Luis Sánchez

Software Engineer · Head of Engineering / Tech Lead · Technical Advisor · Open Source Contributor

Santiago, Chile linkedin.com/in/luissanchezm86 - github.com/LuisSanchez - luissanchez.github.io

PROFESSIONAL SUMMARY

Software engineer with 10+ years building high-performance systems in fintech, fraud prevention, and SaaS. Proficient in Python, Django, React, and cloud infrastructure (AWS, Azure, GCP). Hands-on tech lead who builds platforms from scratch, leads teams, and contributes to open source. Translates complex business requirements into clean, scalable full-stack solutions using modern tooling and AI-assisted practices.

TECHNICAL SKILLS

Languages & Frameworks: Python (7-8 yrs), JavaScript, TypeScript, React, Django, REST APIs, Node.js, Java

Cloud & Infrastructure: GCP, AWS, Azure, Docker (4-5 yrs), Kubernetes, GitHub Actions, CI/CD, Terraform

Data & Databases: PostgreSQL, MySQL, Cassandra, Hive, Hadoop, Redis, DynamoDB, SAP HANA, Azure Synapse, ETL, Data Pipelines, Apache Superset, Polars, Pandas

Streaming & Async Processing: Kafka, RabbitMQ, Celery, Real-Time Streaming, Batch Pipelines

Machine Learning & Analytics: ML Model Optimization, Parallelization, Matrix Computation, Feature Engineering, SciPy, Scikit-learn, Power BI

Payments & ERP Integrations: Stripe, Adyen, PayPal, CyberSource, Apple Pay, PayU, Pix, Transbank, SAP HANA, Azure Synapse, Fraud Prevention

AI-Assisted Development: Cursor, Claude Code, GitHub Copilot, OpenCode + Ollama (local inference), Perplexity

Monitoring & Observability: Datadog, Sentry, New Relic, KPI Design, Performance Monitoring

Practices & Methodologies: Agile, Scrum, TDD, Microservices, SOA, System Architecture, Code Review

PROFESSIONAL EXPERIENCE

Tech Lead / Founding Enginner

May 2026 — Present

UpLabs · United States / Chile - Remote

UP.Labs is a venture studio that partners with global corporations to launch and scale technology startups in the transportation, mobility, and emerging verticals. As a founding engineer and hands-on tech lead within a venture pod, I lead the full engineering lifecycle from zero — architecture decisions, team building, and active development simultaneously.

- Lead, mentor, and manage a fullstack engineering team of 5–10 engineers, setting technical direction, coding standards, and code review culture from day one.
- Design and build scalable fullstack applications end-to-end, owning both backend services and APIs as well as React-based frontends.
- Establish and evolve system architecture for greenfield products, making foundational decisions on infrastructure, data pipelines, and API design under startup constraints.
- Drive AI-first development principles, leveraging frontier tools to accelerate product velocity and quality across the engineering team.
- Collaborate directly with C-level, product managers, and designers to define and execute product roadmaps.
- Implement security best practices, database performance optimization, and high-availability patterns across production systems.

Software Engineering Specialist – Human Data Team

Apr 2026 — Present

xAI · Remote

Elon Musk's AI company building Grok and frontier language models

- Building scalable data pipelines for human feedback collection and curation used to train and align frontier AI models including Grok.
- Designing and maintaining annotation tooling and quality-control workflows that ensure high-signal training data at scale for RLHF pipelines.
- Collaborating cross-functionally with research and model teams to define data requirements and evaluation benchmarks for model improvement cycles.

Head of Engineering / Tech Lead

Jan 2024 — Apr 2026

Revenew · Santiago, Chile

Founded and led engineering at a Chilean B2B SaaS startup managing USD 1.5B in annual food transactions. Built the entire platform from scratch — Django/React monolith, ETL pipelines, and MLOps — while leading a cross-discipline team and acting as the client-facing technical contact.

- Designed and built a multi-tenant Django monolith with tenant isolation, SSO, RBAC, and a client-facing BI/dashboard, owned by a cross-discipline team of engineers and data scientists.
- Designed ETL pipelines to ingest data from client ERP and data-lake sources including SAP HANA, Azure Synapse, and SFTP, enabling reliable scheduled batch processing for downstream ML models.
- Integrated ML-driven elasticity and forecasting models using SciPy, Pandas, and solver libraries; optimized runtime via parallelization and matrix computation efficiency.
- Evaluated and integrated pre-trained inference approaches to speed up predictions; reduced model runtime significantly by changing tooling and workflows.
- Containerized all services with Docker, established CI/CD pipelines to Azure via GitHub Actions, and instrumented monitoring and error tracking with Sentry.
- Used RabbitMQ and Celery for asynchronous task processing, improving system throughput and decoupling long-running background jobs.
- Ran sprint planning and code reviews, set coding standards, mentored engineers and data scientists, and owned deployments and operational readiness.
- Consulted directly with enterprise clients — pricing managers and technical data teams — to design integration architectures and onboard their data infrastructure.
- Embedded AI-assisted development practices (Cursor, Claude Code, Ollama-based local inference) into daily team workflows, increasing engineering velocity and code quality.

