

SAIKUMAR CHINTADA

Backend & Data Platform Engineer | Tech Lead | NIT Calicut

+91 8693858644 • saikumar.chintada.one@gmail.com • smarchone.substack.com
linkedin.com/in/saikumar-chintada-a72671b1 • github.com/smarchone

Summary

Backend and data platform engineer with 8 years, including 3 as Tech Lead. Experienced in distributed backend and orchestration systems in Golang and Python — an open-source platform scheduling 5000+ DAGs, batch compute savings, a zero-downtime TB-scale migration. Depth in event-driven microservices, Postgres, Airflow, Kubernetes and more.

Skills

Languages: Golang, Python, SQL, JavaScript

Data Platform: Airflow, Trino, Hive, dbt, ETL/Reverse ETL, Semantic Layer

Infrastructure: Kubernetes, Docker, Terraform, AWS, GCP, Hetzner

Databases: Postgres, MongoDB, Redis

Experience

Independent R&D | Self-directed | Agentic SDLC & Developer Tooling Apr 2026 – Present

- Building an agentic SDLC workflow (github.com/smarchone/skills) based on hexagonal architecture grounded in a domain ontology.
- Prototyping meeting-driven development, deriving specs from role-based meeting transcripts.

Senior Software Engineer | Sundial, Bangalore Nov 2025 – Mar 2026

- Led developing a semantic layer for exploratory analytics directly over base tables, with a query layer translating metric queries into table queries (Python, Golang, Trino, Kubernetes, Redis).
- Diagnosed and led resolution of a critical concurrency bug in a custom pipeline orchestrator that had caused silent SLA misses for months.

Freelance Consultant | Early-Stage Startup, Remote Jan 2024 – Oct 2025

- Prototyped checkpointing with intent-based speculative scaling, cutting user-facing job latency from 20 minutes to 1 minute.
- Built a cost-efficient k3s-based load testing cluster on Hetzner Cloud.

Tech Lead, Data Platform | Gojek, Bangalore Mar 2022 – Jan 2024

- Led a team of 4 engineers developing an open-source job specification and orchestration layer (Optimus) scheduling 5000+ DAGs for BigQuery ingestion, egress and transformation.
- Cut the release cycle from 3 days to 1 hour by simplifying the plugin system and execution model; added YAML-based plugins and deprecated 90% of legacy plugins.
- Upgraded legacy Airflow clusters, reducing Airflow DB CPU usage by 80% and lowering scheduling lag.
- Achieved 70% cost reduction on batch workloads by fixing issues that blocked Kubernetes node pool scale-down.

Tech Lead, Data Platform | JioSaavn, Bangalore Jan 2020 – Mar 2022

- Built a centrally governed Reverse ETL service and migrated 50+ ad-hoc pipelines onto it, enabling non-technical analysts to run push, email and ad campaigns.
- Orchestrated a zero-downtime MongoDB migration (12 shards, 36 nodes, TB-scale) from AWS to Azure; achieved 10x faster shard operations with a custom chunk migration solution.
- Reduced MongoDB cluster size by 33% through selective archiving, deletion and downscaling.
- Drove Kubernetes adoption by migrating production Trino and Airflow clusters off VMs.

Software Engineer | Coverfox, Mumbai May 2018 – Oct 2019

- Accelerated feature delivery in a finance module from 2 weeks to under an hour by extracting it from the monolith into an event-driven microservice, removing it from the monolith's full regression cycle (Python, Flask, AWS Lambda, Postgres)
- Reduced offline financial reconciliations from weeks to hours by onboarding offline policy workflows and unifying the commission payout system.
- Supervised a non-technical 3-person QA team in building a test suite for complex payout calculations by developing a JSON-based testing framework.

Software Engineer | CouponDunia, Mumbai Jul 2016 – May 2018

- Owned a full-stack sub-product and browser extension; improved coupon auto-apply time at checkout from 30s to 2s.

Education

National Institute of Technology Calicut | B.Tech, Computer Science | **CGPA 7.65** | **2016**