

ADARSH RAJ

Bengaluru, India | adarsh18raj@gmail.com | +91 9206543619 | [linkedin.com/in/adarsh18raj](https://www.linkedin.com/in/adarsh18raj)

PROFESSIONAL SUMMARY

Senior Data Engineer with **6+ years** of experience in architecting and scaling production-grade data platforms across InsurTech, FinTech, Healthcare, E-commerce, and Logistics. Hands-on expertise in **Snowflake, dbt, Apache Airflow, Dagster, Kafka, and Spark** across AWS, GCP, and Azure, building real-time streaming pipelines, cloud data warehouses, modern data stack and analytics infrastructure. Senior engineering experience at YC and Techstars-backed startups, with a track record of cutting cloud costs by **40%**, accelerating data operations **6x**, and compressing report runtime from **6 hours to 10 minutes**.

TECHNICAL SKILLS

Languages: Python, SQL, Bash

Data Warehouses and Databases: Snowflake, Google BigQuery, AWS Redshift, PostgreSQL, MySQL, MongoDB

Orchestration and Transformation: Apache Airflow, Dagster, dbt (Core and Cloud)

Streaming and Big Data: Apache Kafka, Apache Spark, Real-time Event Processing, Event-Driven Architecture

Data Engineering: ETL/ELT Pipeline Development, Dimensional Data Modeling, Schema Design, Change Data Capture (CDC), Performance Tuning, Cost Optimization, SLA Monitoring

API Development: FastAPI, Flask, Django REST Framework, REST API Design, OAuth 2.0/JWT, OpenAPI/Swagger, API Versioning, Microservices, Async Programming

Cloud Platforms: AWS (S3, Lambda, ECS, RDS, Glue), GCP (BigQuery, Cloud Run, Pub/Sub), Azure (ADLS Gen2, Synapse, Fabric, Data Factory)

DevOps and Tooling: Docker, Terraform, Git, GitHub Actions, CI/CD

PROFESSIONAL EXPERIENCE

Senior Data Engineer, CoverWhale

Feb 2025 - Present

- Architected and maintained real-time telematics data pipelines on Snowflake for a US-based commercial trucking insurance platform, ingesting high-volume vehicle and driver event streams.
- Owned the Dagster orchestration repository end to end, leading pipeline architecture, deployment, and production operations across the data platform.
- Designed dimensional data models in Snowflake using dbt for policy, claims, and operational analytics, with version-controlled transformations and automated lineage documentation in dbt Cloud.
- Implemented dbt CI/CD via GitHub Actions, cutting transformation deployment cycles and reducing data incidents reaching production.

Senior Data Engineer, Ramani (YC W20)

Jan 2024 - Jan 2025

- Engineered scalable ETL pipelines using Apache Airflow and Kafka, processing **5TB+** of data across MongoDB and PostgreSQL to support real-time supply chain analytics.
- Established SLA monitoring and pipeline observability across MongoDB and PostgreSQL workloads, reducing incidents and improving operational reliability.
- Mentored a **team of 3** junior data engineers through 1-on-1 sessions and structured workshops on advanced ETL patterns and big data tooling.

Founding Data Engineer, Complya (Techstars)

Apr 2023 - Dec 2023

- Designed and implemented the company's core data, backend, and cloud infrastructure from scratch, laying the foundation for a scalable compliance analytics platform.
- Developed APIs and backend services powering a responsive user experience across the analytics platform.
- Optimized AWS architecture and achieved a **40% cost reduction** while maintaining high availability.
- Mentored a **team of 2** engineers, raising standards in code review and deployment practices.

Senior Software Engineer, Data Engineering, Alopa Infotech

Sep 2020 - Mar 2023

- Pioneered end-to-end Snowflake architecture and data pipelines using Apache Airflow and dbt, establishing the modern data stack and powering real-time analytics for cross-functional product and operations teams.
- Accelerated data operations **6x** and cut data discrepancies by **50%** through query optimization and pipeline re-architecture across Snowflake.
- Led migration of legacy Oracle DB workloads to Snowflake with a team of 4 engineers, cutting daily report runtime from **6 hours to 10 minutes** and storage costs by **25%**.
- Optimized Snowflake architecture for security and performance, reducing system downtime by **70%**.
- Built Python REST APIs (Django REST Framework, FastAPI) for a healthcare software product on AWS and MySQL, scaling to handle **3x user load**.

EDUCATION

Bachelor of Engineering (B.E.), Computer Science and Engineering

KLE Technological University, Hubballi, Karnataka, India

Aug 2016 - Jun 2020