SAIKUMAR CHINTADA

Lead Software Engineer | Backend & Data | NIT Calicut

+91 8693858644 • saikumar.chintada.one@gmail.com • smarchone.substack.com

linkedin.com/in/saikumar-chintada-a72671b1 • github.com/smarchone • smarchone.github.io

SUMMARY

Backend and data platform engineer, 8 years (2+ as Tech Lead), building distributed systems in Golang and Python: an open-source orchestrator scheduling 5000+ DAGs, 70% batch cost savings, a zero-downtime TB-scale migration.

SKILLS

Languages: Golang, Python, SQL, JavaScript
Data Platform: Airflow, Trino, Hive, dbt, BigQuery, ETL/Reverse ETL, Semantic Layer
Infrastructure: Kubernetes, Docker, Terraform, AWS, GCP, Azure, Hetzner
Databases: Postgres, MongoDB, Redis
Practices: Distributed Systems, Event-Driven Microservices, Workflow Orchestration, Agentic SDLC

EXPERIENCE

Independent R&D | Agentic SDLC | https://smarchone.github.io/ Apr 2026 – Present

- Building agentic experiences to improve productivity and accelerate learning.
- Consulting for an early-stage startup, building its AI platform from scratch.

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

- Led a semantic layer for exploratory analytics over JSON log tables (Python, Golang, Trino, Kubernetes, Redis).
- Ended months of silent SLA misses by fixing a major concurrency bug in an in-house data pipeline orchestrator.

Career Break | Freelance consulting for an early-stage startup, Remote Jan 2024 – Oct 2025

- Cut user-facing job latency from 20 min to 1 min via checkpointing and intent-based speculative scaling.
- Delivered a low-cost k3s-based load-testing cluster on Hetzner Cloud.

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

- Led 4 engineers on Optimus, an open-source orchestrator scheduling 5000+ DAGs for BigQuery ETL.
- Cut release cycle from 3 days to 1 hour via a simpler plugin model; moved 90% of legacy plugins to YAML.
- Reduced Airflow DB CPU usage by 80% and scheduling lag by upgrading legacy Airflow clusters.
- Cut batch workload costs by 70% by unblocking Kubernetes node pool scale-down.

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

- Built a governed Reverse ETL service, migrating 50+ ad-hoc pipelines so analysts self-serve campaigns.
- Migrated TB-scale MongoDB (12 shards, 36 nodes) AWS to Azure, zero downtime; 10x faster chunk moves.
- Moved production Trino and Airflow from VMs to Kubernetes; shrank the MongoDB cluster by 33%.

Software Engineer | Coverfox, Mumbai May 2018 – Oct 2019

- Cut feature delivery from 2 weeks to under an hour by extracting a finance module from the monolith into an event-driven microservice (Python, Flask, AWS Lambda, Postgres).
- Reduced offline financial reconciliations from weeks to hours by unifying the commission payout system.

Software Engineer | CouponDunia, Mumbai Jul 2016 – May 2018

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

EDUCATION

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