

Soorya Ravichandran

sooryaprasanna.com

About

I architect and build scalable data platforms. I have 10+ years in distributed systems, data lakes, and real-time processing at enterprise scale. At Microsoft, I build the data infrastructure behind AI-driven security.

Skills

- Java, TypeScript, JavaScript, SQL, Python, JUnit, REST, CI/CD, Git, Azure DevOps, NoSQL.
- Azure, AWS, Flink, HDFS, HBase, Spark, MapReduce, Kafka, Iceberg, Snowflake, Avro, Parquet.
- Distributed Systems, Event-Driven Architectures, Domain-Driven Design, Cloud Computing.

Experience

Microsoft · Senior Software Engineer

Apr 2025 – Present

- Microsoft Security · Exposure Management
 - Leading infrastructure scaling strategy for 10x load increase across HBase, Solr, Hadoop and Azure services to support 180k enterprise tenants, designing critical bottleneck pipelines.
 - Architected Global Insights to Lake initiative end-to-end, defining scalable ingestion architecture and implementing core components including label-derived applicability model and metadata dataset.
 - Implemented guardrails in Microsoft's Discovery services to prevent high-blast radius blocking events, validating safety across 81 million IPs and reducing operational risk.

Goldman Sachs · Vice President

Jan 2021 – Apr 2025

- Data Lake Engineering → Lead for the firm's enterprise Data Lake platform, overseeing core services (ingest, sandbox and simulator) to support reporting and analytics across ~160K unique datasets ingesting ~450K daily data pipelines.
 - Spearheaded the data & operational catalog, and built the ingest framework using HBase, Flink, HDFS, AWS S3 and Hadoop YARN.
 - Developed a production-like simulation test framework that preemptively caught bugs, reducing post-deployment defects by 75% and enabling first-time-right releases.
- Core Data Engineering → Lead Architect for Lake to Lakehouse migration, designing and implementing ingest migration architecture to enhance performance and reduce costs.
 - Designed and implemented workflows to decommission 25% of Lake refiners & datastores, migrating unused data to cold storage and reducing quarterly storage costs by ~\$500K.
 - Engineered parallel ingest pipelines for seamless migration, improved scalability, and enhanced failover reliability.
 - Collaborated with Snowflake Engineers to develop the Snowpark Migration Accelerator, automating refiner categorization and streamlining migration to Lakehouse.

Goldman Sachs · Associate Software Engineer

Jan 2020 – Dec 2020

- Architected and developed key components of Data Lake, including AWS S3 Ingest service, Data Partitioner, Snowflake Ingest service, Avro-Parquet Scanner & Converter utility, ensuring high reliability and on-time delivery.
- Built a Simulator tool that enabled early bug detection, fostered developer mindfulness, and improved performance tuning, resulting in an 80% increase in customer satisfaction.

Goldman Sachs · Analyst Software Engineer

Nov 2017 – Dec 2019

- Contributed to the core backend metadata services for data cataloging, ingestion, and refinement in an early stage enterprise Data Lake; migrated ingest framework from MapReduce to Flink, reducing latency by 25% and memory usage by 33%.
- Built a robust testing framework, standardized for 100% releases, ensuring stable deployments and minimizing post-release firefighting.

The University of Texas at Dallas · Curriculum Developer

Jun 2017 – Oct 2017

- Led the Big Data Club as the primary instructor, designing hands-on projects, course content, and learning pathways in collaboration with the Center for CS Education & Outreach to enhance student engagement, coding skills, and Big Data mastery, earning the CS Outreach Excellence Award for impactful contributions.

Infosys · Software Engineer

Sep 2013 – Jul 2015

- Developed a high-precision Emergency Alert Solution with 100% predictability, warning Australian communities of fire and flood hazards while delivering mission-critical ETL data to Telstra's Network Decision Support Database.

Infosys · Software Engineer Intern

Feb 2013 – May 2013

- Developed an online shopping website that markets seasonal clothes based on customer geolocation, rapidly prototyping new features to validate feasibility and user engagement.

Education**The University of Texas at Dallas**

2015 – 2017

Master of Science in Computer Science · Data Science Major

Anna University, India

2009 – 2013

Bachelor of Engineering in Computer Science

Leadership**Mentor for Emerging Leaders Series · Goldman Sachs**

2024 – 2025

- Mentored sophomore students to cultivate essential industry skills that will help them become better engineers and explore opportunities in Tech.

Mentor for BAI/HLAI Network · Goldman Sachs

2021 – 2025

- Mentored interns, junior and mid-level engineers through the Black Analyst Initiative and Hispanic/Latinx Analyst Initiative to promote diversity within Goldman Sachs.

Americas Campus Recruiting Coordinator · Goldman Sachs

2019 – 2025

- Led the recruiting team to identify, source, and hired 100+ engineering talents across the U.S. and Canada, ensuring alignment with business needs for the Goldman Sachs Engineering Division.

Big Data Instructor · UT Dallas

2016 – 2017

- Developed competency among 500+ students and 300+ industry professionals by conducting hands-on workshops and seminars on Big Data Technologies.
-