

MAITHILI GADDE

Senior Software Engineer

Hyderabad, India • maithili.gadde@gmail.com • maithiligadde.com

PROFESSIONAL SUMMARY

Senior Software Engineer with 10+ years of experience building and scaling large-scale data platforms and distributed systems. At Nielsen, led a full on-prem-to-AWS migration of the Collections & Crediting platform, running both environments in shadow mode for 3–6 months to ensure data consistency, cutting a critical pipeline runtime by ~93% and saving \$900K annually. Led a team of 7 engineers through a high-attrition period, maintaining every SLA while simultaneously rebuilding the team. Currently building production-grade AI-integrated applications and deepening expertise in agentic systems and LLM pipelines. Now seeking senior engineering roles where deep data infrastructure expertise meets modern AI/ML systems.

TECHNICAL SKILLS

Data Processing	Spark, Kafka, Hadoop/MapReduce, Hive, Airflow
AWS Services	EMR, Lambda, Glue, Step Functions, SNS/SQS, API Gateway, ECS, CloudFormation, S3, Redshift, EC2, CloudWatch
Warehousing	Snowflake, Redshift, S3, Redis
Languages	Python, Scala, Java, SQL
AI & Tools	LLM Integration, Prompt Engineering, Anthropic Claude API, Cursor, OpenAI Codex, n8n, Agentic Workflows
Other	Docker, CI/CD, Infrastructure as Code, React Native

PROFESSIONAL EXPERIENCE

Independent Engineering Projects

Jan 2025 – Present

Hyderabad, India

- Built MomPop Care, a React Native + TypeScript caregiver app for elder-care families with a local-first, patient-scoped data model on AsyncStorage using prefix-based isolation. Features include age-aware BMI bands, per-patient threshold overrides, medication adherence tracking, symptom logging, and AI-generated weekly caregiver summaries.
- Integrated Anthropic Claude via a hardened Node proxy for lab-report extraction. Backed by 128 Jest tests, CI pipeline (lint + typecheck), and EAS-built Android/iOS dev clients.

Senior Software Engineer, Lead

Sep 2018 – Dec 2024

Nielsen Media: Audience Measurement • New York, NY / Remote

- Led the full on-prem-to-AWS migration of the Collections & Crediting platform, from architecture design through production rollout, enabling the system to scale to billions of audience records.
- Ensured data consistency during migration by running both environments in shadow mode for 3–6 months, with automated comparison jobs flagging discrepancies in record counts, aggregates, and downstream outputs.
- Cut a critical pipeline runtime from 2h 30m to 10–12 minutes (~93% reduction) by eliminating small-file overhead, resolving data skew via salting, replacing shuffle-heavy joins with broadcast joins, and tuning serialization, saving \$900K annually in compute costs.
- Held the production system together single-handedly during a high-attrition period; kept every SLA intact while recruiting, onboarding, and ramping a replacement team of 7 engineers.
- Led the team of 7 engineers post-rebuild: ran weekly reliability reviews tracking job runtimes, failure rates, and near-misses; mentored engineers on system design, AWS best practices, and production debugging.
- Built and maintained pipelines spanning batch, near-real-time, and streaming workloads using Spark, Kafka, Redis, and Airflow orchestrated on EMR.
- Worked hands-on with a broad AWS stack: Lambda and Step Functions for serverless orchestration, SNS/SQS for messaging, Glue for cataloging, ECS/Docker for containerized services, and CloudFormation for IaC.

- Drove the data integration approach between collection (raw audience signals) and crediting (rule-based preprocessing), ensuring clean handoffs to downstream analytics teams.
- Partnered with product managers, analysts, and cross-functional stakeholders to scope requirements and break them into deliverable engineering work.

Big Data Engineer

Jun 2016 – Feb 2018

Warner Music Group • New York, NY

- Built ETL pipelines processing millions of music-play and billing records daily from platforms like Spotify and iTunes, powering royalty calculations across Warner’s entire catalog.
- Developed hands-on Spark and Hive expertise building end-to-end batch workflows on AWS (EMR, Redshift, S3) supporting daily billing runs across all Warner-owned content.
- Created log analyzers and transformation jobs that normalized usage data from multiple distribution channels into a single billing-ready format.
- Worked closely with business and product teams to translate complex royalty and billing requirements into concrete, shippable engineering tasks.

Systems Engineer

Sep 2012 – May 2014

Infosys Ltd • India

- Contributed to development and support of Infosys Finacle, a core banking platform deployed across 100+ countries, working on transaction processing modules, Java-based application layers, and Oracle database integration for retail and corporate banking clients.
- Gained foundational experience with production systems, structured debugging, Agile delivery workflows, and client-facing requirements gathering across multiple banking engagements.

EDUCATION

Master of Science in Computer Science • New York University (NYU)

2014 – 2016

Bachelor of Technology in Computer Science • JNTU Hyderabad

2008 – 2012