full order lifecycle — 13 states from pending to delivered, with courier stages, disputes and refunds
- Buyer-seller agreement without touching money: the buyer declares payment, the seller confirms it, and an order completes only when both sides have confirmed
- Every transition on a timeline; goods return to stock on cancellation or refund
- Buyer-seller messaging, notifications and reviews

REACH

- Translations of shop content
- Site API for linked shop sites — a live site on its own domain reads the shop and is edited from the Souk dashboard
- Plans and entitlements shared with the rest of Vespera

02 • THE SERVICES (2/3)



Anchor

Assistant & outreach · anchor.vespera.me

188 endpoints 41 test scripts 1,432 checks

A personal command centre: tasks, flexible Rooms, an AI assistant that acts across every Vespera service, and an outreach engine that writes, sends, tracks and reads replies to personal letters at scale.

WHO IT IS FOR Job seekers, researchers applying for positions, freelancers and small teams who write to many people and need every letter to be personal — and anyone who wants one assistant for their whole Vespera account.

THE ASSISTANT

- Chat that acts on the person's data: 43 actions across Anchor, Mirror and Souk, plus 79 generated from Atlas's operation registry
- Guardrails on every instruction: at most 25 changes, 3 deletions and 20 creations; a deletion must name its target ("all" is refused); irreversible order actions count as deletions
- Refusals are explained and counted: after 3 strikes the assistant loses write access until an admin restores it
- Action trail of every change, with undo where the change is reversible
- Long-term memory the person can read and edit; multiple saved chats
- Vision: images in chat, or described for a model that cannot see
- Voice: wake word, speech-to-text, text-to-speech, and read-along highlighting of the spoken word
- Runs on the person's own model, or in the browser through Puter when they have none

WORK & ORGANISATION

- Tasks with priorities, deadlines, focus and notes; mood and energy check-ins feeding a daily balance
- Rooms: spaces the assistant builds on request from 19 field types, 14 layouts and 12 presets — a contact list, a reading log, a budget
- Scout: searches the web and reads the pages itself to find openings, and reports a lead as verified only when its page was actually retrieved

OUTREACH ENGINE

- Import up to 2,000 recipients from a spreadsheet of ready e-mails, a contact file, or leads found by Scout
- One personal AI-written letter per person, 4 written at a time in the background, in the language chosen (interface in English, French and Arabic); machine phrasing detected and rewritten once
- Review, edit and approve every letter; nothing is sent without approval
- Gmail sending with human pacing (1.2 s minimum between messages), a rolling 450-per-24 h cap, and scheduling up to 90 days ahead; rate limits defer a letter, never fail it
- Open and click tracking (scanner clicks and bot opens are filtered), bounce detection
- A reply watcher that reads answers in the Gmail thread and sorts them into 8 kinds — interested, negotiation, question, rejected, unsubscribe, automatic reply, bounce, other — with a summary and a suggested next step
- Orphan-reply matching: answers that arrive outside the original thread (out-of-office notices, new threads) are still tied to their letter
- Kind-reply drafts for every person who answered, and an optional request for feedback after a refusal — always drafts, always approved by the person
- Contact profiles: a mini-CRM of everyone written to, with status, notes, tags and the full timeline
- Insights: the letter-journey map from draft to answer, a 30-day activity chart and the sending capacity left today

MAIL & NOTIFICATIONS

- Casual mail: ordinary untracked e-mail to up to 20 people
- Signature studio with 4 designed layouts, optionally applied to every message
- Inbox reading with sender preferences
- Telegram bot: link an account, ask for replies and status, receive alerts as answers arrive
- Gmail consent split in two: sending needs only the send permission; reading replies is a separate, optional permission



Atlas

Study & thinking space · atlas.vespera.me

83 operations 4 test scripts 144 checks

A study planner, a library that answers questions about its documents, infinite whiteboards, interactive guides, study groups and projects — every operation also drivable from Anchor’s chat.

WHO IT IS FOR Students, study groups, teachers and small project teams who plan, read and think visually.

PLAN

- Subjects, tasks and a calendar of study blocks
- AI-generated study plans
- A home view with today’s tasks, upcoming exams, next sessions, recent boards and guides

READ

- Library of 5 resource kinds — PDFs, links, notes, images, files — with PDF text extraction
- Ask questions of a document and get summaries grounded in the most relevant passages


THINK

- Infinite whiteboards with 10 element types — sticky notes, text, shapes, arrows, pen, images, files, stickers, Mermaid diagrams, frames — and mind maps
- Multi-user sync with per-element revisions
- AI “Full board”: a whole presentation-quality board from a topic — title, mind map, laid-out sections — using web search where the model supports it, with links verified
- Interactive AI-generated guides: quizzes, checklists, steps, callouts, diagrams, progress tracking and “explain this”

TOGETHER

- Study groups with chat, invitations and member roles
- Projects with milestones, AI task splitting across members, and AI roadmap boards
- Sharing by person, group, link with view or edit roles
- 83 operations in one shared registry; 79 of them can be driven from Anchor’s chat

A · System architecture




One server. nginx terminates TLS for every subdomain and routes to each service’s Next.js frontend and Flask backend; all backends import the same platform package and share one PostgreSQL database. External calls go to the person’s own AI provider Google (sign-in and Gmail) Telegram and

03 · CONNECTIONS AND THE AI LAYER


B · Service connection map

WHO ACTS ON WHAT




in-process · no HTTP between services
links in AI-written content: http(s) only

SINGLE SIGN-ON HANDOFF



Anchor's assistant reaches Mirror, Souk and Atlas in-process through guarded actions — no service calls another over HTTP. Atlas's operations live in the shared package, so Atlas's own API and Anchor's chat run the same functions. Right: the single sign-on handoff.

