

# DANIEL JIMÉNEZ RIVERA

Valencia, Spain · [LinkedIn](#) · [GitHub](#) · [dajimenezriv@gmail.com](mailto:dajimenezriv@gmail.com)

## SUMMARY

---

Backend engineer with 7+ years in Go and TypeScript, focused on distributed systems: Kafka ingestion at ~26M events/day, gRPC services, and observability built to alert on work performed rather than liveness. Most recently code owner of a 1M-user, 5-petabyte desktop sync client, including severity-1 incident response.

## EXPERIENCE

---

### Internxt

*Senior Software Engineer*

Valencia, Spain

*Apr 2025 – Jul 2026*

- Code owner of [Internxt Drive for Windows](#) (Electron, Node.js, C++), serving 1M+ users and managing 5+ petabytes of data, contributing to an open-source platform ranked #20 in the Deloitte Technology Fast 50.
- Led a simplicity-first refactor of an unstable codebase, cutting customer-reported issues by 80% and increasing test coverage from 7% → 45% while reducing the codebase from 32k → 15k lines.
- Replaced Node's file watcher with a custom Windows file system watcher written as a C++ native addon (node-gyp), dropping idle CPU consumption from ~30% at 50k files to near 0%.
- Redesigned the synchronization engine around checkpoint-based sync, recovery sync and lazy loading, scaling the client from ~10k files to 500k+ without degradation.
- Owned incident response for severity-1 production issues — data loss, app crashes, memory leaks, auth failures and failed migrations — driving root-cause analysis, permanent fixes and written postmortems.

### Kurita Group

*Backend Engineer*

Remote

*Feb 2023 – Apr 2025*

- Re-architected the sensor ingestion and alerting pipeline of [Kurita Connect360](#) (Go, Kafka, Redis, Snowflake), covering ~18,000 IoT sensors at ~26M events/day, cutting ingest to alert latency from ~10s to ~0.8s.
- Eliminated silent event loss from a watermark poller that skipped rows committed out of `id` order, moving to offset-based Kafka consumption with manual commits, non-blocking retry topics and a replayable DLQ.
- Cut ~\$8k/month of Snowflake compute by eliminating polling loops that kept a warehouse permanently resumed across dev, staging and production, and moving enrichment upstream to a Kafka + Snowpipe load.
- Extracted alerting from the backend into an independent Go service after a schema migration silently killed the in-process goroutine while health checks stayed green; backend reduced to a gRPC control plane.
- Instrumented the pipeline with Prometheus, Loki, Tempo and Grafana, alerting on work performed rather than liveness: event freshness, throughput floor, metric absence, consumer lag, external watchdog.
- Implemented authentication and authorization with Microsoft Entra ID, validating JWTs against rotating signing keys and enforcing role-based access control across services.

### Kenmei Technologies

*Backend Engineer*

Remote

*Sep 2020 – Dec 2022*

- Built a large-scale analytics platform for tier-1 telecom providers (e.g., Telefónica, Vodafone) using Node.js, PostgreSQL and Google BigQuery, processing big data to detect network anomalies in near real time.
- Led the backend migration to TypeScript, generating automated OpenAPI 3 documentation, enforcing end-to-end type safety and cutting integration errors.
- Deployed a distributed observability stack (Loki, Grafana) with Helm and Google Cloud Storage, improving system monitoring, log aggregation and incident diagnosis across environments.
- Implemented a frontend mock server using Mock Service Worker, enabling fast local development by simulating big-data APIs and reducing request latency from 10+ seconds to near-instant feedback.

### Sendmei

*Full Stack Engineer*

Valencia, Spain

*Nov 2018 – Aug 2020*

- Developed a logistics SaaS platform (Django, React, PostgreSQL) integrating shipping carriers like CTT and GLS and ecommerce platforms like Shopify and WooCommerce into a unified system.
- Engineered an event-driven microservices backend using Django REST Framework, GraphQL, Celery workers, RabbitMQ and Django Channels for real-time WebSocket collaboration.
- Deployed on Docker with Nginx reverse proxy (SSL/TLS, HTTP/2, WebSocket upgrade), Unicorn/Uvicorn dual-server setup and multi-environment pipelines (dev/staging/prod).

## EDUCATION

---

### Universitat Politècnica de València

*Computer Engineering, Artificial Intelligence*

Valencia, Spain

*Sep 2018 – Jan 2023*

- Mentored and taught 10 classmates during second- and third-year courses.
- Authored a [paper](#) optimizing the Spectral Clustering algorithm, benchmarked against scikit-learn's DBSCAN, KMeans+ and Spectral Clustering implementations.

## TECHNICAL SKILLS

---

**Languages:** Go, TypeScript, Python, C++.

**Backend:** Node.js, gRPC, REST/OpenAPI, GraphQL, WebSockets, Redis, RabbitMQ, Kafka, PostgreSQL, MongoDB.

**Infra & DevOps:** Docker, Kubernetes (Helm), CI/CD, Google Cloud, Azure DevOps.

**Data & Analytics:** Snowflake, Google BigQuery.

**Observability:** Prometheus, Loki, Tempo, Grafana.

**Desktop:** Electron, Node.js native addons (node-gyp, C++), Windows APIs.

## LANGUAGES

---

**Spanish:** native proficiency.

**English:** full professional proficiency.