
Carlos D. Escobar-Valbuena

Agent OS Architect & AI Engineering Lead

carlosdavidescobar@gmail.com +57 301 7758620 Bogotá, Colombia

broomva.tech linkedin.com/in/broomva github.com/Broomva

SUMMARY

Agent OS architect and AI engineering lead building autonomous, multi-tenant data platforms in production. AI Lead at Stimulus (Databricks-backed agentic procurement SaaS shipping to FIFA 2026 and the Olympics Committee) and, until Jun 2026, Co-founder & CTO at Wedi Pay (agentic cross-border payment orchestration on a Kafka + Databricks backbone with Postgres RLS isolation). Concurrent Data Architect contractor at TEAM International providing oversight on the big-data pipelines I previously led.

Builder of an open-source agent runtime substrate — Life Agent OS, Lago (event-sourced persistence on redb), Vigil (OpenTelemetry-native observability), and Haima (x402 metered agentic payments). Author of the Recursive Controlled Systems (RCS) paper series formalizing LLM-as-controller agents as a 7-tuple Σ with recursive stability budgets across hierarchical control levels.

Strong on production RAG (pgvector, BM25 hybrid retrieval, LLM-as-judge), data governance (Unity Catalog RBAC, Postgres RLS, KMS-per-tenant, lineage via OpenTelemetry), CDC and event sourcing (Dagster, Kafka/Redpanda, append-only journals), and multi-tenant isolation. MSc in Artificial Intelligence at Universidad de los Andes (2026).

SKILLS

Programming Languages: Python, Rust, TypeScript / JavaScript, SQL, Matlab, C++

Databases & Storage: PostgreSQL (incl. pgvector), Delta Lake, Snowflake, Databricks Unity Catalog, AWS RDS, MySQL, MongoDB, Redis, NeonDB

Data Platforms & Lakehouse: Databricks (deep) — Unity Catalog, Zerobus streaming, medallion architecture (bronze / silver / gold), Delta Lake; Snowflake, Spark / PySpark, Iceberg, Trino, Azure Data Factory, AWS Athena / Glue / Lake Formation (architectural), EMR Serverless (architectural), Kafka / Redpanda, Dagster, MLflow

AI Native Tooling: Anthropic SDK (Claude Opus / Sonnet / Haiku), OpenAI, AWS Bedrock, LangChain, LangGraph, LangSmith, Langfuse, MCP, Mastra, N8N, Vercel AI SDK, Generative UI, RAG (BM25 + pgvector / FAISS), evaluation harnesses

ML Frameworks & Tools: PyTorch, PyTorch Lightning, TensorFlow, Hugging Face Transformers (fine-tuning), MLflow, SageMaker, Azure ML, Prophet

Cloud & Infra: AWS (S3, EC2, RDS, Lambda, IAM/KMS, IoT Core, Greengrass), Azure (AKS, ML, Digital Twins, Data Factory, DevOps), GCP, Kubernetes, Terraform, GitHub Actions

Observability & Governance: OpenTelemetry (authored `vigil` Rust crate), CloudWatch, Datadog, data contracts, lineage tracking, Unity Catalog RBAC, Postgres RLS, multi-tenant isolation, KMS-per-tenant

Interfacing: Next.js, React, React Native, Vercel AI SDK, Generative UI, PowerBI

Blockchain / Web3: EVM development, x402, USDC settlement on Base

Languages: English (C1, full professional), Spanish (native), German (basic)

EXPERIENCE

AI Lead *Stimulus — Supplier-Intelligence Procurement SaaS* Nov 2024 – Present
Leading the architecture redesign of an AI-native, Databricks-backed agentic procurement platform serving customers including FIFA 2026 and the Olympics Committee. Optimizing agent UX, ensuring scalable deployment, and guaranteeing enterprise-grade reliability.

- Architected the real-time data substrate around **Databricks Zerobus** for streaming ingestion into a **Unity Catalog**-governed **medallion architecture** (bronze / silver / gold) — eliminating external message-bus dependencies while preserving lineage, auditability, and per-tenant column-level access controls.
- Redesigned the product foundation to leverage best-in-class data governance and managed access to artifacts, with state sharing and agentic flows across the system. Delivered customer-facing features at production scale.

- Owned LLM deployment and tuning across Claude and OpenAI variants. Implemented auto-failover, AKS/EKS comparisons, cost-latency trade-off decisions, and hardened rate limiting and JWT authentication for multi-tenant clients.
- Delivered agent-driven procurement flows that combine multi-document retrieval with predefined policies, task memory, and deterministic tool adapters — yielding accurate and actionable supplier insights, risk checks, and sourcing recommendations.
- Coordinated AI brownbag sessions and internal enablement initiatives. Mentored engineers on prompt patterns, guardrails, and evaluation harnesses.