Open-Source Software Engineer (Contracted)

Nov 2025 — Feb 2026

Apache Superset · Preset · Remote

Short-term hourly contract with Preset to stabilize Apache Superset ahead of the v6.0 release. Contributed 10 merged PRs across the React/TypeScript frontend and Python backend, fixing chart regressions, improving UX, and reducing UI latency by ~550 ms. Worked directly with core maintainers in a 60,000-star open-source codebase.

- Hourly-paid contract to help Apache Superset (60,000+ GitHub stars) migrate and stabilize between major versions, working with the Preset team to prepare the codebase for the v6.0 release.
- Contributed 10 merged pull requests spanning frontend React/TypeScript fixes and backend Python improvements across charts, APIs, and developer tooling.
- Fixed critical frontend regressions: missing filter badge icons on World Map charts, hidden x-axis data points in timeseries charts, value label clipping on horizontal bar charts, and legend alignment regression in v6.0.
- Improved chart UX: removed invalid “Stream” style from Bar charts with backward compatibility, automated axis padding on title add/remove, and fixed Table chart cell bars for NULL-containing columns.

- Added Jinja template support to the backend format_sql API endpoint, enabling parameterized query formatting in SQL Lab.
- Refactored TimezoneSelector with lazy loading and global caching, eliminating a ~550 ms UI delay when opening the scheduling panel.
- Resolved unintended password-manager autofill on search inputs across multiple admin list pages.

Software Engineer

Aug 2022 — Dec 2023

Uber · Santiago, Chile

Built the data engineering layer powering Uber's global fraud prevention ML models. Developed real-time Kafka consumers, day-partitioned Cassandra tables, and Hadoop batch pipelines processing millions of log events per minute, keeping fraud signals low-latency without impacting the user experience.

- Built feature engineering and ingestion pipelines for global fraud prevention, enabling real-time and batch machine learning models operating at high QPS volumes.
- Consumed Kafka topics and built real-time streaming consumers that populated day-partitioned Cassandra tables, providing low-latency feature reads for fraud-detection models that processed millions of logs per minute.
- Implemented batch ingestion pipelines to feed Hadoop-based stores, enabling data scientists to train and deploy global fraud models on enriched historical data.
- Designed table partitioning and storage patterns for low-latency reads; validated that new pipeline changes did not negatively impact production operations.
- Engineered an ETL pipeline to process American Express refunds and disputes: extracted raw financial data, normalized it, and fed it into Uber's internal reconciliation systems.
- Worked with Java and Go microservice components that integrate with Kafka as part of the broader fraud infrastructure.

Backend Engineer

Oct 2020 — May 2023

Cornershop by Uber · Santiago, Chile

Backend engineer on the Payments and Fraud Prevention team at a LATAM grocery delivery platform acquired by Uber. Integrated eight payment providers across four countries, built fraud-prevention mechanisms, and instrumented the payment layer with observability tooling and business KPIs.

- Integrated multiple payment gateways (Stripe, Adyen, PayPal, CyberSource, Apple Pay, PayU, Pix, Transbank) to support millions of transactions across Latin America, handling API differences, reconciliation logic, and regional requirements.
- Designed fraud-prevention mechanisms including sliding-window rate controls with exponential backoff for new card additions and selfie-based identity verification to reduce fraud while minimizing friction for legitimate users.
- Implemented observability tooling and business KPIs using Datadog and Sentry to monitor system health, payment success rates, and operational performance.
- Leveraged Python and Django for all backend services; ranked as a top-10 contributor in the company's private GitHub repository.

Senior Software Engineer

Apr 2012 — Oct 2020

In Motion · Santiago, Chile

Software consultancy delivering enterprise systems for insurance and financial clients across Latin America, including MetLife and Citibank. Worked extensively on complex legacy on-premises codebases, delivering incremental improvements and new integrations while introducing modern practices like TDD, Docker, and CI/CD over nearly nine years.

- Developed software solutions for major insurance and financial clients across Latin America, including MetLife and Citibank Mexico, using .NET, .NET Core, MVC, and Entity Framework.
- Built fast online purchase flows for insurance products (e.g., cell phone protection plans) and a Citibank Argentina integration that streamlined new account onboarding by pulling customer data from their core systems.
- Maintained and upgraded complex legacy codebases for large enterprise clients who preferred on-premises hosting, developing deep experience working within existing systems and delivering incremental improvements.

- Designed and implemented a Security Token Service (STS) to centralize authentication across distributed microservices.
- Containerized services using Docker and Azure DevOps; championed TDD practices using xUnit and FluentAssertions.

EDUCATION

Bachelor's Degree in Informatics Engineering *2003 — 2009*

Universidad Católica Andrés Bello · Caracas, Venezuela

CERTIFICATIONS

Prompt Engineering for Developers *May 2023*

DeepLearning.ai

ScrumManager — Expert Level *June 2015*

ScrumManager.org

LANGUAGES

Spanish — Native

English — Professional Working Proficiency

PERSONAL PROJECTS

Kegy.app — Brewery tracking web application

Built with Python, Django, Kubernetes, and PostgreSQL on Google Cloud Platform. Continuous delivery via CircleCI. Covers inventory management, fermentation tracking, and batch reporting.