

Javier Pontón González

Forward Deployed Engineer

I build AI systems that survive an audit, not just a demo.

Email: javierpontongonzalez@gmail.com

Phone: +34 623 920 307

Location: Asturias, Spain

Time zone: Europe/Madrid

Remote: Fully remote (EU, UK, US)

Availability: Available immediately

Engagement: Independent Consultant (Freelance / B2B Contractor) via Esguil, since November 2023 (through Lando Firm, S.L. since 2026), invoiced from Spain

Work authorisation: EU citizen. Invoices via Esguil (through Lando Firm, S.L. since 2026) as a Spanish B2B contractor; no visa sponsorship required for EU work.

LinkedIn: <https://www.linkedin.com/in/javierpontongonzalez>

GitHub: <https://github.com/sktjpg>

Website: <https://javierpontongonzalez.com>

Book a call: <https://cal.com/javier-ponton-gonzalez-qivh2p>

PROFESSIONAL SUMMARY

Forward Deployed Engineer and AI Engineer. I embed with the customer team, scope the use case against their real data, build on their infrastructure, and stay through production and handover. Seven years of production backend engineering for Iberia, KPN, Mercadona, Inditex, Openbank and Idealista sit under that, plus 2+ years shipping Generative AI to production: AI agents and multi-agent systems (LangChain, LangGraph, MCP, function calling), Retrieval-Augmented Generation with embeddings, hybrid search and cross-encoder reranking, structured extraction with vision-language models, and model serving on AWS Bedrock and Google Vertex AI as readily as fully air-gapped on-premise inference.

What separates my work from prototypes: every system ships with an evaluation harness (golden datasets, field-level precision and recall, regression testing), OpenTelemetry observability, source-level traceability and human-in-the-loop review. Core stack: Python, Java, Kotlin, FastAPI, Spring Boot, PostgreSQL and pgvector, Kafka, AWS, Docker. Independent Consultant (Freelance / B2B Contractor) since 2023 via Esguil (through Lando Firm, S.L. since 2026), with no gaps between engagements, and founder-engineer of four live products.

KEY ACHIEVEMENTS

- Forward Deployed Engineer on two enterprise engagements that ran in parallel: technical owner of the fare pricing domain at Iberia, and sole engineer on an air-gapped multi-agent tender-response platform delivered inside a Saudi systems integrator isolated network, scoring 90%+ field-level precision and recall on the golden dataset before go-live.
- Built a production Model Context Protocol (MCP) server and a RAG service over engineering documentation at Iberia, giving the team governed function calling and natural-language access to the flight pricing platform, and deployed managed model serving on AWS Bedrock and Google Vertex AI to production alongside self-hosted inference, instrumented with OpenTelemetry and Dynatrace.
- Led a zero-downtime Azure to AWS migration of a global logistics platform at Inditex and a Java to Kotlin migration at Openbank inside a 100+ engineer organisation; shipped a consumer app (ZORRO) solo to iOS and Android, reaching 700+ users with 50% D1 retention in three weeks.

TECHNICAL SKILLS

AI & LLM	Large Language Models (LLMs), Generative AI, Applied AI, Machine Learning, Transformers, Vision-Language Models (VLMs), NLP, Prompt Engineering, Structured Outputs, JSON Schema, Function Calling, Tool Use, Hallucination Mitigation, Grounding
AI Agents & RAG	Retrieval-Augmented Generation (RAG), AI Agents, Agentic AI, Agentic Workflows, Multi-Agent Systems, LangChain, LangGraph, Model Context Protocol (MCP), Embeddings, Semantic Search, Hybrid Search, Cross-Encoder Reranking, Vector Databases, pgvector, HNSW
LLM Evaluation & LLMOps	LLM Evaluation, Golden Datasets, Regression Testing, Precision, Recall, F1, LLM-as-a-Judge, Human Evaluation, RAG Evaluation, Hit Rate, MRR, A/B Testing, Shadow Evaluation, Guardrails, Human-in-the-Loop, RAGAS, Promptfoo, LLMOps, OpenTelemetry, Prometheus, Grafana, Dynatrace, Cost Optimization, Latency Optimization
Model Serving & AI Infrastructure	AWS Bedrock, Google Vertex AI, vLLM, Ollama, mistral.rs, OpenAI-Compatible APIs, HuggingFace, Quantization, GGUF, AWQ, GPU/VRAM Optimization, Self-Hosted Inference, On-Premise AI, Air-Gapped Deployment, Qwen3, Qwen2.5-VL, DeepSeek, Claude API, OpenAI API
Backend & Data	Python, Java, Kotlin, Elixir, TypeScript, SQL, Spring Boot, FastAPI, Spring Batch, REST, OpenAPI, Microservices, PostgreSQL, pgvector, MongoDB, Elasticsearch, Redis, DynamoDB, Apache Kafka, AWS SQS
Cloud & DevOps	AWS, Azure, GCP, Docker, Kubernetes, GitHub Actions, Jenkins, Git, CI/CD
Architecture & Engineering	Hexagonal Architecture, Domain-Driven Design (DDD), CQRS, Event Sourcing, SOLID, TDD, API-First, Technical Leadership, Mentoring
Forward-deployed delivery	Forward Deployed Engineering, Customer-embedded delivery, Use-case scoping against real data, Build-versus-buy assessment, Technical discovery with non-engineering stakeholders, Solution architecture, Production operation and handover, Regulated environments (GDPR, Saudi PDPL, NCA, IATA NDC), Direct contracting, no agency layer

