# Rafael C. Delgado

Senior Backend | Full-Stack Engineer | Tijuana, Mexico | (+1) 619-728-8925
co.rafaeldelgado@gmail.com | linkedin.com/in/rafaelcdelgado | https://github.com/rafaelcdelgado | https://rafaelcdelgado.com

## Summary

Senior Software Engineer with 10+ years building high-performance, data-intensive backend systems handling 40M+ records. Specialized in SQL performance optimization, rule-driven financial workflows, and resilient background processing across MySQL, PostgreSQL, and MariaDB environments.

Proven track record reducing multi-second query paths (~30s) to sub-second latency through execution-plan analysis, composite index design, and OpenSearch optimization. Experienced designing idempotent financial pipelines, hierarchical data models, and Redis-backed async systems that maintain correctness under high concurrency.

Hands-on engineer focused on data integrity, performance, and scalability in mission-critical transactional systems. Strong TypeScript experience across Next.js and React applications, enabling rapid ramp into JavaScript/Node ecosystems.

**Backend & Data Systems**

Ruby on Rails and Laravel for production-grade backend services.
Advanced SQL performance tuning across MySQL, PostgreSQL, MariaDB.
Execution-plan analysis (EXPLAIN), composite index strategy, and access-path optimization.
Design of high-volume transactional data flows and hierarchical data models.
Financial workflow automation with strong data-consistency guarantees.
REST API design focused on reliability, observability, and maintainability.

**Infrastructure, Queues & Performance**

Redis-backed caching and Sidekiq background processing pipelines
Idempotent job design and retry-safe async workflows for financial data
Application-level concurrency safeguards to prevent double-processing
Priority-segmented queue architectures and worker pool isolation
Read-replica traffic distribution and high-throughput ingestion systems
Concurrency tuning for SLA-sensitive workloads

**Search & Indexing**

OpenSearch / Elasticsearch at scale (40M+ records)
High-volume indexing pipelines with near-real-time ingestion
Sub-second search performance tuning and query optimization

**Quality & Testing**

Automated testing and unit coverage (RSpec / Minitest) with strong reliability and maintainability practices.

**Frontend (Supporting)**

TypeScript (production usage across Next.js and React applications)
React, Tailwind CSS, ShadCN

## Experience

### Independent Consultant — Senior Backend / Full-Stack Developer
Remote | May 2025 – Present

Built an AI-assisted financial intake pipeline that extracts, validates, and routes invoice and purchase order data into accounting workflows.

Designed rule-driven pricing matrix overriding fixed profit logic based on product price tiers, with deterministic rounding and separation of canonical vs computed pricing.

Implemented time-window controls across Accounts Payable and Accounts Receivable workflows with idempotent safeguards for audit consistency.

Engineered high-speed PDF → structured data ingestion system with validation rules to ensure financial data accuracy.

Designed and optimized self-referential chart-of-accounts schema (parent/subaccount model) with targeted indexing for fast traversal and aggregation.

Implemented Stripe subscription billing and automated AR ledger generation ensuring transactional integrity across the billing lifecycle.

Reconciled payment data across Clover POS, internal systems, and QuickBooks, improving financial visibility and reducing discrepancies.

Designed scalable backend APIs and data pipelines emphasizing correctness, observability, and performance.

## Cabal — Senior Software Engineer

San Francisco Bay Area (Remote) | Nov 2023 – May 2025 | Full time | Contract

Engineered high-throughput background processing pipelines handling tens of millions of records using Redis-backed Sidekiq workers.

Reduced multi-second query paths (~30s) to sub-second latency through targeted index redesign, execution-plan analysis, and access-path optimization.

Designed idempotent job workflows and concurrency safeguards to prevent duplicate processing of financial and transactional data.

Implemented priority-segmented queue architecture (critical / important / default) with dedicated worker pools, improving SLA reliability.

Tuned Sidekiq concurrency, bulk enqueue, and batching strategies to increase queue throughput under heavy load.

Built and maintained OpenSearch indexing pipelines supporting 40M+ searchable records with near-real-time ingestion.

Partnered cross-functionally with Product and GTM to deliver data-heavy features in a fast-moving startup environment.

## Tyrannosaurus Tech — Senior Laravel & Ruby on Rails Developer

Atlanta, GA (Remote) | Nov 2022 – Sept 2023 | Full time | Contract

Designed and implemented scalable backend services and REST APIs supporting high-volume transactional workloads.

Diagnosed and resolved database performance bottlenecks through query refactoring, composite indexing, and query-plan optimization.

Investigated complex production incidents and implemented durable fixes to improve system reliability.

Mentored junior engineers and enforced structured code review practices.

Contributed to system architecture decisions focused on long-term scalability and predictable performance.

## Grupo AFAL — Senior Web Developer

Tijuana, Baja California, Mexico | Mar 2021 – Sep 2022 | Full time

Led Ruby on Rails development initiatives delivering scalable web solutions aligned with dynamic business needs.

Applied Test-Driven Development (RSpec) to ensure long-term code reliability and maintainability.

Enforced CI/CD and code review practices supporting consistent deployments.

Proactively monitored application health using AppSignal and resolved production issues.

Mentored team members and served as technical liaison with clients.

## Education
Universidad Autónoma de Baja California — Ingeniería en Computación (2007–2012)

## Languages
Spanish (Native), English (Professional)

## Professional Strengths

- Cross-functional collaboration with Product, Design, and GTM teams
- Strong ownership of reliability, performance, and data correctness
- Mentorship and code review leadership in distributed teams
- Clear technical communication with technical and non-technical stakeholders

## Certifications
Data Science and Machine Learning: Making Data-Driven Decisions — MIT IDSS (2024)