

Amando Martínez Vivancos

Backend · Data & AI Engineer

Spain · CET amandomvwork@gmail.com linkedin.com/in/amando-martínez-vivancos
amandomv.dev

PROFILE

I'm a backend engineer who ended up in data and AI because that is where the interesting failure modes are. I joined Datamaran as a junior in 2022 and grew into the senior role there, which in practice meant owning systems end to end: the pipelines that read 20,000 news items a day and cover 193 legal sources across 130 jurisdictions — one of them 85% cheaper to run than when I inherited it, the other saving the legal team 70% of its report-processing time — plus the auth and sharing layers underneath the product and the observability that tells us when any of it breaks.

I build and automate with AI in the loop — not as a demo, but as the way the work gets done. That means LLM agents wired into real pipelines (a scraper that repairs itself instead of paging someone), evaluation harnesses so a model or prompt change is an experiment with a number attached, and AI-assisted development spread across repositories and delivery processes.

EXPERIENCE

Joined as Software Engineer (Oct 2022) · promoted to Senior in January 2026

- Built the news intelligence pipeline that ingests, deduplicates, clusters and classifies 20,000+ news items a day with LLMs on Airflow 3 over Kubernetes — and cut its annual LLM bill by 85% — a six-figure sum down to a five-figure one — by ordering stages by cost and moving them to smaller models once evaluation showed no quality loss.
- Built the regulatory pipeline's discovery layer: a generator that writes fetchers behind a seven-stage deterministic harness and a human review, plus the health job that detects silent breakage and hands the worst sources back to it. 193 of ~290 catalogued sources across 130 jurisdictions are covered by 153 modules, and report processing takes the legal team 70% less time.
- Shipped the application's user management and authentication system, and the resource-sharing layer that makes live collaborative work possible across accounts.
- Introduced and spread the practices a 20-developer engineering org now works by: DDD, hexagonal architecture, TDD, shared linting and formatting standards, and proper database connection-pool management.
- Led the adoption of AI-assisted development — agents and evaluation harnesses wired into pipelines, repositories and delivery processes, not just editor autocomplete.
- Owns observability for the systems I build: Grafana dashboards and Sentry alerting, so failures surface as signal rather than as a customer email.
- Works directly with stakeholders to turn customer needs into features, from scoping through to production.

SELECTED PROJECTS

News intelligence pipeline

Twenty thousand items a day flow through unattended, and the pipeline's annual LLM bill came down by 85% — from a six-figure sum to a five-figure one. Clustering is what did most of it: paying once per story instead of once per article, then letting evaluation move stages onto smaller models wherever the numbers said quality held.

Python 3.12 · Airflow 3 · Kubernetes · pgvector · PostgreSQL · MongoDB · LLMs · AWS

Self-healing regulatory pipeline

193 of the catalogued sources are covered, 153 fetcher modules live in the repository, and report processing takes the legal team 70% less time than before. A daily health job marks a fetcher broken after three consecutive failures, or degraded after three zero-yield runs against its own positive baseline, and hands the worst two to the healer — which opens a reviewed pull request instead of paging anybody.

Python 3.12 · Airflow 3 · Kubernetes · Anthropic LLMs · PostgreSQL · MongoDB · Docling

Gestión Guarda Rural

The work gets recorded once, where it happens, and monthly invoicing went from an afternoon of manual work to a few clicks — including multi-recipient client notifications and a PDF invoice layout that survives contact with a real printer.

Python · PostgreSQL · Neon · Railway · Vue

Fantasy Voley

A working league built end to end — ingestion, scoring engine and front end — kept in sync automatically through the season, on a data model that survived a federation site redesign without a rewrite. Public launch is next.

Python · PostgreSQL · Neon · Railway · Vue

SKILLS

Languages: Python · TypeScript · JavaScript · SQL

Backend:

FastAPI · asyncio · SQLAlchemy · Pydantic · REST API design · Connection pooling

Building & automating with AI:

AI-assisted development (agentic workflows, Claude Code) · Self-healing pipelines driven by LLM agents · Evaluation harnesses for LLM stages · Prompt and model cost/accuracy tuning · Custom tooling & MCP integrations

Data platform:

Airflow 3 · Kubernetes · Large-scale scraping · Clustering & deduplication · ETL

Architecture: Domain-Driven Design · Hexagonal architecture · TDD · SOLID

Storage: PostgreSQL · MongoDB · Elasticsearch · Redis

Platform & observability: AWS · Docker · GitHub Actions · Grafana · Sentry

Front end: Vue 3 · Astro · Tailwind CSS

EDUCATION

BSc in Computer Engineering — Computation specialisation

Universitat Politècnica de València · 2018–
2022

LANGUAGES

Spanish: Native

English: Professional working proficiency

Last updated: 2026-09