PROFESSIONAL EXPERIENCE

Esguil (own AI engineering consultancy) (Lando Firm, S.L. since 2026) – Forward Deployed Engineer

Nov 2023 – Present · Remote, Europe and Middle East

Independent B2B since November 2023 through my own consultancy, Esguil (through Lando Firm, S.L. since 2026). Five back-to-back enterprise engagements with no gaps and no agency layer, contracted directly with the client.

- End-to-end ownership.** Own each engagement end to end from inside the client team: use-case scoping against real data, build-versus-buy assessment, architecture, delivery, production operation and handover, across Spain, the Netherlands and Saudi Arabia.
- Stakeholder discovery.** Run technical discovery directly with non-engineering stakeholders (finance, operations, legal and compliance), translating regulatory and audit constraints into system design decisions.
- Practice inside client teams.** Standardise engineering practice inside client teams: testing strategy, hexagonal architecture and DDD conventions, and AI-assisted development guidelines adopted by the engineering group.

Confidential (IT systems integrator) (client under NDA) – Forward Deployed Engineer

Jul 2026 – Aug 2026 · Remote, Saudi Arabia

Air-gapped multi-agent AI platform for public-tender (RFP) response. Sole engineer on the engagement, which ran in parallel with the Iberia engagement below; both have since ended.

- **Multi-agent AI system.** Architected a two-agent solution for tender response. A financial agent performs structured extraction from vendor quotations in heterogeneous formats (PDF, Word, Excel, email, supplier portals) into a costing model; a technical agent applies RAG and Arabic OCR over RFPs of hundreds of pages, extracts scope of work and bill of quantities, and drafts the technical proposal from a fixed corporate template.
- **Human-in-the-loop and traceability.** Engineered source-level traceability from every extracted figure back to its document and page, mandatory human review gates, and visual flagging of AI-generated content with no vendor source behind it, as required by tender confidentiality.
- **Measured before go-live.** Scored 90%+ field-level precision and recall on the golden dataset of historical tenders before the platform went live, with regression testing on every prompt or model change.
- **LLM evaluation.** Defined the evaluation harness before the build: a golden dataset of historical tenders scored on field-level precision and recall, regression testing on every prompt or model change, and OpenTelemetry traces per agent step for latency, token cost and failure analysis.
- **On-premise AI infrastructure.** Designed a model-agnostic architecture against OpenAI-compatible APIs, benchmarked Qwen3, DeepSeek, Falcon-H1 Arabic and ALLaM, sized GPU and VRAM for vLLM and Ollama serving, and met Saudi PDPL and NCA requirements in a fully isolated network.

Technologies: Python, LangChain, LangGraph, vLLM, Ollama, Qwen3, DeepSeek, RAG, pgvector, Arabic OCR, OpenTelemetry, Docker

Iberia, S.A. – Forward Deployed Engineer (FDE)

Feb 2026 – Sep 2026 · Remote, Spain

Flight pricing platform. Technical owner of the fare pricing domain.

