

AI Career Intelligence Agent

Essential Product Management Portfolio Document Pack

PRODUCT	STATUS	ROLE
Human-in-the-loop career automation	In development	Product Manager / Solution Architect

	Portfolio note This document presents the product strategy, product requirements and delivery plan for an in-development personal AI career platform. Metrics are defined as targets unless explicitly marked as measured outcomes. No unverified business impact is claimed.
--	--

Document guide

This portfolio pack combines the essential product artifacts a hiring team would expect from an AI Product Manager: product strategy, research framing, PRD, user experience, AI governance, roadmap, metrics, launch plan and product decision log. It is designed to be read as a coherent case study rather than a collection of disconnected templates.

Section	What it demonstrates
1. Executive product brief	Problem framing, product judgment and concise communication
2. Product strategy	Vision, positioning, goals, non-goals and differentiation
3. User discovery	Personas, jobs-to-be-done, pain points and assumptions
4. Product requirements	Prioritization, user stories, requirements and acceptance criteria
5. AI and solution design	Human approval gates, RAG grounding, model routing and architecture
6. Roadmap and delivery	Phased execution, dependencies and scope control
7. Metrics and evaluation	North-star thinking, guardrails and experimentation
8. Risk and launch plan	Security, privacy, reliability and rollout readiness
9. Portfolio case study	A polished narrative suitable for a portfolio website or interview

1. Executive product brief

One-line product statement

An AI-powered career operations platform that finds relevant roles, explains fit using verified career evidence, prepares personalized application materials with approval, and supports the user through interviews and offers.


1.1 The problem

A serious job search is a fragmented workflow spread across job boards, email alerts, resumes, cover letters, spreadsheets, calendars and personal notes. The user repeatedly performs the same high-effort tasks: finding fresh roles, deciding whether a job is relevant, rewriting application materials, tracking follow-ups and preparing for interviews. Existing automation tools often optimize for application volume rather than application quality, which can create inaccurate applications and damage trust.

1.2 Product opportunity

The product opportunity is to automate the repetitive work while preserving human ownership of consequential decisions. The system should operate as a career intelligence assistant rather than an autonomous applicant. It can discover, organize, analyze and draft; the user decides, approves and submits.

Human-in-the-loop AI Career Workflow



Automation handles discovery, analysis, drafting, reminders and preparation; the user controls every consequential action.

Figure 1. The end-to-end product experience keeps the user in control at every consequential step.

1.3 Target outcome

Reduce the time required to identify and prepare high-quality applications while improving relevance, consistency and confidence. The initial target is to turn a multi-hour daily process into a focused review workflow without sacrificing accuracy or personal judgment.

1.4 Product principles

Principle	Product implication
Evidence before eloquence	The system may only claim skills, experience and outcomes supported by approved career evidence.
Human approval by default	No resume, cover letter, application, follow-up, calendar event or offer decision is finalized automatically.
Quality over volume	Optimize for a smaller number of relevant, well-prepared applications rather than mass application.
Transparent reasoning	Show why a role matches, what evidence supports the recommendation and where gaps remain.
Cost-aware intelligence	Use Nano for lightweight classification, Mini for normal generation and Premium only by explicit request.
Privacy as a feature	Keep parsing and embeddings local where practical; minimize data sent to external models.

2. Product strategy

2.1 Vision

Create a trusted personal career operating system that continuously converts a user's verified experience into better decisions, better application materials and better interview preparation - without taking control away from the user.

2.2 Mission

Help professionals spend less time on repetitive job-search administration and more time choosing opportunities, presenting their real strengths and preparing for meaningful conversations.

2.3 Strategic goals

Goal	Success signal
Increase opportunity relevance	Most surfaced jobs meet configured role, location and work-mode preferences.
Reduce preparation effort	The user can move from selected job to reviewable application package in minutes, not hours.

Goal	Success signal
Protect professional credibility	Unsupported claim rate remains near zero and all major claims have provenance.
Build a reusable career memory	New agents can retrieve the same approved experience without duplicating or contradicting facts.
Support the complete journey	The product remains useful after application submission through interview and offer stages.

2.4 Non-goals

- Mass-applying to every available vacancy.
- Logging into or scraping authenticated job platforms in violation of platform rules.
- Automatically submitting forms, sending email or messaging recruiters.
- Inventing experience, achievements, metrics, company facts or compensation benchmarks.
- Replacing the user's judgment on career, legal, tax, immigration or financial decisions.
- Building a multi-tenant commercial SaaS product in the initial personal-use release.

