

MEGHNA NAG

720-878-0666 | meghnanag2@gmail.com | [linkedin.com/in/meghnanag](https://www.linkedin.com/in/meghnanag) | github.com/meghnanag2 | San Francisco, CA

SUMMARY

Software Engineer with 6.5 years in distributed systems, cloud engineering, and full stack development using Java, Python, and SQL. Strong foundation in object-oriented design and Agile delivery. MS in Data Science, University of Colorado Boulder.

LANGUAGES AND TECHNOLOGIES

Programming Languages: Java, Python, TypeScript, JavaScript, SQL, Go, Scala, R, Bash

Backend & APIs: Spring Boot, Spring MVC, FastAPI, Flask, RESTful APIs, gRPC, Protobuf, GraphQL, OAuth2, JDBC, JPA, JMS

Frontend: AngularJS, AJAX, React, jQuery, HTML, CSS

Databases: PostgreSQL, PostGIS, MySQL, Oracle, Redis, BigQuery, NoSQL

Data Engineering: Apache Spark, Hadoop/HDFS, MapReduce, Kafka, RabbitMQ, MinIO

Cloud Tech: AWS (AppFlow, Glue, Lambda, S3, IAM, CloudWatch, CodePipeline, CloudFormation, RDS), GCP (Pub/Sub, Cloud SQL, Dataproc)

DevOps & Tools: Docker, Kubernetes (GKE), CI/CD, GitHub Actions, Dagger, Prometheus, Grafana, Git, JIRA; SDLC

AI/ML & Visualization: Scikit-learn, PyTorch, Pandas, NumPy, Random Forest, MLP, CNNs, NLP, Transformer Models, Cross-Attention, Embeddings, Computer Vision, LLMs, SHAP, spaCy, NLTK, Tableau, Matplotlib, Seaborn, SciPy

EXPERIENCE

Software Engineer, Persistent Systems

Jan 2022 - May 2024

Team: Amazon Web Services - AppFlow & Glue Team

- Built a configurable, multi-tenant AppFlow Connector SDK enabling new SaaS integrations through connector-specific configuration, reducing per-integration development time by 66%; maintained regulatory compliance by updating integration logic to reflect SaaS-side schema and API contract changes across connectors
- Built the AppFlow ingestion framework used across 79 SaaS integrations, enabling low-latency data transformation and transfers of up to 100 GB per flow into Amazon S3, eliminating the need for custom ETL pipelines
- Engineered a scalable AWS Glue ETL framework with data modeling and schema evolution across native and connector-based integrations, standardizing pipelines and eliminating custom ETL jobs across multiple data sources
- Mentored 7 engineers as primary technical point of contact for requirements gathering and architectural design in an Agile environment; drove a 50% improvement in delivery velocity within 3 months

Software Engineer, Tata Consultancy Services

Dec 2016 - Dec 2021

Team: Solidarity Insurance / New India Assurance

- Reduced dashboard load time by 94% (180s to under 10s) through materialized views, stored procedures, and optimized SQL queries
- Architected an async document generation pipeline post policy approval, generating 7 policy documents in under 10 seconds using multithreaded processing (CompletableFuture), replacing slow sequential batch jobs
- Automated insurance premium recalculation for vehicle make/model updates, replacing manual uploads with a UI-integrated backend workflow, reducing manual errors by 50%
- Built and automated ETL data pipelines replacing manual batch processes for staging table loads, configuring logging mechanisms to monitor and alert on data transfer failures
- Migrated a legacy EJB application to a microservices architecture using Spring MVC and AngularJS, building user interfaces with JUnit test coverage, and enforcing code quality standards using SonarQube in a CI/CD environment

EDUCATION

Master of Science in Data Science, University of Colorado Boulder

Aug 2024 - May 2026

Relevant Courses: Data Structures and Algorithms, Datacenter Scale Computing, Machine Learning, Deep Learning

Bachelor of Technology in Information Technology, West Bengal University of Technology

Aug 2012 - Jul 2016

Relevant Courses: Data Structures and Algorithms, Operating Systems, Database Management Systems, Pattern Recognition

PROJECTS

- ElectroHub (Scalable Cloud Marketplace):** Built a distributed electronics marketplace on event-driven microservices using consistent hashing with virtual nodes for PostgreSQL sharding, token bucket rate limiting via atomic Redis Lua scripts, Kafka KRaft consensus for activity streaming, RabbitMQ for async notifications, and Redis Pub/Sub for real-time WebSocket chat. (Python, FastAPI, Kafka, Redis, PostgreSQL, Kubernetes)
- Secure Routes & Risk Prediction for FIFA World Cup 2026:** Developed a geospatial analysis dashboard using GeoPandas, PostGIS, and OSMnx to analyze crime data and generate risk-adjusted route recommendations with interactive map visualization.
- FetalVisionAI (Prenatal CHD Risk Screening):** Built a multimodal deep learning system fusing fetal ultrasound image and clinical tabular data via cross-attention; applied Focal Loss and Platt scaling achieving 0.814 AUROC and 77% sensitivity. Deployed with a multi-agent RAG system, drift monitoring, and auto-retrain pipeline. (PyTorch, Go, gRPC, FastAPI, Kafka, Airflow, MLflow)
- Predicting Interview Performance with Explainable ML:** Built an interview scoring system using MLP and Random Forest achieving 1.9x better accuracy than GPT-2 prompting; integrated SHAP and LLM-based reasoning for per-prediction explainability.

PROFESSIONAL HONORS

HighFive Awards (Sep 2022, Dec 2023) | On the Spot Award (Nov 2022) | Bravo Team Awards (Jun 2022)