

SUMMARY

Backend engineer focused on performance-critical data systems, distributed storage, and AI-powered dynamic pricing. I build systems that need to be fast, scalable and maintainable, often migrating from the legacy version.

EXPERIENCE

Smartness – Backend Engineer

Apr 2022 – Present · Rovereto, Italy

Backend engineer on **Smartpricing**, the AI-powered RMS inside the Smartness platform (4,000+ properties in Europe).

- Rewrote the **core pricing engine** that produces all dynamic prices for **4,000+ hotels**; migrated gradually with no downtime.
- Built from scratch an **event scraping and deduplication pipeline**, using geospatial clustering and connected-components graph algorithms to keep stable identifiers across runs.
- Co-developed the **competitor price scraper**, its APIs, and a **ScyllaDB** storage layer holding **~300 GB** compressed at a **~100x compression ratio**.
- Led multiple **legacy-to-modern migrations** of backends and infrastructure with **zero-downtime** cutovers.

SELECTED PROJECT

Competitor Pricing – Compression-Friendly Storage Layer

2024 – 2025

ScyllaDB-backed storage for scraped competitor prices, designed from the schema up for extreme compression while keeping reads fast for the pricing engine.

- Stores **~300 GB** of competitor pricing data at a **~100x compression ratio**, serving **~100 QPS** to the pricing engine and downstream services.
- Designed the **partition and clustering key layout** so that the partitions scale horizontally and the rows are compression-friendly. The biggest factor is the custom row format, which is a binary vector-like encoding of the prices.
- Built the **ingestion path** from the scraper into Scylla via **batching, Kafka and stream processing**, keeping write amplification low and absorbing bursts.
- Migrated from the old format to the new one, **cutting storage by 90%**.

TECHNICAL SKILLS

Languages	Python, TypeScript, SQL
Backend / APIs	FastAPI, Fastify, REST, async / event-driven services
Data / Storage	PostgreSQL, Redis, ScyllaDB, Cassandra, Kafka, Neo4j
Data analysis & visualization	Pandas, NumPy; interactive data visualization with Plotly, Seaborn, but also custom frontend+backend tools with Vue/Svelte and FastAPI/Fastify
Infra & Ops	Docker, Kubernetes, Helm, GCP, Grafana, CI/CD
Practices	Performance optimization, zero-downtime migrations, parameter- and data-driven design, observability

EDUCATION

Università di Trento – M.Sc. Computer Science, 110/110 cum laude

2019 – 2022

Master of Science in Computer Science and Information Technologies, with a specialization in Data Science.