2.5 Positioning

Alternative	Limitation	Product differentiation
Job alerts	High volume and weak personalization	Structured filtering, deduplication and evidence-based analysis
Generic AI chat	No persistent verified career context	Grounded retrieval from an approved career knowledge base
Auto-apply bots	Prioritize volume and may submit inaccurate content	Approval-first workflow with no automated submission
Spreadsheets / Notion trackers	Useful for tracking but not intelligent preparation	Integrated analysis, document generation, reminders and interview support
Resume tools	Focus on one artifact	End-to-end lifecycle from discovery to offer evaluation

3. User discovery and experience

3.1 Primary persona

Attribute	Description
Persona	Technical professional transitioning toward AI Product Management / AI Solution Architecture
Context	Applies to remote or international roles while managing a broad technical and product background
Goals	Find relevant opportunities, present the most applicable experience and improve interview readiness
Pain points	Repetitive searching, fragmented records, resume rewrites, generic cover letters and inconsistent tracking
Trust concerns	Does not want AI to invent achievements, submit applications or expose confidential career data
Preferred interaction	Concise Telegram decisions with secure web previews for detailed review

3.2 Jobs to be done

Situation	Job to be done	Desired outcome
When new roles are published	Find and filter jobs that match my target roles and constraints	Review only the most relevant opportunities
When I am evaluating a role	Explain how my verified experience matches the requirements	Make an informed apply / skip decision
When I decide to apply	Prepare a truthful, role-specific resume and cover letter	Submit a strong application without rewriting everything
After I apply	Track status and remind me of appropriate next actions	Avoid losing opportunities through poor follow-up
Before an interview	Turn my real project experience into relevant preparation	Answer confidently without inventing stories
When I receive an offer	Normalize, compare and prepare a negotiation response	Make a deliberate decision with clear trade-offs

3.3 Key pain points and product responses

Pain point	Product response
Too many irrelevant jobs	Configurable search profiles, deterministic filters and deduplication
Unclear job fit	Requirement extraction and evidence-backed match mapping
Generic application documents	Role-specific drafts generated from approved career evidence
Fear of AI hallucination	Claim provenance, truth audits and human approval gates
Scattered application tracking	Application CRM, status history, tasks and reminders
Weak interview recall	Evidence-grounded STAR outlines and preparation packs
Unclear offer value	Normalized compensation, missing-term detection and comparison

3.4 End-to-end user journey

Stage	User action	System action	Trust mechanism
Discover	Sets search profile	Collects, normalizes and filters jobs	Source transparency and no restricted scraping
Evaluate	Queues selected jobs	Maps requirements to career evidence	Citations, gaps and deterministic score
Prepare	Approves tailoring	Creates resume and letter drafts	Preview, diff, truth audit and versioning
Apply	Submits manually	Creates application record after confirmation	No automated external action
Follow up	Approves a draft	Creates a concise message	No automatic sending
Interview	Requests preparation	Generates questions and STAR outlines	Evidence provenance and confidentiality controls
Offer	Records and confirms terms	Evaluates and drafts negotiation	Source references and explicit final confirmation

4. Product Requirements Document (PRD)

4.1 Product objective

Build a self-hosted, end-to-end career intelligence platform that automates repetitive work, produces grounded recommendations and documents, and maintains explicit human approval for every external or irreversible action.

4.2 MVP definition

MVP outcome

A user can upload an approved resume, configure target roles, receive relevant job alerts, analyze a selected job against verified career evidence, approve a tailored resume and cover letter, and receive a ready-to-submit application package.

4.3 Functional requirements

ID	Capability	Requirement	Priority
FR-01	Career knowledge ingestion	Upload PDF, DOCX, TXT or Markdown; extract, chunk, embed and index with metadata.	Must
FR-02	Master resume approval	Create a structured master resume draft and require explicit approval before tailoring.	Must
FR-03	Job discovery	Collect jobs from supported public sources, email alerts and manual submission.	Must
FR-04	Normalization and deduplication	Normalize source payloads and prevent duplicate job records and alerts.	Must
FR-05	Preliminary filtering	Apply target title, keyword, location, remote and seniority rules.	Must
FR-06	Job match analysis	Extract requirements, retrieve evidence and calculate a transparent 0-100 match score.	Must
FR-07	Human interest decision	Allow Save, Ignore, Analyze and Interested actions through Telegram.	Must
FR-08	Resume tailoring	Generate a structured, ATS-safe draft only after user approval.	Must
FR-09	Resume preview and revision	Show diff, truth report, ATS report, DOCX and PDF before approval.	Must
FR-10	Cover letter generation	Create full letter or short application note after explicit approval.	Must
FR-11	Application package	Bundle approved resume, letter, links and metadata without submitting.	Must
FR-12	Application tracking	Record manual submission, status history, notes and follow-up reminders.	Should
FR-13	Interview preparation	Generate role-specific questions and evidence-backed STAR outlines after approval.	Should

