Vishal Srivastava

vishalgosrivastava@gmail.com | +1 949-795-8058 | LinkedIn | GitHub

Software Engineer with **3+ years** of expertise in the finance domain, specializing in **large-scale data pipelines, distributed systems, and cloud-based automation on AWS**. Proficient in Python and SQL, with a strong track record of delivering high-performance, reliable data solutions. Passionate about **Machine Learning, Agentic AI, and modern data architectures**. Skilled in optimizing workflows, designing scalable architectures, and implementing cost-efficient cloud solutions in compliance-focused environments.

EDUCATION

University of California, Irvine — *Sep 2024– Dec 2025*
MS in Data Science (GPA: 3.75) — *Irvine, United States*

SRM Institute of Science & Technology — *July 2017–June 2021*
Bachelor of Technology in Computer Science & Engineering (GPA: 3.9) — *Chennai, India*

SKILLS:

Programming Languages Python, Java, C++, R

Cloud & Distributed Systems : AWS(Glue,Kinesis,EC2,S3,Lambda,DynamoDB),Docker, Kubernetes,Concurrency,Fault Tolerance

Data Engineering & Databases : MySQL, MongoDB, Vector Databases, Hadoop, Spark, Kafka, ETL Pipelines

Machine Learning & AI :PyTorch, TensorFlow, Scikit-Learn, LLMs, Generative AI, NLP, AI Agents, Langchain, MCP, LangGraph

Developer Tools & Libraries: Flask, FastAPI, Git, CI/CD, HuggingFace

EXPERIENCE

Software Developer - Data | BARCLAYS , India — *Feb 2022 – June 2024*

- **Built and migrated Distributed data platform** at Barclays from legacy on-prem to AWS, implementing fault-tolerant orchestration (Step Functions) and scalable ETL (Glue/Lambda) to improve data availability and reduce operational complexity.
- Designed and ran an event-driven notification system for 250,000+ customers, using Kinesis for streaming ingestion, EventBridge for routing, SQS for buffering/backpressure, and Lambda for stateless processing
- **Built automated data-quality and schema enforcement** across 40,000+ attributes, achieving 99.8% compliance and cutting downstream data errors by 60% via rule-based validation, standardized exception handling, and failure routing for fast triage and recovery.
- **Implemented Runbook Automation**, embedding OpsItem creation for proactive production issue handling.

***Tech Stack**: AWS (Glue, Step Functions, Lambda, IAM, S3, DynamoDB, SQS, Kinesis, EventBridge, CloudFormation, Bedrock), Python, SQL, Tableau, Jenkins, Bitbucket, Gitlab, JIRA, Django*

Software Developer | FORD MOTOR COMPANY , India — *Oct 2021 – Feb 2022*

- **Created reusable post-deployment health check framework**, reducing incidents. Automated iPlan creation on Confluence, saving 7–8 PD/month across domains.
- **Knowledge Graph & Ontology Automation** – Automated data workflows, reducing 90% manual effort. Led multiple high-reliability production deployments, providing fast bug resolutions.

Graduate Student Researcher | University of California, Irvine — *June 2025 – Sept 2025*
UCI Environmental Health Lab | Advised by Prof. Jun Wu

- Improved object detection accuracy by 23% by implementing a customized CenterNet-ResNet50 architecture integrated with **CNN-based Super-Resolution (ESPCN)**, identifying trucks and containers in low-resolution satellite imagery of Southern California with 94% precision and 89% recall.
- Enhanced model interpretability and deployment efficiency by integrating **Residual Deblurring modules** and implementing automated data preprocessing pipelines using PyTorch, reducing inference time by 35% while maintaining model accuracy for environmental health research applications.

PROJECTS

FlashDB | Consistent Hashing, Synchronization, Multithreading, Replication — GitHub

- Implemented an LSM-tree storage engine with an in-memory MemTable and on-disk SSTables to improve write throughput and support efficient lookups.
- Implemented a background compaction process that merges SSTables in a merge-sort style and removes tombstones, running in parallel while the database continues to serve reads and writes.
- Designed a consistent hashing scheme with virtual nodes to distribute data across nodes and reduce reshuffling when the cluster scales.

Reddit Moderation Framework | Kafka, Flask, AI Agents, NLP, LoRA Finetuning — GitHub

- Architected a **Multi-Agent NLP** System utilizing **Mistral-7B** (Triage Agent) to normalize noisy text (slang/acronym expansion) and **Llama-3** (Repair Agent) to diagnose model failures via K-Means semantic clustering, enabling an Adapt Agent to automate **PEFT/LoRA fine-tuning** on subreddit-specific data.
- Engineered a **scalable real-time pipeline** featuring a **Flask interception proxy** that injects toxicity overlays directly into the DOM, supported by a **Kafka streaming architecture** for high-throughput data ingestion and a Streamlit dashboard for live monitoring and feedback collection.