- **MCP server.** Built and deployed an internal Model Context Protocol server exposing pricing platform tooling to LLM clients, enabling governed function calling and agentic queries over fares, providers and configuration instead of ad-hoc scripts.
- **RAG service.** Delivered retrieval over internal specs and domain documentation with embeddings, hybrid search and cross-encoder reranking, cutting the time engineers spent locating pricing and NDC domain answers, instrumented with OpenTelemetry and Dynatrace like any other production service.
- **Managed and self-hosted model serving.** Deployed managed model serving on AWS Bedrock and Google Vertex AI to production alongside self-hosted inference, choosing per use case on latency, cost, data privacy, scalability and freedom of model choice, integrated with the MCP server and the RAG service under OpenTelemetry and Dynatrace observability.
- **AI-assisted software engineering.** Introduced Claude Code for service scaffolding, Karate and JUnit test generation and large refactors, driven by project-specific SKILL.md and CLAUDE.md conventions I defined and rolled out to the team.
- **Fare pricing services.** Engineered the services that calculate flight prices by orchestrating calls to external fare-calculation providers, normalising and aggregating heterogeneous responses under IATA NDC for internal channels and distribution partners; optimised latency and cut provider calls with Redis caching on the highest-traffic pricing flows.
- **Event-driven integration.** Delivered Kafka and PostgreSQL integration on AWS (ECS, SQS, S3) in hexagonal architecture with DDD, owning resilience of the pricing path (timeouts, retries and fallbacks) against third-party provider degradation.

Technologies: Python, Java, Kotlin, FastAPI, Spring Boot, MCP, AWS Bedrock, Google Vertex AI, Redis, Kafka, PostgreSQL, AWS, Karate, Claude Code

KPN – Senior Full Stack Engineer

Mar 2025 – Feb 2026 · Remote, Netherlands

Fibre installation tracking platform, ODF International Team.

- **Generative AI in the delivery workflow.** Introduced LLM-assisted development across the international team (test generation, legacy refactors, PR review) with shared prompt and review conventions, so AI output was always human-verified before merge.
- **Retrieval assistant.** Built semantic search with embeddings over runbooks and ServiceNow incident history, cutting the time to find the right precedent when triaging field-operations incidents.
- **Platform delivery.** Delivered a large-scale fibre installation tracking platform running in three European countries under international SLAs, building microservices and field-operations features and owning production incident resolution.
- **Data migration.** Migrated historical production records from PostgreSQL to DynamoDB with chunk-based storage, offloading the operational database and keeping the hot read path fast as volume grew.

Technologies: Python, Java, Kotlin, Spring Boot, Kafka, PostgreSQL, AWS (DynamoDB, Lambda), Embeddings, Semantic search, Vue.js, Docker

Mercadona, S.A. – Lead Backend Engineer

Mar 2024 – Mar 2025 · Remote, Spain

Product analytics platform.

- **Ingestion pipeline.** Engineered a Spring Batch ingestion pipeline processing millions of product records per day from heterogeneous sources, with restartable jobs and idempotent writes.
- **API-first services.** Designed API-first REST services (OpenAPI) in hexagonal architecture with DDD; sustained over 85% automated test coverage under real TDD.

Technologies: Java, Kotlin, Spring Boot, PostgreSQL, Kafka, AWS, GCP, Docker, Flyway

Inditex, S.A. – Lead Backend Engineer

Nov 2023 – Mar 2024 · Remote, Spain

Global logistics, garment-sorting event platform. Technical lead for the Azure to AWS migration.

- **High-concurrency platform.** Engineered a serverless platform on Azure and AWS processing thousands of garment-sorting events per minute across the logistics network.
- **Zero-downtime migration.** Led the Azure to AWS migration with zero downtime; built async REST APIs in hexagonal architecture with DDD and introduced Karate integration testing alongside the existing JUnit suite.

Technologies: Java, Kotlin, Spring Boot, MongoDB, PostgreSQL, Kafka, AWS, Azure, Docker, Liquibase

Openbank (Open Bank, S.A., Santander Group) – Senior Backend Engineer

Oct 2022 – Nov 2023 · Madrid, Spain

Investment automation platform. Technical lead for the Java to Kotlin migration.

- **Serverless batch at scale.** Implemented serverless batch pipelines on AWS Lambda processing millions of financial transactions per day on top of the T24 core banking system, inside a 100+ backend engineer organisation.
- **Java to Kotlin migration.** Led the migration across services and Lambdas; delivered internal tech talks on DDD and hexagonal architecture to around 50 engineers.