ID	Capability	Requirement	Priority
FR-14	Offer evaluation	Record and confirm terms; normalize cash, equity and benefits; identify missing information.	Could
FR-15	Negotiation drafts	Generate approved, evidence-backed negotiation content without sending.	Could

4.4 Non-functional requirements

Category	Requirement
Security	HTTPS, authenticated endpoints, internal API keys, Telegram allowlist, protected files, SSRF and path traversal prevention.
Privacy	Local parsing and embeddings where practical; minimum necessary model context; no prompt logging by default.
Reliability	Idempotent callbacks, retries for transient failures, lock expiry, health checks and backups.
Explainability	Every material positive claim links to approved source evidence.
Performance	Personal-use workloads should support daily job ingestion, up to 20 analyses and limited concurrent document generation.
Cost	Tiered model routing, token limits, daily/monthly budgets and no automatic Premium escalation.
Maintainability	Centralized settings, model abstraction, prompt versioning, database migrations and exportable n8n workflows.
Accessibility	Readable, selectable documents with standard headings and ATS-safe layouts.

4.5 Key user stories and acceptance criteria

ID	User story	Acceptance criteria
US-01	As a job seeker, I want a daily list of relevant roles so I do not spend hours searching.	Only active profiles are used; duplicate jobs are not re-notified; each result includes source and preliminary reason.
US-02	As a user, I want to understand why a role matches me before I apply.	Analysis contains score, recommendation, strengths, gaps, confidence and valid evidence references.
US-03	As a user, I want AI to ask before changing my resume.	Tailoring cannot start without an approved master resume and explicit TAILOR decision.
US-04	As a user, I want to preview every generated document.	Preview, change summary and truth status are available before approval; approved version is immutable.
US-05	As a user, I want to control application submission.	System has no automated submission path; application becomes APPLIED only after explicit confirmation.
US-06	As a user, I want trustworthy interview preparation.	Every candidate claim is linked to approved career evidence; missing metrics are called out rather than invented.
US-07	As a user, I want to evaluate offers without misleading calculations.	Guaranteed, target and one-time compensation are separated; equity is never presented as guaranteed cash.

4.6 Scope boundaries

In scope	Out of scope
Public job sources, email alerts and manual links	Authenticated scraping or CAPTCHA bypass
Evidence-grounded drafts with approval	Unsupervised resume, cover-letter or message sending
Manual application status confirmation	Automatic form filling or Easy Apply
Offer analysis and negotiation support	Legal, tax, immigration or financial advice
Personal self-hosted deployment	Multi-tenant billing, public marketplace or enterprise admin

5. AI product and solution design

5.1 Human-in-the-loop decision architecture

The product distinguishes between reversible automation and consequential action. The system may search, classify, retrieve, calculate, draft and remind automatically. It may not submit, send, schedule, accept, decline or represent an unverified fact without explicit authorization.

Action type	Automation level	Example
Low-risk, reversible	Automatic	Normalize a job title, detect duplicates, retrieve relevant evidence
Interpretive	Automatic with explanation	Generate a match report or ATS warning
Content creation	Draft then approve	Tailored resume, cover letter, follow-up or negotiation message
External action	Manual only	Submit application, send email, message recruiter
Irreversible decision	Two-step confirmation	Accept or decline an offer

5.2 RAG and career evidence

The career knowledge base separates raw documents, structured resume facts and vectorized evidence. Job requirements are converted into targeted retrieval queries. Retrieved chunks are deduplicated, reranked and classified as direct, transferable, educational, weak or missing evidence. The generation layer receives only the smallest relevant context rather than the entire career history.

