Rahul Jeyasingh Software Engineer

rahuljeyasingh.com rascimat@gmail.com +91 8838933383 linkedin.com/in/rahuljeyasingh

Profile

Backend engineer with 4.5 years building production microservices and data systems on Google Cloud. Designed a config-driven record-linkage service running fuzzy matching at scale on Cloud Run. Co-author of **CANalyse**, presented at **Black Hat Arsenal, Las Vegas, 2022**. Interested in distributed systems, developer tooling, and clean API design.

Skills

Languages: Python , SQL

Backend: FastAPI, REST APIs, microservices, async I/O, Pydantic, Redis

Cloud & Infra: GCP (Cloud Run, BigQuery, Pub/Sub, Dataproc, IAM), Terraform, Docker, Kubernetes

Data Systems: Apache Airflow, dbt, PySpark, Polars, Pandas

Other: Git, CI/CD, system design, Linux, embedded hardware (ESP32, Raspberry Pi, Arduino)

Professional Experience

Infosys

Digital Specialist Engineer — Aug 2025 – present, Chennai, India

- Refactored 5+ core dbt ELT models in BigQuery, cutting weekly engineering maintenance overhead by 30% while maintaining full data correctness across downstream consumers.
- Built and deployed PySpark jobs on GCP Dataproc processing 5+ TB/day, reducing job completion time by 30% through partition tuning and shuffle reduction.

Digital Specialist Engineer — Sep 2022 – Jul 2025, Berlin, Germany

- **Designed and shipped a config-driven record-linkage microservice** (Python, deployed on GCP Cloud Run) that performs generic similarity matching between any two datasets via YAML configuration. Combines trigram + TF-IDF candidate retrieval with Jaro-Winkler and weighted-ratio scoring; processes claims data with **12× throughput** over the previous workflow.
- Implemented a dependency-resolution layer using **topological sort over a DAG** of processing stages, enabling parallel execution of independent steps and clean failure-recovery semantics.
- Owned the service's infrastructure-as-code in **Terraform** (Cloud Run, IAM, Pub/Sub, networking), reducing manual provisioning time by 70% and lifting deployment reliability to 99.9%.

Digital Specialist Engineer — Aug 2021 – Sep 2022, Chennai, India

- Designed and developed Apache Airflow ETL pipelines ingesting **100+ GB** of daily API and SFTP data into Google BigQuery. Reduced operational cloud costs by **40%**.
- Delivered **15+** data models adopted by **30+** cross-functional stakeholders, reducing report generation time from 3 days to under 5 minutes.

Open Source Contributions

CANalyse — CAN bus reverse-engineering toolkit, Co-author · Presented at Black Hat Arsenal, Las Vegas, 2022
https://github.com/canalyse/CANalyse-2.0

- Python tool for analysing, filtering, and replaying CAN protocol traffic from vehicle ECUs. Built session-packet capture, SQL querying over packets, and a replay-attack module. Featured at Black Hat USA Arsenal as one of the showcased security tools of 2022.

Education

B Tech, Amrita School of Engineering Jun 2017 – Jun 2021