

AHMED SOLIMAN

Software Engineer / Architect

U.S. Citizen | New York City Metropolitan Area | +13477554426 | [LinkedIn](#) | ahmedsoliman.com

PROFESSIONAL SUMMARY

Experienced Software Engineer and Architect in financial technology with 6+ years building and operating enterprise-grade distributed data platforms. Deep expertise in Apache Spark job orchestration, Kubernetes-native deployments across multi-cluster environments, and end-to-end CI/CD pipeline automation. Proven track record designing real-time and batch data ingestion systems processing high-volume financial datasets. Skilled in containerized microservice architecture using Docker, Helm, and cloud-native tooling, with strong proficiency in Elasticsearch and data lake technologies (Iceberg, Nessie). Passionate about AI/ML innovation, currently building AI-powered agents and exploring LLM-driven workflow automation.

PROFESSIONAL EXPERIENCE

BNY (Bank of New York Mellon) — Software Engineer Full Stack Engineer

New York City Metropolitan Area | 2022 – Present

Platform Engineering & Data Processing

- Designed and implemented core components of a Global Data Platform including the Spark processing engine, Vertica analytical loader, event-driven trigger service, and data ingestion framework REST APIs — from early design through production deployment
- Operated distributed Spark workloads across multi-cluster Kubernetes environments spanning development, QA, and production with high availability and fault tolerance
- Built and maintained Elasticsearch indices for file tracking, event metadata, and data certification; wrote Painless scripts for bulk metadata transformations across the platform
- Led the convergence initiative between two major data platforms, aligning architecture, deployment processes, and operational practices to reduce redundancy and improve maintainability
- Upgraded core platform components to Java 17 and managed Apache Iceberg table maintenance and Nessie catalog operations to ensure data lake reliability

DevOps & CI/CD Automation

- Constructed a unified GitLab CI/CD pipeline for all platform components with integrated change management, multi-project triggers, and structured Maven release versioning
- Automated Docker image build and publish workflows across a private container registry, managing deployment artifacts for 6+ distributed services including Storm, ZooKeeper, and Spark Operator
- Managed Helm chart releases for Kubernetes-based microservices and configured RBAC policies for secure Spark job submission across namespaces
- Implemented end-to-end deployment automation for legacy and modern components including artifact download, image build, verification, and cloud-native deployment

- Integrated automated security vulnerability scanning into CI/CD pipelines, reducing the vulnerability backlog and strengthening the organization's security posture

Risk Management & Production Support

- Developed and executed compliant backout plans with rollback procedures for Kubernetes deployments and cloud-hosted services
- Diagnosed and resolved production issues using Spark driver/executor logs, Splunk queries, and Elasticsearch analytics across platform indices
- Proactively remediated platform vulnerabilities in critical services by patching, rebuilding, and redeploying containerized images
- Managed Kubernetes secrets, Spark UI debugging, and StatefulSet/Pod lifecycle operations across production clusters
- Supported SailPoint Certification reporting, SoD violation remediation, and metadata rebranding automation

Innovation & Collaboration

- Built AI agent integrations using an enterprise AI platform SDK (Python), creating JWT-authenticated agents for automated data querying and operational insights
- Developed custom API automation modules enabling AI agents to query platform metadata, event types, and access roles via RESTful endpoints
- Explored LLM-driven workflow automation for AI-assisted alert triage, log analysis, and intelligent pipeline monitoring
- Shared DevOps best practices through documentation, code reviews, and knowledge-sharing sessions on CI/CD, Kubernetes, Helm, and Maven workflows
- Collaborated on Snowflake historical data extraction and cross-team technical initiatives

Previous Roles — Software Engineer

2016 – 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)

Master's Degree — M.S. Computer Science 2022 – 2024

City University of New York (CUNY)

Bachelor's Degree — B.S. Computer Science 2016 – 2022**Relevant Coursework:**

- Advanced Programming Techniques (*Honors*)
- Data Structures (*Honors*)
- Cryptography and Crypto Analysis
- Discrete Structures and Applications to Computer Science
- Analytic Geometry and Calculus I & II
- Data Processing at Scale
- Applied Cryptography

TECHNICAL SKILLS

Category	Technologies
Languages	Java 17+, Python, JavaScript/Node.js, Bash/Shell, PowerShell, SQL
Big Data & Processing	Apache Spark, Apache Storm, Apache Iceberg, Nessie, Vertica, Snowflake, Apache Kudu
Search & Metadata	Elasticsearch, Painless Scripting, Elastic SQL
Cloud & Infrastructure	Microsoft Azure, Kubernetes (multi-cluster), Helm, Docker, ZooKeeper
DevOps & CI/CD	GitLab CI/CD, Maven, JaCoCo, Docker Image Pipelines, Automated Change Management
APIs & Microservices	RESTful APIs, JWT/OAuth, Spring WebFlux, Microservice Architecture
Monitoring & Observability	Splunk, Elastic Stack, Spark History Server
AI/ML & Automation	AI Agent Development, Prompt Engineering, LLM Integration, Workflow Automation
Security & Governance	SailPoint, Vulnerability Scanning, Threat Modeling, Kubernetes RBAC, SoD Compliance
Other	React Native, IBM Blockchain, Google Analytics, Jira, Git, Insomnia/Postman

LANGUAGES

- **English** — Native / Bilingual Proficiency
- **Arabic** — Native / Bilingual Proficiency
- **Spanish** — Elementary Proficiency