5.3 Model routing

Tier	Primary use	Control
Nano	Classification, extraction, normalization and lightweight metadata	Default for high-volume low-complexity tasks
Mini	Job analysis narrative, resume/letter drafting, interview and offer synthesis	Default generation tier
Premium	Optional final review for high-value cases	Disabled by default; explicit approval; never an automatic fallback

5.4 Architecture overview




Figure 2. Product architecture and service responsibilities.

5.5 Core product decisions and trade-offs

Decision	Why	Trade-off
Self-host n8n on AWS VPS	Lower recurring platform cost and direct control of workflows	Requires operational maintenance and security ownership
Use local embeddings	Reduces cost and limits external exposure of career documents	Uses VPS memory and may be less accurate than premium embedding services
Separate PostgreSQL and Qdrant	Structured lifecycle data and semantic evidence require different storage models	More operational components
Keep submission manual	Protects professional reputation and avoids platform-policy risk	The user still performs the final application step
Generate structured documents first	Supports validation, provenance, consistent rendering and versioning	More engineering work than free-form text generation
Use Telegram as primary interface	Fast decisions and notifications with minimal UI scope	Detailed review still needs secure web previews

6. Roadmap and delivery plan




Figure 3. The roadmap separates foundational risk from progressively higher-value workflows.

6.1 Phase outcomes

Phase	Theme	Definition of done
0	Infrastructure foundation	Secure Docker deployment, Nginx/TLS, PostgreSQL, Qdrant, Redis, n8n and backups
1	Career knowledge base	Document ingestion, local embeddings, RAG retrieval and model routing
2	Job discovery	Source adapters, search profiles, normalization, deduplication and Telegram digest
3	Job match analysis	Requirement extraction, evidence mapping, deterministic score and recommendations
4	Resume tailoring	Master resume approval, truth/ATS validation, preview, revision and export
5	Cover letter and package	Full letter / short note, approval, secure files and ready-for-application package
6	Application and interview	CRM, reminders, follow-up drafts, interview records and preparation
7	Offer and negotiation	Offer parsing, evaluation, comparison and approved negotiation drafts
8	Production hardening	Security, observability, backups, E2E tests, cost controls and release readiness

6.2 Prioritization rationale

The roadmap follows a dependency-first sequence. Career evidence and trustworthy retrieval are built before match analysis; analysis is built before tailored content; approved documents are built before tracking and interview workflows. Offer intelligence is delayed until the core application loop works. Production hardening is a distinct final phase to avoid mixing reliability work with feature delivery.

6.3 MVP versus later scope

Release	Included	Deferred
MVP	Phases 0-5: discover, analyze, approve documents and prepare application package	Tracking, interview, offer intelligence and advanced operations
Personal beta	Phases 0-6 with reminders and interview preparation	Offer comparison, negotiation and full hardening

Release	Included	Deferred
Production-ready personal release	Phases 0-8 with security, observability, cost limits and tested recovery	Commercial multi-user features

7. Metrics and evaluation

7.1 North-star metric

North-star metric

High-quality application packages approved per week for jobs the user intentionally selected.

This metric combines product value and user intent. It avoids rewarding raw job volume, model output volume or automated applications. A package is counted only when the job was analyzed, the user expressed interest, the documents passed truth checks and the user approved the final materials.

7.2 Product metrics

Metric group	Metric	Why it matters
Discovery	Relevant jobs surfaced / total jobs collected	Measures filtering quality and reduces review burden
Analysis	User agreement with Apply / Skip recommendation	Measures decision support quality
Grounding	Unsupported material claim rate	Primary trust and safety quality metric
Preparation	Median time from Interested to approved package	Measures workflow efficiency
Documents	Resume and cover-letter approval rate	Measures first-draft quality
Engagement	Jobs reviewed and decisions completed	Measures whether alerts lead to action
Tracking	Follow-ups completed before due date	Measures operational value after submission
Outcome	Interview and offer rates by match band	Helps evaluate scoring, with sample-size warnings
Operations	Workflow success rate and recovery time	Measures reliability
Cost	Average AI cost per approved application package	Keeps the product economically sustainable

7.3 Guardrail metrics

- Human approval-gate bypass rate: target 0%.
- Invalid evidence citation rate: target 0%.
- Automatic Premium escalation: target 0%.
- Confidential data exposure in notifications or logs: target 0 incidents.
- Duplicate job notification rate: below an agreed threshold.
- Applications marked APPLIED without explicit confirmation: target 0.
- Monthly AI spend must remain below the configured hard limit.

