
EDUCATION

Indian Institute of Technology (IIT), Bhubaneswar, India
Bachelor of Technology (Electronics and Communication Engineering)

July 2017 - June 2021
GPA - 7.97

EXPERIENCE

Convozen.ai (By Nobroker) & Nobroker.com, Bengaluru India
Senior Engineer (Backend & Data)

Oct 2022 to Present

Senior Member of Convozen.ai by Nobroker.com, a B2B SaaS conversation analytics platform.

- Designed and implemented comprehensive canned dashboards on Convozen, **leading a team of three** and **leveraging NoBroker's data lake** to deliver actionable insights and enhance user experience.
- Led a team of 2 developers to integrate innovative call highlights tagging technology into Convozen.ai by merging three projects into one, **reducing tagging time by 90%** in some cases and improving code maintainability and efficiency.
- Engineered an automated ingestion pipeline that pulls data from client SFTP servers via cron jobs, integrating it into the product pipeline and **reducing manual effort by 99%**.
- Implemented a **robust CI/CD pipeline for ML models on Kubernetes**, automating workflows, ensuring scalability, and **reducing deployment time by 80%** compared to VMs.
- Built a bulk-job platform with real-time tracking by migrating processes from Django to Spring Boot, enhancing performance, increasing throughput, and **reducing escalations by 50%**.

Data Engineer

Managing data operations for one of the largest property-tech platforms in India, handling around 130 TB of data and ensuring accessibility to end users with minutes-level latency

- **Designed and implemented end-to-end data ingestion pipelines** for NoBroker's Data Lake, integrating diverse sources such as MySQL, MongoDB, Elasticsearch, and custom data systems using **Apache Hudi, Kafka, and Strimzi Kafka Connectors**.
- Spearheaded the development of a monitoring and alerting system for NoBroker's data pipelines, **cutting end-user complaints by 90%** and **reducing issue identification from hours to minutes**, enhancing operational efficiency.
- Transformed a manual data pipeline deployment process into a streamlined, UI-driven workflow, **cutting deployment time from 2-3 hours to just 2-3 minutes**, improving productivity across the team.
- Developed a cloud data pipeline management framework, optimising spark job scheduling and resource allocation, leading to a **20% reduction in cloud costs**.
- Successfully migrated all data pipelines to **Apache Hudi v0.13** from **v0.5**, improving data ingestion speed, query performance, and storage efficiency while ensuring backward compatibility and minimal downtime during the upgrade.

Sigmoid, Bengaluru, India

June 2021 to Oct 2022

Software Development Engineer - 1

Core member of the Structured Data Engineering team at Goldman Sachs, processing 500M+ records daily from 50+ sources

- Led the strategic migration of table joins from MapReduce to Spark, **cutting runtime by 80%** and **reducing production issues by 90%**, resulting in a more stable, efficient, and scalable pipeline.
- Led efforts to optimise data storage by implementing innovative data compression techniques in **Hbase**, resulting in cost savings and **improved storage utilisation by 30%**.
- Designed and implemented a robust ingestion framework in **Scala** to onboard and process ~1 TB of data daily from multiple data feeds
- Onboarded and integrated over **50 diverse data feeds**, transitioning them from an outdated pipeline to a more robust system, streamlining operations and ensuring scalability.

SKILLS

- **Programming Languages** - Python, Java, JavaScript, C++, SQL, Scala
- **Frameworks** - Spring boot, Django, Flask
- **Big Data & Data Engineering** - Apache Hudi, Flink, Spark, Kafka, Hadoop, MapReduce, Hive, HBase, Presto, Trino
- **ORMs and Databases** - Django, Elasticsearch, MongoDB, MySQL, Redis, SQLAlchemy
- **Tools** - Google Cloud Platform, Docker, Kubernetes, Kafka, Git