D · The AI layer



ROLES BESIDE CHAT

Vision · image → text Speech · text → speech Hearing · whisper-large-v3 → whisper-large-v3-turbo Search · native web search

Keys are encrypted per person and only ever sent to the address they were saved with.

A request walks the person's routes top to bottom and each route's models in order. Easy mode is one of the routes. Failures are sorted: provider-level problems skip the route, model-level problems put that model on a shared cooldown.

05 · THE OUTREACH ENGINE


Hundreds of personal letters, and every answer read.

The most developed part of the platform, and the one most directly useful to a business: it turns a list of people into personal letters, sends them through the person's own Gmail at a human pace, and reads what comes back.

01 Import Spreadsheet, contact file or Scout leads · up to 2,000 people	02 Compose One letter per person, 4 at a time, machine phrasing rewritten	03 Approve Review, edit, approve or schedule	04 Send Gmail, 1.2 s apart, 450 per 24 h, deferred on limits	05 Track Opens, clicks, bounces · bots filtered	06 Read Reply watcher · 8 kinds · orphan replies matched	07 Act Reply drafts, feedback requests, contact profiles, Telegram alerts
--	--	---	---	--	---	--

The letter journey (Insights graph)

ILLUSTRATIVE NUMBERS



Answered by a person = interested + question + negotiation + rejected


A recreation of Anchor's Insights flow. Band width is the number of letters. The numbers are illustrative, chosen for legibility — they are not real campaign data.

Consent Sending asks only for Gmail's send permission. Reading replies is a second, optional permission.


Honesty Opens are shown as hints, not facts; scanner clicks and bot opens are stored apart and not counted.

Pacing Rate limits and daily caps defer a letter with back-off; they never fail it.

Respect "Do not contact me" is honoured: no reply draft, no feedback request, ever.



The product's own flow chart, as it renders in Anchor.



Replies — sorted, summarised, with a next step and a kind-reply draft one click away.

06 · FOR ACQUIRERS & PARTNERS

What you would be taking on.

WHAT IS INCLUDED

- Source code of all four services — backends and frontends
- The shared platform package: identity, billing, AI layer, encryption, design vocabulary, Atlas operations
- Infrastructure: nginx configuration, systemd units, gunicorn settings, deploy and server-setup scripts
- The test suites — 71 scripts, 2,379 checks
- Documentation: build contracts, theming and internationalisation notes, subscriptions, deployment and DNS guides
- The Vespera brand, the four service marks and the vspera.me domain

WAYS TO WORK TOGETHER

- Outright acquisition** — The platform, code, brand and domain transfer to you, with a handover.
- Licence** — Use the platform or one service under licence, in your own infrastructure.
- White-label** — Your name and look on a service — the design system is token-based and the services were built to be renamed.
- Co-development** — Build a fifth service, or a vertical version of one, together.
- Sponsorship or pilot** — Fund continued development, or run a pilot with your users and keep what you learn.

A FIRST 30 DAYS COULD LOOK LIKE THIS An offer, adjusted to you — not fixed terms.

WEEK 1

Technical walkthrough A live tour of every service, the codebase, the data model and the deploy, with your engineers.

WEEK 2

Sandbox access A private instance seeded with demo data so your team can use and break things.

WEEKS 3-4

Handover plan Agreed scope, documentation gaps closed, and a written plan for running it on your side.

Honest current status. Production-ready, running on a single server, in private beta by invitation, validated with a small circle of users. No revenue, user numbers or customers are claimed here, because there are none to report yet.

ABOUT THE BUILDER

Vespera was designed, built and is run by one engineer: product, backends, frontends, AI layer, infrastructure and operations. Ghalmi Ahmed Nour Eddine is a biomedical engineer — Master 2 in Biomedical & Hospital Informatics from the Department of Biomedical Engineering, Université Abou Bekr Belkaid, Tlemcen — and was a research intern at LaTIM (INSERM UMR 1101) in Brest, working on foundation and vision-language models for retinal imaging.

He is now seeking funding for a PhD in medical-imaging AI. Acquiring, licensing or sponsoring Vespera directly supports that doctoral research — and puts a platform built with care in hands that can take it further. Portfolio and CVs: ahmed.vespera.me

QUESTIONS

Is it live?

Yes. All four services run in production at their own subdomains of vspera.me, by invitation during the private beta.

Can I see it?

Yes. Ask for a walkthrough and you will get a live demonstration and, if useful, a sandbox account with demo data.

What is the stack?

Python (Flask, SQLAlchemy, Alembic) with PostgreSQL behind each service; Next.js 15, React 19 and TypeScript in front; nginx, systemd and gunicorn on one Linux server.

How is AI handled, and what about privacy?

Each person brings their own provider and key; keys are encrypted at rest and never sent back to a browser. The platform holds no AI keys of its own by default. Gmail access is split into separate send and read permissions, and personal data is not sold or used for advertising.

Can it be white-labelled?

Yes. Colours, type and chrome are design tokens with light and dark themes, and Atlas was built with its name and look meant to be swapped in one place.

What would you need to take it over?

One Linux server (or your own cloud), PostgreSQL, a domain, a Google Cloud project for sign-in and Gmail, and optionally a Telegram bot token and a Cloudinary account. AI is per user. The deploy scripts, documentation and test suites come with it.

Let's talk about what Vespera could become with you.

E-MAIL ghalmiahmednouredine@gmail.com
 WEB info.vespera.me · vspera.me · ahmed.vespera.me

Open to offers — let's talk.