Co-founder & CTO *Wedi Pay — Agentic Global Payment Orchestration*

Oct 2024 – Jun 2026

Led platform architecture and AI/ML engineering for an AI-native, multi-tenant B2B cross-border payment orchestration platform unifying traditional banking infrastructure with blockchain rails.

- Designed and shipped a modular Turborepo monorepo: Next.js frontend + API routes, FastAPI Python backend, LangGraph-based agentic services, NeonDB / Postgres data layer with Row-Level Security for stringent multi-tenant isolation, and a shared agent/UI state backbone.
- Architected an event-driven backbone on Kafka / Redpanda integrated with Databricks asset bundles, enabling gold-table ETLs that feed both application models and analytics — providing real-time insights into payment flows.
- Defined Wedi Agents, an in-product assistant that interprets historical Wedi activity and, under strict policy, consent, and scoped credentials, executes payment actions (e.g., buy or send USDC) with full audit trails.
- Helped design and validate the comprehensive revenue model (subscription tiers, volume-based fees, dynamic session quotes), benchmarked against leading LATAM players (Global66, DolarApp, Wise) and mapped to CFO/finops reporting needs.
- Signed and shipped payment-rail integrations on a provider-agnostic integration layer I designed: **Prometeo** (open-banking API for direct B2B account-to-account payments in LATAM), **TruBit Business — now VelaFi** (USDC on/off-ramp operations across multiple countries), and **Thirdweb** (web3-native experiences).

TEAM International — multiple concurrent contracts

Nov 2020 – Present

Microsoft Gold Partner; named among IAOP Global Outsourcing 100 and Software 500 companies.

Data Architect (Contract)

Nov 2024 – Present

- Architectural oversight on the big-data pipelines I originally designed; supporting and enhancing them as workloads scale.
- Architecting next-generation iterations to meet new business and technical requirements (higher throughput, new data sources, advanced analytics).
- Overseeing migrations and platform updates with minimal downtime and full data integrity.

Technical Interviewer (Contract)

Nov 2024 – Present

- Conducting technical interviews for senior data engineering, ML, AI, and modern cloud-platform roles.
- Designing assessment frameworks rigorously evaluating distributed data processing, predictive modeling, generative AI, and cloud-native architecture.
- Collaborating with hiring managers to align interviews with project needs and team composition.

Sr. Machine Learning Engineering Tech Lead — Digital O&G Solutions

Sep 2022 – Nov 2024

- Owned the Time Series Engine product end-to-end — a multi-tenant streaming + batch telemetry engine on Databricks + MLflow + PyTorch + Prophet under TDSP-aligned MLOps. Anchor customer: a **NYSE-listed US natural-gas operator with ~70,000 wells** across multiple US basins (Appalachian + Central regions). Production lakehouse pattern with **medallion-style bronze / silver / gold tables since 2021**.
- Led the Time Series Engine team — backlog definition, estimation, risk / feasibility assessment, architectural specification across multi-quarter roadmaps.
- Led development of a specialized ML package for forecasting and anomaly detection in O&G time series (alarm threshold estimation, predictive maintenance) with PySpark + MLflow distributed learning. Productized as the *SCADALite Well Manager* offering with bilingual (EN/ES) capability decks and tiered pricing.
- Designed a digital-twin framework via Azure Digital Twins for O&G processes, enhancing asset management and operational insights across thousands of wellsites.
- Shipped and operated a multi-tenant telemetry solution handling **10,000+ records/sec**, with concurrency, workload distribution, and customizable ETL logic across stream / batch pipelines.
- Conducted technical interviews for junior and senior data engineering, big data, ML, and Python roles.

Machine Learning Engineering Tech Lead

Aug 2021 – Sep 2022

- Led big data streaming solutions: distributed processing/storage, batch and real-time inference for forecasting and anomaly detection (PySpark, Prophet) for O&G production assessment, on-site safety, and key-metric evaluation across temperature, pressure, and volume of pipelines, tanks, and extraction systems.
- Completed development of a PySpark streaming engine for real-time data with business-specific logic interfaces and dashboards for O&G production volumes at thousands of sites.
- Owned the Time Series Engine product backlog and architectural roadmap; integrated packaged solutions for data processing and inference into portfolio products under structured Agile.
- Designed and implemented an orchestration engine for dynamic task creation and execution, enhancing data product development and monitoring (Azure DevOps for packaging, versioning, delivery).
- Led knowledge-sharing sessions, webinars, and tech talks; participated in C-level architectural pathways, product prioritization, and technical specification meetings.

