AHMED SOLIMAN

Software Engineer / Architect

U.S. Citizen | New York City Metropolitan Area | +1 347-755-4426 LinkedIn | ahmedesoliman.com | GitHub

PROFESSIONAL SUMMARY

Software Engineer and Architect at the intersection of Enterprise Data Strategy and Generative AI, with 4 years at BNY designing and leading enterprise-grade distributed data platforms in financial technology. Currently leading the enterprise-wide Storm-to-Snowflake migration at BNY and architecting the development framework for the Eliza Mesh data ingestion platform. Deep expertise in Apache Spark orchestration, Azure Kubernetes Service (AKS) multi-cluster deployments, and end-to-end CI/CD automation. Proven track record designing real-time streaming and batch ingestion systems processing high-volume financial datasets. Skilled in building production-grade AI agents leveraging retrieval-augmented generation (RAG) and LLM-driven workflow automation. Focused on technical leadership — architecting solutions that are compliant, scalable, and built for long-term operational efficiency.

EXPERIENCE

BNY (Bank of New York Mellon) — Software Engineer / Full Stack Engineer
New York City Metropolitan Area | 2022 – Present

- Led and co-led major platform initiatives at BNY, including the enterprise-wide Storm-to-Snowflake migration (Lead), the Eliza Mesh ingestion framework (Co-Lead), and the production RAG Agent Platform (Lead); accountable for solution architecture, rollout strategy, validation testing, and cross-team coordination.
- Architected and deployed core Global Data Platform components: Spark processing engine, Vertica loader, event-driven trigger service, and REST APIs.
- Operated and optimized distributed Spark workloads on multi-cluster AKS, ensuring high availability, fault tolerance, and production monitoring.
- Automated CI/CD pipelines with GitLab, Docker image management, and Helm chart releases; integrated security scanning and SonarQube quality gates.
- Built and maintained Elasticsearch indices and dashboards for ingestion monitoring and event metadata analytics.
- Developed API automation and integrations for RAG agents, enabling role-aware data discovery, alert triage, and log summarization.
- Designed and built an Angular dashboard that provided real-time visibility into ingestion operations, rejection analytics, and processing statistics to accelerate troubleshooting.
- Implemented compliant backout plans, production support processes, and vulnerability remediation; shared DevOps best practices across teams.
- Contributed to the convergence initiative between two major data platforms, aligning architecture, deployment processes, and operational practices to reduce redundancy and improve maintainability.

RF CUNY (Research Foundation of CUNY) — Software Engineer

New York City Metropolitan Area | 2020 – 2022

- Developed full-stack web applications and RESTful APIs with modern frontend frameworks, responsive design, and intuitive UX
- Developed web applications and backend services using Java, Python, and JavaScript, following best practices for code quality, testing, and maintainability
- Built backend services handling authentication, data persistence, and third-party API integrations
- Designed and deployed production websites with a focus on performance, accessibility, and SEO (Google Analytics, conversion tracking)
- Built and maintained CI/CD pipelines for automated testing, code quality checks, and deployment to cloud platforms
- Gained foundational expertise in software architecture, relational databases, and agile development methodologies

EDUCATION

Arizona State University (ASU)
M.S. Computer Science, 2022 – Present

City University of New York (CUNY)
B.S. Computer Science & Information Security, 2016 – 2022

Relevant Coursework: Data Processing at Scale, Discrete Structures, Advanced Programming Techniques *(Honors)*, Data Structures *(Honors)*, Cryptography and Crypto Analysis, Applied Cryptography, Computer Systems and Networks, Software Engineering, Database Systems, Operating Systems, Computer Architecture, Information Security Principles

TECHNICAL SKILLS

- **Languages:** Java 17+, Python, JavaScript/Node.js, Bash/Shell, PowerShell, SQL
- **Big Data & Processing:** Apache Spark, Apache Storm, Apache Airflow, Apache Kafka, Apache Iceberg, Nessie, Vertica, Snowflake, Apache Kudu
- **Search & Metadata:** Elasticsearch, Painless Scripting, Elastic SQL
- **Cloud & Infrastructure:** Microsoft Azure, Azure Kubernetes Service (AKS), Helm, Docker, ZooKeeper
- **DevOps & CI/CD:** GitLab CI/CD, SonarQube, Maven, JaCoCo, Docker Image Pipelines, Automated Change Management
- **APIs & Microservices:** RESTful APIs, JWT/OAuth, Spring WebFlux, Microservice Architecture
- **Frontend Frameworks:** React, React Native, Angular, Responsive Design, UI/UX
- **Monitoring & Observability:** Splunk, Elastic Stack, Spark History Server
- **AI/ML & Automation:** AI Agent Development, RAG, Prompt Engineering, LLM Integration, Workflow Automation
- **Security & Governance:** SailPoint, Vulnerability Scanning, Threat Modeling, Kubernetes RBAC, SoD Compliance
- **Databases:** PostgreSQL, Vertica, Snowflake, SQL

- **Other:** React Native, IBM Blockchain, Google Analytics, Jira, Git, Insomnia/Postman

---

LANGUAGES

- English (Native/Bilingual)
- Arabic (Native/Bilingual)

PROJECTS

- **Storm-to-Snowflake Migration (BNY):** Led migration of streaming and batch pipelines from Apache Storm to Apache Spark and Snowflake, improving data latency and maintainability across the enterprise. Designed data migration strategies, validation checks, and backout plans used during phased rollouts.
- **Eliza Mesh Data Ingestion (BNY):** Co-led the design and delivery of a scalable, governed ingestion framework that standardizes onboarding for Data Lake tenants. Implemented resilient retry semantics, schema evolution handling, and automated onboarding pipelines.
- **RAG Agent Platform (BNY):** Architected and delivered a production retrieval-augmented generation platform used for automated alert triage, log summarization, and intelligent data discovery. Integrated vector search, LLM orchestration, and role-aware access controls.
- **RAG Demo App**: Retrieval-Augmented Generation demo (FastAPI backend, FAISS vector store, React frontend) showcasing embeddings, vector search, and LLM-driven Q&A.
- **ML Data Pipeline**: Python library for composable machine learning data transformations, featuring 8 transformer classes and config-driven pipelines.
- **Distributed Task Queue App**: Python/FastAPI + React system for async job processing, real-time status tracking, and distributed worker pools using Redis and RQ.
- **Real-Time Analytics Dashboard**: Full-stack dashboard streaming live metrics via WebSockets, with React/Recharts frontend and FastAPI backend.
- **K8s Deployment Platform**: Kubernetes management platform with deployment CRUD, pod scaling, YAML manifest generation, and a modern React UI.
- **Dev CLI Toolkit**: Typer-based Python CLI for Docker/Kubernetes operations, project scaffolding, and system health checks.
- **Pcap Analyzer**: Python tool for analyzing and visualizing network packet capture (PCAP) files, supporting protocol breakdown and traffic statistics.
- **LongInt Project**: C++ library implementing arbitrary-precision integer arithmetic for cryptography and numerical computing.
- **Stock Trader App**: Interactive stock trading simulator with real-time price updates, portfolio management, and order simulation built with Vue.js.