Mario Adrián Chasco Frauca

AI Engineer | LLM products, agents and automation systems

Pamplona, Spain · Remote, CET · mariochasco@flowaraai.com · flowaraai.com
github.com/mariochasco78557 · linkedin.com/in/mario-adrian-chasco-frauca · youtube.com/@mariochasco_AI

ABOUT ME

I build and operate AI systems that run unattended for small businesses. A hair salon in Pamplona has been using my voice and WhatsApp booking agent since September 2026: 407 calls and 212 bookings by late October. I trained as a mechanical engineer, so I think first about how a system can fail and I test it before I trust it. How I work: code decides; the model reads and drafts. Flowara AI is my public work: the systems I built for real businesses and the templates I extracted from them.

SELECTED AI SYSTEMS

Voice and WhatsApp booking agent

Hair salon in Pamplona · in production since September 2026

Answers the salon's phone in the owner's cloned voice and books into its real appointment system; a WhatsApp agent on Meta's official API handles confirmations, changes and reminders. The model only picks from options the code gives it, and every tool call must pass a strict Zod schema.

Numbers: 407 calls and 212 bookings by late October 2026. First two days: 63 calls, 34 bookings, about 1.5 min per call, 0 technical errors. Production system: over 5,700 tests. Public template: 4,767 + 277 tests, all passing.

Stack: TypeScript, Fastify, Zod, ElevenLabs, Meta WhatsApp Cloud API, Docker
github.com/mariochasco78557/ai-booking-agent-template

Read-only crypto analytics platform

Investment firm in London · since June 2026

Pulls exchange history and live account state through read-only keys, reconciles it against the ledger and shows Sharpe, Sortino, Calmar and drawdown. It never trades, withdraws or transfers.

Numbers: 1.26M rows ingested, live capture every minute. Commissions matched the ledger 117 of 117 days. Two independent sources agree within 0.2% over 40 days.

Stack: Python, FastAPI, PostgreSQL, Next.js, Docker
Private, demo on request

AI Knowledge Feeds and MCP server

Open source (MIT) · runs daily on a Debian VPS

Syncs official docs, changelogs and video transcripts every morning and serves ranked fragments, never whole documents, to coding agents over MCP. Official docs rank above opinion.

Numbers: 31 sources synced daily at 08:00. Router right first time went from 44% to 70% on a golden set, 80% within two picks. Injects at most 1,200 characters in 0.25 s.

Stack: Python, BM25, MCP, yt-dlp, pytest
github.com/mariochasco78557/ai-knowledge-feeds

Self-hosted meeting transcription stack

Public template · runs 24/7 on my own server

A SQLite queue with explicit job states runs transcription, speaker diarization and an LLM summary, then sends a Markdown note back to Drive. The model only summarises; code decides where the note goes and when to retry.

Numbers: 21 modules, 161 tests including a golden-file test. 1,300-line architecture doc. Memory and speed benchmarked on a small VPS.

Stack: Python, SQLite, faster-whisper, sherpa-onnx, LLM APIs, Google Drive API
github.com/mariochasco78557/transcription-stack-template

Claude Code toolkit

Open source (MIT) · the setup I build everything else with

Hooks, handoff skills and subagents picked by task complexity; a verifier reviews each change without editing it. A Stop hook blocks any turn that promises work without starting it. I added it after my server sat idle for 3 hours.

Numbers: 4-tier subagent routing (junior, plumber, senior, verifier). Per-prompt router with --dry-run. Every generated change goes through tests and review.

Stack: Claude Code hooks, subagents, skills
github.com/mariochasco78557/claude-code-toolkit

ALSO BUILT

- Unattended video pipeline for a travel agency: cloned-voice narration, ffmpeg assembly and delivery to the client's Google Drive. A new video every 3 days since October 2026.
- Lead intake system: web form, LLM classifier, router and Google Sheets. Client websites in Astro and Next.js on Vercel.
- Privacy-first AI secretary for teachers: public template on synthetic data, with a fail-closed privacy gate and 11 unit tests.

