

Gokul Sreekumar

Bengaluru, Karnataka, India

gokulsreekumar2@gmail.com — +91 8547058985 — github.com/gokulsreekumar — linkedin.com/in/gokul-sreekumar-63581b174

Summary

Software engineer with experience building trading, booking, and data platforms across equities and derivatives. Skilled in backend services, event-driven architectures, and AI-enabled data workflows using Java, Python, Spring Boot, Kafka, and cloud infrastructure.

Experience

Goldman Sachs

Bengaluru, Karnataka, India

Associate (Software Engineer II) — Asset & Wealth Management

Nov 2024 – Present

- Led end-to-end delivery of a cross-market equities trading and booking platform, enabling UHNW sales desks to trade and book equities across the top 10 Asian markets, extending an existing US-only system to multi-market scale.
- Owned requirements, architecture, development, testing, regression sign-offs, and production rollout, partnering with Sales, Global Markets, Execution Services, and Books & Records teams.
- Implemented FIX-based integrations for order placement, executions, partial fills, and average-price allocations, and built booking workflows ensuring correct ledgers, cash journals, commissions, and Asia-specific charges.
- Developed backend services using Spring Boot with a Kafka event-driven architecture, contributed to React-based order-entry and blotter UI, and implemented observability including health checks, end-of-day reports, and PagerDuty alerts across markets and time zones.
- Leveraged AI-assisted analysis and development using LLMs and internal tooling for requirements analysis, debugging, and generating exhaustive instruction files to enable minimal hands-on (automated) Spring migration from a proprietary GS build framework.
- Contributed to internal AI-driven data platforms ingesting Confluence documentation and GitLab repositories, enabling semantic search, contextual retrieval, and LLM-assisted SQL query generation for data exploration.

Analyst (Software Engineer I) — Asset & Wealth Management

Jun 2022 – Oct 2024

- Modernized derivatives booking pipelines by eliminating legacy flows and migrating to JMS-based ingestion with robust filtering, validation, and fault handling for OTC trades.
- Built a MongoDB-backed staging layer to improve traceability, debugging, and downstream corrections; implemented complex cash and equities movement workflows.
- Re-architected a file-based booking ingestion system integrating with a secure firmware platform for external institutional clients; designed an extensible framework supporting Excel, CSV, TSV, and proprietary formats with plug-and-play parsers and a normalized trade abstraction.

Summer Analyst — Consumer & Wealth Management

Jun 2021 – Jul 2021

- Designed and implemented a pre-trade disclosure document delivery workflow using Java and Spring, integrating REST services and messaging components.
- Gained exposure to front-office systems and regulatory constraints through close collaboration with senior engineers and business teams.

Education

National Institute of Technology Calicut

Kerala, India

Bachelor of Technology in Computer Science and Engineering (CGPA: 8.95)

Apr 2018 – Jun 2022

Technical Skills

Languages: Java, Python, JavaScript, C/C++, SQL

Backend & Distributed Systems: Spring Boot, REST APIs, Apache Kafka, JMS, Event-driven architectures, FIX protocol

Databases & Storage: MongoDB, SQL, DynamoDB, GemFire (OQL)

Cloud & Infrastructure: AWS (EC2, S3, IAM, VPC), Docker, Kubernetes, Argo CD

AI / Developer Tooling: GitHub Copilot (GPT, Claude), OpenAI APIs, Large Language Models (LLMs), LangChain, Model Context Protocols (MCP)

Observability & DevOps: GitLab CI/CD, Grafana, PagerDuty, JIRA

Certifications

Generative AI with Large Language Models — DeepLearning.AI (AWS Team), June 18, 2025

Deep Learning Specialization — DeepLearning.AI (Andrew Ng), August 16, 2020

Projects

DataDistillerAI — Knowledge & Data Exploration Engine

Python, FastAPI, Kafka, FAISS, LangChain, PostgreSQL, MinIO

- Built an end-to-end RAG system that ingests unstructured documents, performs cleaning and semantic chunking, and indexes content for retrieval-based analysis.
- Designed an async ingestion pipeline using FastAPI + Kafka queues, enabling non-blocking uploads and horizontal scaling via worker-based processing (ingestion and embedding stages).
- Implemented embedding + vector search retrieval with FAISS to ground LLM responses in source passages, improving factuality and reducing hallucinations.
- Built modular prompt/retrieval workflows for Q&A, summarization, and exploratory analysis; integrated multiple LLM backends (OpenAI/Ollama/Gemini/Claude) via LangChain.
- Added production-grade infrastructure for persistence, including PostgreSQL for job metadata, S3-compatible object storage (MinIO).

NITCBase — Educational Relational Database System

C / C++

- Designed and implemented an educational relational database system from scratch to demonstrate core RDBMS concepts including storage layout, indexing, query execution, and transaction handling.
- Focused on system design trade-offs and data structures underlying elementary database engines to help students reason about internals.
- Authored comprehensive documentation and step-by-step implementation guides using Docusaurus; project adopted into the Database Management Systems curriculum at NIT Calicut.

Kerberos Authentication Protocol — RFC-Based Implementation

Java, Networking

- Implemented the Kerberos authentication protocol by following the RFC specification, including ticket-granting flows, message formats, and cryptographic exchanges.
- Built a multi-VM test setup to simulate protocol participants and validated correctness through controlled attack scenarios.

eXpOs — Experimental Operating System

C, OS Internals

- Built an experimental operating system for a simulated machine (XSM), implementing kernel components including interrupts, system calls, paging, and file/disk I/O.
- Designed kernel routines and exception handling mechanisms to manage process execution, memory, and user authentication.