

Mohammed Kaish Ansari

| iiitu17131@gmail.com | [LinkedIn](#) | [Github](#) | [Leetcode](#)

PROFESSIONAL SUMMARY

Software Engineer with 5+ years of experience designing and maintaining scalable, high-performance data pipelines and AI-integrated developer tooling. Experienced in building LLM-driven agentic automation pipelines and taking AI capabilities from experiment to production in fintech environments, including multi-tenant GenAI platform infrastructure and agentic workflow systems. Proficient in Scala, Python, C++, AWS, Microservices, SQL, and NoSQL databases. Skilled in competitive programming and dedicated to delivering efficient, reliable solutions.

EDUCATION

Indian Institute of Information Technology Una <i>Bachelor of Technology in Computer Science and Engineering</i>	H.P., IN <i>Jul 2017 – May 2021</i>
JPMorgan Chase & Co. <i>Quantitative Research Virtual Experience Program</i>	Virtual <i>June 2025</i>

TECHNICAL SKILLS

Languages: Scala, Python, C++, Java, TypeScript, Go
AI & LLMs: LLM Integration, Agentic Systems, RAG, Prompt & Context Engineering, LangChain, LangGraph, Pydantic AI, OpenAI SDK, Model Context Protocol (MCP), GitHub Copilot
AI & Developer Tools: VS Code Extensions, LLM Observability & Evaluation, Experiment-to-Production AI Pipelines
Development: REST API, Amazon Web Services (EMR Serverless, EC2, S3, Athena, SQS, RDS, IAM), Kafka, Docker, Kubernetes, Solace
Tools and Frameworks: Airflow, Hadoop, Spark, Sqoop, Hive, Snowflake, dbt, Terraform, Prometheus, Grafana, Node.js, Git
Databases: Postgres, MySQL, Oracle, DynamoDB, Redis, HBase, Impala
Authentication: ADFS, SAML SSO, ROPC, OAuth 2.0, Kerberos

EXPERIENCE

Software Engineer II - (Liquidity Risk, Big Data) <i>JPMorgan Chase & Co.</i>	Feb 2024 – Present <i>Glasgow, UK</i>
Applied AI & Developer Productivity - VS Code Extension: Architected and developed an enterprise VS Code extension integrating 11+ MCP (Model Context Protocol) servers to expose internal tooling (Bitbucket, Confluence, Jira, Swagger, Jules CI, Igloo/Databricks) as GitHub Copilot tools, enabling LLM-driven agentic automation workflows for enterprise engineers. Built custom MCP servers in TypeScript/Node.js including Knowledge MCP (semantic search across 100+ documentation repositories), Terminal MCP (stateful command execution with guardrails), Skynet MCP (AWS EMR Serverless monitoring and S3 log analysis), and Memory MCP (persistent cross-conversation context storage). Established best practices for AI observability, reliability, and context engineering across all agentic pipelines. Took multiple AI capabilities from prototype to production, iterating on reasoning quality and output reliability with structured evaluation frameworks. Implemented unified ADFS authentication layer supporting SAML SSO, ROPC, and cookie-based flows across all MCP servers with lockout prevention. AWS EMR Serverless Migration: Led migration of critical Stress Analytics workflows from EC2 EMR to AWS EMR Serverless, achieving 30–40% cost reduction and eliminating cluster management overhead. Migrated 15+ Spark-based pipelines including Stress Display Rules (RLAP 100/366), DBox Transformations, and Finance Pre-calculations. Optimized Spark configurations transitioning from YARN-based memory overhead to EMR Serverless memory overhead factor, improving resource utilization. Implemented production parallel testing strategy running EC2 and Serverless pipelines side-by-side to validate data parity before cutover, with comprehensive sanity validation across 1M+ records. Asset Class Model Integration and Stress Alignment: Led transition of Stress data pipelines to AWS utilizing Spark and Snowflake for efficient data migration. Coordinated comprehensive testing and reconciliation processes to ensure seamless alignment with strategic datasets. Streamlined datatype changes and deprecated columns, enhancing data consistency and operational efficiency across platforms.	

- **Performance Tuning and Cost Reduction:** Optimized large-scale Spark jobs for Liquidity Risk Stress calculators, resulting in 20% reduction in processing time for daily stress calculations, enabling timely regulatory reporting to the Federal Reserve. Implemented partition optimization and Spark SQL tuning improving query performance by 25% and reducing daily ETL runtime from 120 minutes to under 90 minutes.
- **Upgraded Stress API Application:** Upgraded Stress API's Moneta Boot, Scala, and Java versions to latest JPMC-approved versions. Developed an automated EMR cluster management system using custom APIs and OAuth 2.0, reducing cluster provisioning time by 60% and optimizing resource utilization across multiple Spark workloads.

Software Engineer - (Postpaid, Big Data)

July 2021 – Feb 2024

Airtel Digital (Bharti Airtel Ltd.)

Gurugram, IN

- **Real-Time Fraud Prevention and Dedupe System:** Architected and implemented a real-time Fraud Prevention and Deduplication Engine processing 550M+ customer records, leveraging Apache Spark for batch processing, Kafka for real-time data ingestion, and developing highly concurrent REST APIs in Go that achieved sub-second response times for fraud detection queries.
- **Data Reconciliation System:** Built an end-to-end Data Reconciliation System using Apache Spark, Golang, and Python, orchestrating data integrity checks with Apache Airflow to automatically identify and resolve data discrepancies, while implementing a real-time notification system through Telegram bots that reduced Mean Time To Resolution (MTTR) by 40%.

INTERNSHIPS

Bumblebee Data Inc. | *Python, Spark, NLP, RAPIDS*

- Written and modified programs, primarily using Python, Apache Spark, using Agile development methodology.
- Identified opportunities for Machine Learning functionalities and integration in Bumblebee's Software.

Applied Science Lab - IIT Mandi | *Python, ETL Pipeline, Machine Learning and EDA*

- On-site internship in Farmer Zone Project under Dr. Srikant Srinivasan, Learned and implemented various aspects such as Outlier Detection, various Techniques of Data Cleaning and Developed a script that can be used for cleaning Data in Real-time (Data Pre-processing).

HONORS & AWARDS

Systematic Trading Competition Award 2025

Secured 3rd Place

JPMorgan Chase & Co.

Nov 2025