Shitij Mathur

shitijkmathur@gmail.com | +1 (623) 277-7854 | Tempe, AZ | linkedin.com/in/shitijmathur

EDUCATION

Master's In Computer Science **August 2024 - May 2026**
Arizona State University, Tempe, AZ GPA: 4.0
Coursework: Cloud Computing, Statistical Machine Learning, Data Mining, Data Processing at Scale

Bachelor of Technology in Electronics & Communication Engineering **July 2015 - July 2019**
Manipal Institute Of Technology, Karnataka, India

SKILLS

Languages: Python, SQL, JavaScript, Bash
Backend: REST APIs, WebSockets, Microservices, Event-Driven Architecture
Data: Apache Airflow, Pandas, NumPy, PostgreSQL, DynamoDB, Redshift, OpenSearch
ML/AI: LangChain, RAG Pipelines, Vector Databases, PyTorch, SageMaker
Tools: Docker, Kubernetes, Git, AWS (Lambda, ECS, SQS, S3, API Gateway)
Certification: AWS Machine Learning - Specialty, AWS Certified Solutions Architect - Associate

PROFESSIONAL WORK EXPERIENCE

Quantiphi **Boston, MA (Remote)**
Platform Engineer Intern (Backend & Data Systems) **May 2025 - January 2026**

- Built reusable infrastructure modules in Python using AWS CDK to provision data platform components for multiple domain teams. Enabled config-driven onboarding where new domains self-served by updating configuration files. *(CDK, MWAA, Glue, Step Functions)*
- Wrote Python scripts using Redshift Data API to automate DDL deployments, enabling schema changes to propagate across environments without manual intervention.

Quantiphi **Mumbai, India**
Senior Platform Engineer (Backend & ML Systems) **January 2021 - May 2024**

- Designed and built a RAG-based document retrieval system in Python that indexed 500+ technical PDFs. Implemented chunking logic, vector indexing, and prompt construction to serve queries via a conversational interface, reducing search time by 95%. *(OpenSearch, LangChain, AWS Bedrock)*
- Built an event-driven document processing pipeline in Python where file uploads triggered async workers to parse, chunk, and index documents. Implemented retry logic and dead-letter queues to handle failures during concurrent processing. *(S3, SQS, ECS)*
- Implemented comprehensive testing strategies, including unit tests, mocking, and performance testing to ensure reliability under load. *(API Gateway, Lambda, pytest, JMeter)*
- Built a real-time chat backend in Python using API Gateway WebSockets for persistent connections and DynamoDB for session state. Integrated with Amazon Lex for NLU, automating 80% of support queries. *(WebSockets, Lambda, DynamoDB)*

Quantiphi **Mumbai, India**
Platform Engineer (Backend & Cloud Systems) **July 2019 - January 2021**

- Built an event-driven pipeline in Python to process financial risk reports asynchronously. Implemented retry handling and dead-letter queues using SQS, improving reliability and reducing manual processing effort by 70%. *(Lambda, DynamoDB, PostgreSQL)*
- Wrote ETL scripts in Python using Pandas to extract, transform, and validate large datasets. Designed Airflow DAGs to orchestrate daily batch jobs processing millions of records. *(Airflow, Redshift)*
- Developed RESTful APIs in Python with authentication and rate limiting to provide secure access to ML models and data systems. *(API Gateway, Lambda, Cognito, WAF)*
- Deployed ML models as REST endpoints using Flask and Docker containers on SageMaker, enabling scalable inference for downstream applications. *(Flask, Docker, SageMaker, ECS)*

PROJECTS

Distributed Graph Processing Pipeline **November 2025**

- Implemented a real-time distributed graph processing system using Kafka, Kubernetes, and Neo4j to ingest streaming events, construct graph-based data models, and execute graph analytics (PageRank, BFS) for streaming transportation data.

HONORS & ACTIVITIES

Zoom Fellowship - ASU Next Lab **November 2025 - Present**

- Selected for the Zoom Fellowship, a competitive innovation program focused on developing next-generation educational applications on Zoom's developer platform.

AWS Road to re:Invent 2025 Hackathon **November 2025**

- Selected among 50 developers globally for a live-streamed on-the-road hackathon from LA to Las Vegas.
- Built "Lucky Loo," a multi-agent AI system where LLM sub-agents analyze facial expressions and deliberate to decide whether to reveal the nearest restroom location; implemented with AWS Bedrock (Claude), Strands Agents SDK, Groq (TTS), and a Flask API.

Hack SoDA 2024 Hackathon | Winner - 2nd Place | Amazon's Online Safety Track [Link] **October 2024**

- Built PassGen, a Chrome extension that generates unique, strong passwords for each user account using multi-layered hashing.