EXPERIENCE

Founder and AI Engineer · Flowara AI — Pamplona, Spain / Remote · 2025 to present

- AI automation studio for small businesses. I own the full cycle: discovery, architecture, build, deployment on my own VPS, monitoring and support.
- I build with Claude Code and Codex as pair programmers; every generated change goes through tests and review.
- 38+ short build-in-public videos on YouTube.

Backend and DevOps Collaborator · Investment firm based in London — Remote · June 2026 to present

- Built and operate the read-only analytics platform above: ingestion, reconciliation, live capture, dashboard, Docker deployment, PostgreSQL, logging and validation.

AI Engineer Trainee · Selective remote AI engineering program — Remote · Aug 2025 to Sep 2025

- Selected at about 10% acceptance. Delivered three end-to-end projects with n8n, LLMs and API integrations: (1) a Google Workspace personal assistant (Gmail, Calendar, Contacts), (2) a RAG chatbot over a movie script, (3) a natural-language-to-SQL marketing analyst.
- In the last challenge I linked Claude Code to n8n over MCP. That was my first time using Claude.

Technology Teacher · Government of Navarra — Pamplona, Spain · Sep 2023 to present

- I teach AI, cybersecurity and programming in English, to mixed levels.

EARLIER ENGINEERING

Mechanical Engineer · Berry Global: Computer-vision quality inspection (INTRAVIS, In-Sight) and commissioning of robotic systems on molding machines. Good parts were rejected when they drifted out of the camera's view; I proposed adding AI, and the company declined. 2018 to 2020

Laboratory Technician · Schneider Electric: Product compliance testing to UNE and IEC standards. 2016 to 2017

Engineering Intern · CENER: Python and MATLAB multibody code for floating offshore wind turbines. 2016

EDUCATION

Master's Degree in Secondary Education and Teaching · Universidad de Navarra — 2021

Master's thesis (June 2021): motion of a projectile in a dense granular medium. I wrote MATLAB and Python code to fit its equation of motion and validated it against simulation data from the university's physics department, plus an IB Physics teaching proposal.

MSc in Astronomy and Astrophysics · VIU, Universidad Internacional de Valencia — 2017 to 2019

Master's thesis: "Automatic classification of stellar spectra" (2019).

Bachelor's Degree in Mechanical Engineering · Universidad Pública de Navarra (UPNA) — 2012 to 2016

Final degree project (June 2016): reduced dynamic model for floating offshore wind turbines. I extended a MATLAB multibody code from 2 to 3 degrees of freedom on a CENER code base, co-supervised by UPNA and CENER.

TECHNICAL SKILLS

LLMs and agents	Anthropic and OpenAI APIs, tool calling with strict schemas (Zod), agentic workflows, MCP servers, prompt and context design, evals and golden sets, n8n, Claude Code, Codex
Retrieval and data	RAG with BM25 and fragment-level routing, ingestion and reconciliation pipelines, PostgreSQL, SQLite, Prisma
Voice and messaging	ElevenLabs, Twilio, Pipecat, Deepgram, Meta WhatsApp Cloud API, faster-whisper, sherpa-onnx
Backend and infra	Python, FastAPI, TypeScript, Node.js (Fastify), Next.js, Tailwind, Docker, Caddy, systemd, Debian VPS, Vercel, Google Workspace APIs, yt-dlp
Testing	pytest, Vitest, golden-file tests, health checks

LANGUAGES

- **Spanish:** Native.
- **English:** Cambridge C1 Advanced (CAE), working at C2 level; I teach in English.

HOW I BUILD

- Least AI possible: code decides identity, money, routing and sending; the model reads and drafts.
- Autonomy is earned per task type and proven on a golden set of real cases, with code checks first.
- No silent failures: log, counter, canary.
- A named human owns every agent, and every mistake becomes a rule in a file.