CARLOS ADRIAN GARCIA REYES

Lead Software Architect | AI Platform Engineer

Mexico City, MX · Remote (U.S. time zones) | carloss_dev@hotmail.com | (+52) 55 40 67 59 15
linkedin.com/in/carlos-adrian-garcía-reyes-b5b288171 | carlossdev.com

PROFESSIONAL SUMMARY

Lead Software Architect and AI Platform Engineer with 10+ years taking enterprise systems from zero to production and owning them once live, including at two $1B+ venture-backed startups (Kavak, Clip). Builds distributed platforms for high demand and high concurrency, with availability and reliability as product requirements. Currently leads the platform behind a K-12 literacy ecosystem with 200K+ daily active users across 300+ U.S. districts, heading a ~20-engineer platform organization. Pioneered company-wide AI adoption (governance, PII guardrails, MCP tooling, AI agents), cutting AI costs by 60% and refactoring 23 core microservices in 45 days. Business-first problem solver who removed delivery bottlenecks to take production releases from every 5 weeks to every 3 days.

PROFESSIONAL EXPERIENCE

GlobalLogic / 95 Percent Group

Dec 2022 – Present

Lead Software Architect | AI Platform Engineer · Remote (USA / Mexico)

System architect and technical lead for the One95 Literacy Ecosystem, a K-12 science-of-reading platform used across 300+ U.S. districts with 200K+ daily active users, owning platform engineering, AI strategy, and delivery end to end.

- Cut production release cycles from **every 5 weeks to every 3 days** (2 weeks at the midpoint) by building platform tooling and delivery pipelines that remove engineering bottlenecks and bureaucratic hand-offs; track **DORA metrics** in an executive dashboard reported to the CTO and VP.
- **Reduce AI costs by 60%** through model right-sizing, temperature tuning, and prompt optimization, backed by token-cost agents, per-role usage limits, and consumption monitoring.
- Turn new reports from **a 2–3 week sprint into a 3-day catalog change** in production with a reporting agent system, plus product-facing AI agents that raise the reporting team's delivery velocity.
- **Architect the core system**, including the authentication and authorization layer, so new engineers ship business features from day one without touching identity, permissions, or other cross-cutting concerns.
- Pioneer and scale **enterprise AI adoption**, owning the AI governance model (role-based access, usage policies, acceptable-use standards) and equipping every developer with AI tooling: custom MCP tools, modular LLM skills, and deployment pipelines for AI agents.
- **Secure AI usage** with PII-handling guardrails inside the AI gateway, plus access controls and rate-limiting that protect sensitive data and internal systems.
- Orchestrated the AI-accelerated refactoring of **23 core microservices in just 45 days**, after leading client discovery workshops to scope and align the initiative.
- Lead the cross-team **platform organization of ~20 engineers** across multiple pods in a distributed U.S./Mexico team, with full ownership and autonomy: recurring 1:1s, roadmap and code-review standards, and a high-performance team culture.

KAVAK

Jan 2022 – Dec 2022

Lead Engineer (promoted from Senior Software Engineer) · Mexico

Mexico's first unicorn (US$8.7B valuation in 2021), serving millions of users across Mexico, Argentina, and Brazil; joined during post-funding hypergrowth.

- Shipped the **panic button and theft-reporting feature**, dispatching real-time alerts in Mexico City that helped recover stolen vehicles and **raised customer satisfaction by 10%**; still live today, its impact earned a promotion from Senior Software Engineer to Lead Engineer.
- **Led a 12-person squad**, owning technical direction, code reviews, and delivery planning; took the team on-site to observe operations firsthand and turned that business insight into product improvements.
- Drove the **monolith-to-microservices transition** with decoupled, cloud-native Node.js and AWS Lambda services, keeping the platform scalable while engineering headcount grew rapidly.
- **Bridged product and engineering**, turning business logic into technical user stories and shipping through GitLab CI/CD across DEV, QA, and Production.

Clip
Aug 2021 – Jan 2022

Infrastructure & Platform Engineer · Mexico

Mexican fintech unicorn (~US$2B valuation, 1,000+ employees) processing payments for hundreds of thousands of merchants.

- Built the internal **RabbitMQ library** that Clip developers use for messaging, part of a set of core Node.js libraries, shared services, and SDKs that standardized integrations across product squads.
- **Treated developers as internal customers**, shifting the team to a platform mindset and shaping tooling around their workflows to remove friction and speed up delivery.
- **Owned and operated AWS infrastructure** (EC2, ECS, Kubernetes, CloudFormation) and its CI/CD pipelines, with Jenkins and SonarQube quality gates and CloudWatch-driven production debugging.

ClearMechanic
Feb 2021 – Aug 2021

Software Engineer · Mexico

- **Led third-party integrations**, mapping the platform's API data schema to external systems so new clients and suppliers could connect their services to the platform.
- **Built and shipped internationalized automotive software** with Node.js (Serverless) and React, owning backend logic, API security, servers, and releases across DEV, QA, and Production.

Earlier Experience

Softtek · Software Engineer — Feb 2020 – Feb 2021

- Built internal Angular management tools for a banking client and implemented organizational security policies.

Secretaria de Finanzas · Software Engineer — Jun 2019 – Feb 2020

- Digitized municipal reporting and internal scheduling on self-managed servers, including authentication.

Dritdi · Software Engineer — Jun 2017 – Jun 2019

- Co-built a medical-education startup (3D training app licensed to schools), owning the full stack and VPS deployments.

TECHNICAL SKILLS

AI Engineering: AI Agents & Agent Harnesses, RAG (Retrieval-Augmented Generation), Evals, Fine-Tuning, Embedding Systems, Vector Databases, Model Context Protocol (MCP), Custom LLM Skills & Agentic Tooling, Foundation Model APIs & LLM Integration, Prompt Engineering, AI-Assisted Code Modernization.

AI Governance, Security & Cost: AI Governance & Usage Policies, Responsible AI, Role-Based Access & Usage Controls, LLM Guardrails & Safety, Data Protection for AI Workloads, Rate-Limiting, Token & Cloud Cost Optimization, Consumption Monitoring.

Programming Languages: Node.js, TypeScript / JavaScript, Go, Python, SQL.

Cloud & Infrastructure: AWS (S3, EC2, ECS, Lambda, SQS, SNS, CloudWatch, Secrets Manager), Azure (Blob Storage, Virtual Machines, AKS, Functions, Monitor, Key Vault, Azure AD / Entra ID, Data Factory, Azure Batch), Kubernetes, Containerization, Helm, Terraform, Microservices, REST & GraphQL APIs, System Design, Event-Driven Systems & Streaming (Kafka), SQL & NoSQL Databases (PostgreSQL, MySQL, MongoDB, Redis).

Observability: Metrics, Logging, Distributed Tracing, Real-Time Monitoring & Alerting, SLOs, Dashboards, Incident Response.

CI/CD & Delivery: Argo CD, Canary & Blue-Green Deployments, Jenkins, SonarQube, Git / GitFlow.

Spoken Languages: Spanish (Native), English (Advanced, C1).

Leadership: End-to-End Ownership & Autonomy, Platform Team Leadership (~20 engineers), Client-Facing Solution Engineering, Technical Discovery & Scoping.

EDUCATION

Universidad Tecnologica de Tabasco · Mexico — 2017 – 2019
B.S. in Information Technology Engineering

Universidad Tecnologica de Tabasco · Mexico — 2015 – 2017
Associate Degree in IT, Electronic Commerce (TSU TI)