Technologies: Java, Kotlin, Spring Boot, PostgreSQL, AWS (Lambda, S3), Elasticsearch, Docker, MockK, T24

Idealista, S.A. – Backend Engineer

Mar 2021 – Oct 2022 · Madrid, Spain

Digital contract-signing platform.

- **Contract-signing platform.** Built a digital contract-signing platform for Spain, Portugal and Italy: REST APIs plus Kafka event streams for real-time signing, with event sourcing keeping an auditable, replayable history of every contract.

Technologies: Java, Kotlin, Spring Boot, PostgreSQL, Elasticsearch, Kafka, Docker, Event sourcing, DDD, TDD

Playboard, e-commerce search configuration platform.

- **Search configuration platform.** Built the configuration platform enterprise clients (Kroger, Carrefour, Inditex) use to tune their Elasticsearch-backed search engines; executed zero-downtime migrations: Java 8 to 11, GCP to AWS, monolith to API Gateway.

Technologies: Java, Spring Boot, MongoDB, Elasticsearch, Docker, Kubernetes, GCP, AWS, JUnit

SELECTED PROJECTS

Facturias – Multimodal invoice-processing SaaS

Website

Founder and sole engineer · Live · Web, iOS, Android

Structured extraction from invoices with a self-hosted vision-language model (Qwen2.5-VL) plus RAG over Spanish tax regulation to classify and validate entries, with confidence scoring, a labelled evaluation set tracking field-level precision and recall across model versions, and human review on low-confidence fields. VeriFactu compliant, multi-tenant FastAPI backend on PostgreSQL 17 with row-level security and pgvector.

- **Vision-language extraction.** Self-hosted Qwen2.5-VL reads invoices as documents rather than as OCR text dumps, preserving table structure, line items and stamps that flat OCR loses.
- **RAG over Spanish tax regulation.** Retrieval over the Spanish tax code classifies and validates each entry, so a deduction is justified by a retrievable rule rather than by model intuition.
- **Confidence and human review.** Every field carries a confidence score; anything under threshold is routed to human review instead of being written silently.
- **Measured, not assumed.** A labelled evaluation set tracks field-level precision and recall across model versions, so a model upgrade is a measurable decision.
- **Multi-tenant by construction.** PostgreSQL 17 row-level security isolates tenants at the database, not in application code. pgvector holds the regulation embeddings.

Technologies: Python, FastAPI, SQLAlchemy (async), PostgreSQL 17, Row-level security, pgvector, Qwen2.5-VL, Docker · **Keywords:** invoice OCR, VeriFactu, vision language model, RAG, multi-tenant SaaS, Spanish tax

ZORRO – Consumer dating app for the gay and queer community, iOS and Android

Website · App Store · Google Play

Founder and sole engineer · Live · iOS, Android

Shipped solo to both app stores in June 2026, reaching 700+ users with 50% D1 retention in three weeks. Semantic matchmaking with embeddings and pgvector plus cross-encoder reranking, and self-hosted vision-language moderation of user photos handling GDPR special-category data with no third-party processors. Polyglot backend in Kotlin/Spring Boot, Elixir and Python.

- **Semantic matchmaking.** Profiles are embedded and retrieved with pgvector, then reordered by a cross-encoder reranker: the same retrieve-then-rerank architecture as a production RAG pipeline, applied to people instead of documents.
- **Self-hosted photo moderation.** A vision-language model screens user photos on my own infrastructure. Sexual orientation data is GDPR special-category data, so it never reaches a third-party processor.
- **Polyglot backend.** Kotlin and Spring Boot for the domain, Elixir for realtime chat and presence, Python for the AI services.
- **Traction.** 700+ users with 50% D1 retention within three weeks of launch.

Technologies: Kotlin, Spring Boot, Elixir, Python, PostgreSQL, pgvector, Cross-encoder reranking, AWS SES, Hetzner, Docker · **Keywords:** semantic matching, embeddings, pgvector, content moderation, GDPR, mobile app

Apunta – Multi-tenant SaaS for shooting clubs, from custom PCB to mobile app

Website · App Store

Founder and sole engineer · Live · iOS, Android, Web, Hardware

Multi-tenant SaaS for shooting clubs owned end to end: custom NFC PCB (ESP32-C6, PN532, PoE) and C/ESP-IDF firmware, Spring Boot backend and React Native app. Range access, training sessions, scoring and club administration on one stack.

