Luis Sanchez

Software Engineer · Head of Engineering · Open Source Contributor

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

PROFESSIONAL SUMMARY

Software engineer with 10+ years of experience building high-performance, scalable systems across fintech, fraud prevention, payments, data visualization, and SaaS. Hands-on CTO with a track record of leading architecture redesigns, establishing DevOps practices, and delivering open-source contributions to large-scale projects. Adept at translating complex business requirements into clean, maintainable backend and full-stack solutions using Python, Django, React, and cloud-native tooling.

TECHNICAL SKILLS

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

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

Data & Databases: PostgreSQL, MySQL, Hive, Cassandra, SAP HANA, Azure Synapse, ETL, Data Engineering, Data Pipelines, Apache Superset, Polars

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

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, KPI Design, Performance Monitoring

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

PROFESSIONAL EXPERIENCE

Head of Engineering / CTO   *Jan 2024 — Present*

Revenew · Santiago, Chile

- Lead the technical strategy and architecture of the platform, overseeing a full migration from a monolithic system to a service-oriented architecture built on Python, Django, and React.
- Established CI/CD pipelines on Azure and AWS using GitHub Actions, and containerized all services with Docker, reducing deployment time and improving environment consistency.
- Implemented Power BI dashboards for real-time KPI monitoring, giving business stakeholders clear visibility into key revenue and operational metrics.
- Redesigned authentication and authorization flows, significantly improving application security posture.
- Designed and implemented ETL workflows and scheduled batch jobs to connect to clients' enterprise data warehouses (SAP HANA, Azure Synapse) and ERP systems, enabling reliable data pipelines for downstream analytics.
- Consulted directly with enterprise client stakeholders — from business managers to technical data teams — to design integration architectures and onboard their data infrastructure.
- Optimized existing machine learning models by introducing parallelization and more efficient matrix computation using Polars, significantly reducing processing time for complex analytical workloads.

- Improved asynchronous process handling to increase system throughput and enhance the end-user experience.
- Embedded AI-assisted development practices into the team's daily workflow using Cursor, Claude Code, and Ollama-based local inference, increasing engineering velocity and code quality.

Open-Source Software Engineer   *Dec 2025 — Feb 2026*

Apache Superset · Apache Software Foundation · Remote

- Contributed 10 merged pull requests to Apache Superset, a leading open-source data visualization and BI platform with 60,000+ GitHub stars, working directly with the core maintainer team.
- Fixed critical frontend bugs: missing filter badge icons on World Map charts, hidden x-axis data points in timeseries charts, and value label clipping on horizontal bar charts.
- Improved chart UX: removed invalid "Stream" style from Bar charts with backward compatibility, corrected legend alignment regression in v6.0, and automated axis padding based on title presence.
- 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 delay when opening the scheduling panel.
- Fixed Table chart cell bars not rendering for columns containing NULL values.
- Resolved unintended password-manager autofill on search inputs across multiple admin list pages.

Software Engineer   *Aug 2022 — Dec 2023*

Uber · Santiago, Chile

- Developed high-performance feature engineering rules for global fraud prevention, operating at some of the highest query-per-second (QPS) volumes in the industry.
- Built and maintained real-time data pipelines for user risk evaluation, integrating with Uber's large-scale data infrastructure.
- 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.
- Proactively identified emerging fraud patterns by analyzing large-scale logs and translating business requirements into scalable, production-grade anti-fraud solutions.

Backend Engineer   *Oct 2020 — May 2023*

Cornershop by Uber · Santiago, Chile

- Developed payment integrations with Stripe, Adyen, PayPal, CyberSource, Apple Pay, PayU (Colombia), Pix (Brazil), and Transbank (Chile), supporting millions of transactions across Latin America.
- Built fraud prevention tools that reduced financial losses while maintaining low friction for legitimate users.
- Leveraged Python and Django for backend services; implemented observability tooling with Datadog and Sentry; designed KPIs to track system performance and business health.
- Ranked as a top-10 contributor in the company's private GitHub repository.

Senior Software Engineer   *Apr 2012 — Oct 2020*

In Motion · Santiago, Chile

- Delivered end-to-end enterprise software for major insurance clients (MetLife, Citi) using the Microsoft stack: .NET, .NET Core, MVC, and Entity Framework.
- Modernized legacy web interfaces with CSS3, HTML5, Bootstrap, and JavaScript, improving usability and consistency.
- Containerized services using Docker and Azure DevOps pipelines.
- Championed TDD practices using xUnit and FluentAssertions, increasing code coverage and reducing regression risk.

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.