Data Engineer

Nov 2020 – Aug 2021

- Migrated an O&G company's legacy ERP into SAP using best-practice design and complex SQL development (MSSQL + T-SQL with SAP).
- Designed, tested, and implemented a Snowflake data warehouse for streaming IoT device data — a unified source of truth with high responsiveness for O&G analytics.
- Defined the migration guidelines from legacy DWH to Snowflake using Talend and Azure Data Factory ETL.
- Led an R&D team applying AI techniques to raw telemetry data, building Anomaly Detection, Alarm Threshold Estimation, and Time Series Inference engines (decoupled, scalable, batch or streaming) using Azure ML + scikit-learn + TensorFlow + PyTorch + PowerBI.
- Engineered a Snowflake pipeline transforming semi-structured IoT data into structured DWH models.

AI Engineer *Independent Engagements & Selected Project Deliverables*

Aug 2023 – Present

Cross-industry agentic, data-engineering, and ML pipeline work delivered for clients, partners, and high-bar evaluation cycles. Anonymized references below; case studies available on request.

- **AWS Premier Consulting Partner (US-based AI advisory):** Designed and prototyped a production-grade agentic RAG over an internal SageMaker engineering knowledge base — pgvector + FAISS hybrid retrieval, three switchable interaction modes (graph-based agentic RAG, tools agent, conversational RAG), LangServe backend with Chainlit UI. Delivery accepted at evaluation stage.
- **Tier-1 LATAM e-commerce + fintech platform:** Built a Polars-based ML ranking pipeline for the recommendation carousel — prints / taps / payments event-stream feature engineering, GridSearchCV-tuned classifier and regressor variants, encapsulated as a class-based pipeline with MLflow autolog for full experiment reproducibility. Architected for batch and incremental retraining at scale.
- **Public-sector procurement intelligence platform:** Both a TF-IDF baseline and a LangGraph agentic RAG (FAISS + retriever tools) to match upcoming government project tenders against qualified subcontractor capabilities and certifications. Bilingual processing pipeline.
- **Enterprise insurance analytics (LATAM):** Read-only LangGraph SQL agent over insurance contract data with schema-aware multi-step query refinement, multi-LLM backend (OpenAI + Groq Llama-3), and Chainlit streaming UI.
- **LLM-enhanced Change Data Capture pipeline** (Dagster + Supabase + Airtable + LangChain + TogetherAI + Twilio): auto-syncs records, triggers scraping and NER + summarization, and exposes a WhatsApp endpoint for direct LLM-agent interaction.
- **Spark SQL LLM Agent** on Databricks Delta (LangChain + Streamlit + OpenAI), implemented under **Unity Catalog RBAC** for role-scoped, secure agent-driven analytics.

Digitalization Leader *INSA Ingeniería S.A.S*

May 2019 – Nov 2021

Siemens Solution Partner. Digitalization, automation, and instrumentation technological developments and integrations.

- Developed solutions for O&G, energy, and discrete manufacturing industries integrating industrial automation equipment from multiple vendors (Siemens, Rockwell). Built digital solutions for IIoT, cybersecurity, and analytics, end-to-end from field to cloud.
- Accelerated deployment of a comprehensive digitalization strategy on **AWS IoT Core, AWS Greengrass, AWS EC2, ThingsBoard**, Siemens IoT suite (Simatec IoT, S7), with cybersecurity baseline following ISO 27001 — achieving robust IT/OT integration.
- Completed digitalization of an analog O&G field with 50+ wells across 3 sites: improved remote monitoring, analytics, and decision-making; **reduced transportation costs for well monitoring by 50%** and **avoided 2 overfills/spills** that would have

caused significant environmental damage and cleanup costs.

- Designed industrial network topologies, PLC programming, and closed-loop control system tuning (Siemens S7-1200/1500, TIA portal); pump control for separators, pipelines, and tanks.
- Enabled remote monitoring and predictive maintenance for gas-electricity generators (Remote Generation Units Condition Monitoring).
- Participated in holistic project design, achieving **cost reductions up to 35%**; fieldwork and resolution in remote locations with scarce resources.
- Implemented a touchless temperature measuring system (Siemens HMIs + industrial Raspberry Pi + GCP Voice API + Cloud Storage) for COVID-19 health and safety protocols.

Mechatronics Engineer Intern *Universidad de San Buenaventura*

Jun 2017 – Nov 2017

Industrial automation laboratory; Allen-Bradley and Siemens PLC technologies.

- Reviewed, analyzed, and designed a preventive maintenance route for the industrial automation lab.
- Plotted methodology for testing, inspection, and correction of PLCs, instruments, and actuators.
- Designed electronic cards for recovering broken-down devices (Eagle, Proteus for PCB design and simulation).
- Organized and structured predictive maintenance and repair of systems and electronic hardware.