7.4 Evaluation approach

Evaluation	Method
Requirement extraction	Curated job descriptions with labeled required and preferred requirements
Evidence mapping	Manual review of direct, transferable, educational, weak and missing classifications
Score consistency	Repeat analysis with identical inputs and compare deterministic outputs
Resume truthfulness	Validate generated claims against approved master resume and source chunks
Prompt injection resistance	Adversarial content embedded in jobs, resumes, company pages and offers
Document usability	ATS parsing checks, selectable PDF text and human review
End-to-end reliability	Mocked E2E journey from job ingestion to offer decision

7.5 Target metrics for beta

Target	Initial beta goal
Relevant jobs in daily digest	At least 70% judged worth reviewing by the user
Unsupported material claims	0 approved outputs
First-draft resume approval	At least 60% with minor or no revision
Median application preparation time	Under 20 minutes of user review time
Daily workflow success	At least 95% excluding external source outages
AI cost per approved package	Within user-defined budget, measured rather than assumed

8. Risk management, privacy and launch

8.1 Product risk register

Risk	Severity	Mitigation
Hallucinated experience	High	Truth audit, provenance, approved master resume, blocking validation
Job-source policy violation	High	Use public APIs/RSS/email alerts; prohibit authenticated scraping and CAPTCHA bypass
Sensitive data exposure	High	Local parsing/embeddings, context minimization, protected files and redacted logs
Automatic action without consent	High	Explicit approval gates and release-blocking tests
Duplicate or stale jobs	Medium	Canonical URLs, content fingerprints, update rules and expiry jobs
Model availability changes	Medium	Logical model tiers, environment-configured model IDs and validation script
Unexpected AI cost	Medium	Token limits, daily/monthly budgets, no expensive fallback and warnings
VPS outage or data loss	Medium	Health checks, automated backup, restore runbook and integrity checks

Risk	Severity	Mitigation
Overreliance on match score	Medium	Explain score components, confidence and missing evidence; keep user decision separate
Biased or incomplete recommendations	Medium	Transparent evidence, user feedback, evaluation set and conservative language

8.2 Privacy design

- Resume and project documents are treated as personal or confidential data.
- Text extraction and embeddings run locally where practical.
- Only the minimum retrieved evidence is sent to an external chat model.
- Prompts and responses are not logged by default.
- Generated files are stored in protected paths with expiring download tokens.
- Confidential employer names use approved generic labels.
- The user can export or delete records, with documented backup limitations.

8.3 Launch plan

Stage	Users	Focus	Exit criteria
Internal alpha	Single owner	Functional flow and data correctness	Core Phase 0-5 journey works with mocked and real sample jobs
Personal beta	Single owner in daily use	Relevance, document quality and Telegram usability	Stable weekly use, no unsupported approved claims, manageable cost
Production-ready personal release	Single owner	Security, backup, observability and recovery	Phase 8 release-blocking checks pass
Future private beta	Small invited cohort	Multi-user privacy and onboarding discovery	Requires separate product strategy and tenant isolation

8.4 Rollout checklist

- Approve the canonical master resume and career evidence.
- Configure at least one search profile and verify job sources.
- Verify Telegram bot identity, allowlist and callback security.
- Run a complete mocked end-to-end test.
- Run an actual low-volume job analysis and document-generation smoke test.
- Confirm daily and monthly AI cost limits.
- Confirm PostgreSQL, Qdrant and file backups; test restore in an isolated environment.
- Review production-readiness report and resolve all release-blocking findings.

9. Product decision log

Decision	Selected approach	Rejected / deferred approach	Reason
Primary user interface	Telegram + secure web previews	Full custom frontend	Faster MVP while preserving detailed review
Application submission	Manual	Auto-submit	Trust, accuracy and platform-policy safety
Career memory	Qdrant + approved structured resume	Prompting from one large resume file	Better retrieval, provenance and reuse
Workflow engine	n8n	Custom scheduler from day one	Visual orchestration and faster integration delivery
AI model strategy	Tiered OpenRouter models	Single premium model for all tasks	Cost control and fit-for-purpose routing

Decision	Selected approach	Rejected / deferred approach	Reason
Document generation	Structured JSON then deterministic rendering	LLM-generated binary or unstructured document	Validation, versioning and layout consistency
Deployment	Self-hosted AWS VPS	Managed n8n and separate SaaS services	Cost control and ownership; accepts operational burden

10. Portfolio case study narrative

10.1 Challenge

I wanted to make my job search more systematic without using a mass-application bot. The process was repetitive, but the most important decisions - whether a role was right, how to position my experience and what to send - still needed human judgment. I defined the product as a human-in-the-loop career intelligence platform rather than an autonomous applicant.