- **Custom hardware.** Designed the NFC access PCB around an ESP32-C6 and a PN532 reader with Power over Ethernet, and wrote the firmware in C on ESP-IDF.
- **Backend and multi-tenancy.** Spring Boot services with per-club tenancy, membership, range booking, training sessions and scoring.
- **Mobile.** React Native app shipped to the App Store and Google Play for shooters and club administrators.

Technologies: Kotlin, Spring Boot, React Native, PostgreSQL, C, ESP-IDF, ESP32-C6, PN532, Docker · **Keywords:** IoT, NFC, embedded firmware, multi-tenant SaaS, React Native

Grabia – Self-hosted meeting intelligence, fully local AI

Founder and sole engineer · Private inference node · Self-hosted

Local pipeline with WhisperX transcription and pyannote diarisation feeding a LangChain and LangGraph agent that produces structured summaries, decisions and action items, with RAG over the meeting archive. Served from a private node (Ryzen AI MAX+ 395, 128 GB, ROCm) running Qwen3-30B under Ollama and mistral.rs; nothing leaves the machine.

Technologies: Python, LangChain, LangGraph, WhisperX, pyannote, Ollama, mistral.rs, ROCm, Qwen3 · **Keywords:** local LLM, speech to text, diarisation, agentic summarisation

icekar – Distributed scraping and sub-second search over 100,000+ car listings

Website

Founder and sole engineer · Live · Web

Distributed scraping and sub-second Elasticsearch search over 100,000+ car listings indexed nightly, with an agentic LLM layer that rewrites its own scraper rules when sites change markup.

Technologies: Python, Elasticsearch, LLM agents, Docker · **Keywords:** web scraping, self-healing scrapers, Elasticsearch

Llave Asturias – Buy-renovate-sell house-flipping brand, plus its back-office and trade tools

Website

Lead engineer, client engagement via Esguil · Live · Web

Renovation and house-flipping brand in Asturias, Spain (legal entity LANDO FIRM, S.L.): the client finds a home to reform, Llave Asturias buys it, renovates it and sells it back move-in ready. Esguil built the marketing site and lead-capture funnel (Astro, zero JS by default, with a Go API for forms backed by SQLite, AWS SES for mail and Stripe for reservations), an internal back-office panel gated by a bcrypt-hashed key for managing subscribers, publishing opportunities and sending alerts, and Claude Code agentic workflows that plan, generate and validate the site's SEO and GEO content against a closed specification with build-time checks for stale dates, near-duplicate FAQ answers and unverified claims. A companion product, Presupuestador, is a separate offline-first iOS and web app for budgeting plumbing, electrical and HVAC renovation work.

- **Astro site and Go API.** Zero-JS-by-default Astro marketing site and lead funnel, backed by a Go API with SQLite storage, AWS SES for mail and Stripe for reservations, deployed behind nginx on a dedicated VPS.

- **Internal back-office.** A separate Astro admin panel, excluded from indexing and gated by a bcrypt-hashed key, for managing subscribers, publishing opportunities and sending alerts.
 - **Agentic SEO/GEO content pipeline.** Claude Code agents implement content and structured data against a closed specification; the build fails on stale dates, near-duplicate FAQ answers, unverifiable price claims or missing regulatory disclosures.
 - **Companion budgeting app.** Presupuestador, a separate offline-first Expo app and MCP server, prices plumbing, electrical and HVAC renovation work against Spanish market rates.
- Technologies:** Astro, Go, SQLite, TypeScript, AWS SES, Stripe, nginx, Docker, Claude Code · **Keywords:** house flipping, home renovation, buy renovate sell Asturias, SEO automation, GEO AEO, agentic content workflows

EDUCATION

M.Sc. in Artificial Intelligence Research (official) Sep 2026 – Sep 2027
UIMP / AEPIA
Specialisation in Machine Learning and Data Science. Started September 2026, part-time alongside consulting work.

B.Sc. in Computer Science and Software Engineering
University of Oviedo
Oviedo, Spain.

LANGUAGES

- **Spanish.** Native (CEFR C2)
- **English.** Professional working proficiency (CEFR B2)