OPEN SOURCE & RESEARCH

Personal portfolio canon and continuously evolving substrate at broomva.tech and github.com/Broomva.

Life Agent OS ([broomva/life](#))

Rust agent runtime kernel. Multi-tenant gateway with TLS 1.3, JWKS, KMS abstraction (Vault Transit / AWS KMS / GCP KMS), per-tenant rate limiting (token-bucket per-user / per-IP), event-sourced persistence layer (Lago, on redb v2), and OpenTelemetry-native observability (Vigil). 15+ crates published on crates.io.

Haima — Agentic Finance Engine

Six-crate Rust workspace for x402 metered machine-to-machine payments with USDC settlement on Base. secp256k1 wallet, per-task billing, full audit trails. Distribution rail for the Life Runtime metered platform.

bstack — Agentic Development Harness ([broomva/bstack](#)) + Skills ([broomva/skills](#))

Open-source, control-native harness metalayer for AI-native development: twenty irreducible automation primitives (P1–P20) — conversation bridging into a knowledge graph, policy gates on destructive operations, parallel worktree fan-out, knowledge bookkeeping with promotion gates, CI productive-wait, persistent cross-context loops, dependency-chain and state-snapshot disciplines, and a cross-model adversarial review gate — that turn an agent-driven workspace into a self-operating system. Companion skills monorepo: **48 portable agent skills** (agentskills.io spec) compatible with Claude Code, Codex, Cursor, Gemini CLI, Goose, and Copilot; distributed via `npm skills add`.

Agentic Control Architecture — Autonomy / Search / Intelligence Research

Control-systems-native architecture for autonomous workflows: governance expressed as typed plant/action/trace schemas, machine-readable policy gates, and hook-enforced safety shields, with verification signals causally independent of the acting model ($h \perp U$) — closing the loop that prompt-level instruction leaves open. Active research program on autonomy, search, and intelligence over the RCS substrate: hierarchical stability budgets (L0–L3), rule-of-three crystallization of recurring patterns into machine-checkable primitives, and empirical falsification harnesses that publish negative results at equal prominence to positives.

Recursive Controlled Systems (RCS) — 5-Paper Series

Formalization of LLM-as-controller agents as a 7-tuple $\Sigma = (X, Y, U, f, h, S, \Pi)$ with recursive stability budgets across hierarchical control levels (L0–L3). Foundations paper (P0) complete; 4 papers in progress. Provides the mathematical substrate for the agent OS architecture.

Arcan AI ([arcan.la](#))

AI / web3 tooling platform for decentralized customization and enhancement of AI agents — bridging AI and blockchain for personalized, user-owned agents.

Vortex AI

Versatile Orchestrated Execution Engine for Data & AI Pipelines. LLM-powered smart agent for workflow development and automation on top of Spark/Databricks.

Spanish/Wayuu Translation Model ([HuggingFace](#))

Fine-tuned BART, T5, MT5, and Llama-2-7B for low-resource indigenous-language translation. Public demo space.

Broomva Book RAG Agent ([book.broomva.tech](#))

Chainlit endpoint (LangChain + FAISS) over a vectorized corpus of personal notes and course content; serves as an interaction layer for the underlying knowledge graph.

Hyperopt Prophet, MLFlowOps, Databricks Session

Open-source Python packages for distributed hyperparameter optimization on Prophet (Pandas/PySpark), streamlined MLflow operations, and Databricks session creation.

EDUCATION

MSc, Artificial Intelligence *Universidad de los Andes* 2026

BSc, Mechatronics Engineering *Universidad de San Buenaventura* 2013 – 2018
GPA 4.1/5.0

- Thesis: Designed and validated a control system for rejection of failures and disturbances in a paint-drying oven.
- Speaker, IV International Congress on Advanced Mechatronics Technologies, Design, and Manufacturing: *On Cascade ADRC for Time-Delay Systems with Subsystems Faults* (2018).
- Best GPA Academic Distinction (2016).

CERTIFICATIONS & TRAINING

Diplomado Machine Learning & Data Science Avanzado (Universidad Nacional). Building Transformer-Based NLP Applications (Nvidia); Development of Predictive Algorithms with Machine Learning (Fedesoft Colombia); Quantum Computing (IBM); Big Data: Acquisition and Storage (Universitat Autònoma de Barcelona); GCP Big Data and Machine Learning (Google); Machine Learning with Big Data (UC San Diego); Data Science with Python (IBM); Getting Started with AWS Machine Learning (AWS); TensorFlow 2.0 and Deep Learning (SuperDataScience); Kinematics: Describing the Motions of Spacecraft (University of Colorado Boulder).