10.2 My role

I worked across product strategy, workflow design and solution architecture. I framed the problem, defined the product principles, separated the roadmap into dependency-driven phases, specified approval gates and designed the AI architecture around evidence retrieval, model routing and deterministic validation. I used Codex as an implementation partner and n8n as the workflow orchestration layer.

10.3 Product approach

1. Start with trusted career evidence, not job automation. The system first needs an approved source of truth for experience, projects and skills.
2. Automate discovery and analysis before document generation. A role should be worth applying to before the system spends time tailoring content.
3. Treat every generated claim as data with provenance. A polished sentence is not acceptable if its source cannot be verified.
4. Use human approval gates for consequential actions. The platform prepares and recommends; the user submits, sends and decides.
5. Deliver in phases so infrastructure, data quality and trust controls are solved before advanced features.

10.4 Key design outcome

The final product concept covers the complete career journey: discovery, fit analysis, resume and cover-letter preparation, application tracking, follow-up, interview preparation, offer evaluation and negotiation support. The distinguishing feature is not the number of agents; it is the trust architecture that connects every recommendation and generated claim to approved evidence and a human decision.

10.5 Intended value and measurement

The project is currently positioned as an in-development portfolio and personal productivity product. The intended value is reduced preparation time, higher relevance and stronger consistency. These outcomes must be measured through real usage; the portfolio deliberately avoids presenting target metrics as achieved results.

10.6 What this project demonstrates

Product capability	Evidence in the project
Product strategy	Vision, target user, positioning, principles, non-goals and phased scope
Product craft	End-to-end user journey, approval states, previews, versioning and acceptance criteria
AI product judgment	RAG grounding, model routing, prompt-injection defenses and truth audits

Product capability	Evidence in the project
Technical fluency	AWS, Docker, n8n, FastAPI, PostgreSQL, Qdrant, Redis and OpenRouter architecture
Risk management	Privacy, platform-policy boundaries, cost limits, backups and release-blocking checks
Outcome thinking	North-star metric, guardrails, evaluation dataset and rollout plan

11. Future opportunities

- A lightweight web dashboard for job review, document comparison and application analytics.
- Learning from user feedback to improve search profiles and match-weight configuration.
- A verified skills-gap plan that recommends portfolio projects or learning activities without overstating capability.
- Additional approved job-source adapters and company career-page monitoring.
- Multi-language application support with independent truth and style validation.
- A future multi-user product only after tenant isolation, privacy, onboarding and commercial validation are redesigned.

Appendix A. Example match report

Field	Example output
Role	AI Product Manager at Example Company
Overall match	84 / 100 - APPLY
Strongest direct evidence	AI customer assistant, fraud detection, fintech and enterprise backend architecture
Transferable evidence	Senior Java delivery mapped to technical product and solution architecture requirements
Missing evidence	Large-scale LLM evaluation and formal MLOps ownership
Recommended positioning	Emphasize the combination of AI product ownership and enterprise engineering depth
Confidence	88%, based on evidence coverage and requirement clarity

Appendix B. Example resume approval flow

6. The user marks an analyzed job as Interested.
7. The system asks whether to tailor the resume, use the master resume or skip.
8. If Tailor is selected, the system generates a structured draft from approved evidence.
9. Truth and ATS validation run before the draft becomes preview-ready.
10. The user reviews the preview, change summary and unsupported keywords that were intentionally excluded.
11. A revision creates a new version; approval makes the selected version immutable.
12. Only the approved version is available to the cover-letter and application-package workflow.

Appendix C. Portfolio summary

Portfolio-ready description

Designed a self-hosted AI Career Intelligence Agent that automates job discovery, evidence-based role matching, resume and cover-letter preparation, application tracking, interview readiness and offer evaluation. The product combines n8n orchestration, FastAPI services, PostgreSQL, Qdrant RAG, local embeddings, OpenRouter model routing and Telegram approval workflows. I structured delivery across eight phases and designed truth validation, provenance, privacy, cost controls and human approval gates so the platform improves job-search efficiency without sacrificing credibility or user control.

END OF PORTFOLIO PACK