

Nkosi Felix

Orlando, FL / Remote | U.S. Citizen | No Sponsorship Required | Eastern Time | 404-932-5956 | njfelix11@live.com | linkedin.com/in/nkosifelix | gitlab.com/ZorillaTrace

Senior Data Platform Engineer | Governed Ingestion · Azure · Kafka · Microsoft Graph

PROFESSIONAL SUMMARY

Senior Data Platform Engineer building governed ingestion pipelines and streaming services where document quality, compliance boundaries, and auditability matter. Specializes in Python ingestion harnesses, Microsoft Graph/SharePoint synchronization, and Kafka-to-PostgreSQL consolidation, with validation, provenance, replay safety, and observability built into the intake path. Partners with engineering, Legal, Cyber, Finance, and business owners to produce reviewable data for analytics, automation, and AI systems.

CORE COMPETENCIES

Data Platforms & Ingestion: Apache Kafka, PostgreSQL, parser orchestration, delta-sync, data validation routing, provenance tracking, ETL/ELT, dimensional modeling, idempotent consumers, schema validation, data lineage, replay safety, canonical bundle design

Cloud & Integration: Microsoft Azure, Microsoft Graph API, SharePoint/OneDrive, Azure Databricks, Power BI, Power Automate, Logic Apps, REST APIs, event-driven integration

Programming & Storage: Python, Java 17, Spring Boot 3, SQL, JPA/Hibernate, Spring Cloud Stream, XMPP, OpenAPI, YAML, JSON, PowerShell

Observability, Testing & DevOps: Grafana, Prometheus, Loki, correlation-ID tracing, SLO/SLI design, GitLab CI/CD, Docker, JUnit 5, Mockito, Maven Surefire, contract testing, deployment gates, structured logging, architecture decision records (ADRs)

PROFESSIONAL EXPERIENCE

Lead Software Engineer (W2 Contract) | Universal Creative / NBCU | Orlando, FL | 02/2026 – 05/2026

- Replaced manual SharePoint document handoffs — which lacked a consistent provenance, replay, and audit mechanism — with a Python ingestion harness that detects file type, routes inputs through parsers, generates timestamped intake manifests and provenance records, and packages canonical bundles for downstream review by Legal, Finance, Estimating, and Cyber teams.
- Advanced a working OneDrive proof path into a checkpointed Microsoft Graph delta-sync that tracks file changes with eTags and content hashes, runs idempotently, lands files atomically, and records file-change deltas and run evidence so each synchronization could be reviewed and replayed within tenant-permission boundaries.
- Implemented ACCEPT/QUARANTINE/SYSTEM routing with parser provenance and ambiguity holds so unsafe, ambiguous, or duplicate files were classified and held with structured failure reasons rather than silently reaching downstream AI or analytics workflows; produced run records for each processing outcome to support review and replay.
- Validated a controlled Legal document-renaming workflow on copied SharePoint files, protecting originals while parser provenance tracked rename decisions and idempotent reruns confirmed no additional changes; documented duplicate holds, ambiguous cases, and unresolved taxonomy questions for owner review.
- Built guarded ClickUp API automation with dry-run and apply separation, preflight traces, and status updates disabled by default so task-state changes required explicit operator approval; supported comment-only access modes and a manual fallback path to keep automation bounded to audit-support use while preserving business-owner approval over task-state changes.

Software Engineer II (Data Engineer) | 4C Strategies North America | Orlando, FL | 05/2025 – 10/2025

- Designed and built an idempotent Kafka-to-PostgreSQL consolidation service to merge scatter-gather mission data — email, chat, and radio traffic — into a unified, queryable dimensional read model; used as a baseline template that gave later teams a repeatable starting point for onboarding and technical review.
- Implemented idempotent consumer logic with message-replay handling and deduplication, DTO validation at the consumer boundary, JPA/Hibernate persistence with schema-consistent write paths, and paginated REST APIs with optional filters for downstream mission-data retrieval; refactored to remove ad-hoc query patterns and reduce the read-model API surface to a stable, queryable contract.
- Built a Grafana/Prometheus/Loki observability stack with multi-layer instrumentation covering JVM health, ETL latency, schema-violation rates, and SLO compliance; correlation-ID tracing propagated across pipeline stages cut distributed debugging cycles by approximately 30%, enabling faster root-cause isolation and reducing dependency on cross-team escalation during pipeline investigation.
- Authored contract and unit tests for Kafka consumer and normalizer services using JUnit 5, Mockito, and Spring Cloud Stream test binder — including JSON serialization validation — and integrated the suite into GitLab CI deployment gates, catching schema-drift regressions and serialization failures during CI before deployment review.

- Built a deterministic XMPP acceptance harness with XML-driven sender/receiver, in-band registration support, and templated conversation flows for repeatable local air-gapped validation without external service dependencies; delivered runbooks, OpenAPI notes, security profiles, and ADRs as reference artifacts for air-gapped or compliance-sensitive mission environments.

Senior Cloud Solutions Architect | Accenture, Microsoft Advanced Azure Engagement | Atlanta, GA | 05/2024 – 12/2024

- Supported Microsoft Advanced Azure customer-success programs by tracking cloud adoption, usage trends, account priorities, and engagement risks through Power BI and operational reporting.

- Prepared and participated in weekly and monthly customer business reviews, translating Azure services and early Copilot capabilities into clear use cases, recommendations, risks, and next steps.

- Partnered with Microsoft, Accenture, and customer stakeholders to surface adoption blockers, coordinate follow-ups, and route deeper technical issues to the appropriate product or engineering teams.

- Created reusable account-review materials, executive updates, and engagement documentation that kept customer priorities, technical questions, and cloud-adoption actions visible across stakeholder groups.

Customer Success Manager | Kearney | Atlanta, GA | 10/2021 – 12/2022

- Formalized customer-success playbooks for AI/ML startup engagements, improving consistency across onboarding, progress reviews, resource handoffs, and executive business reviews.

- Built Voice-of-Customer analytics in Excel and Power BI to surface adoption barriers, satisfaction drivers, churn signals, and recurring engagement risks.

- Consolidated Salesforce/CRM reporting into account-health, adoption-blocker, escalation-history, retention, and business-development views for Sales, Product, and Data Analytics teams.

- Built a CRM-backed advocacy and reference program that organized customer proof points, case-study material, and success-story evidence for retention and new-business conversations.

Senior Data Analyst / Data Engineer (W2 Contract) | Flexport | Atlanta, GA | 10/2020 – 04/2021

- Analyzed international freight routing patterns, carrier performance, and SLA attainment data using Python, SQL, and statistical methods to surface systemic delays, capacity constraints, and route optimization opportunities across the supply chain.

- Built Power BI and Excel operational dashboards providing operational visibility across 100+ concurrent shipments, enabling dispatch teams, logistics managers, and exception handlers to make data-driven decisions on live freight status and delay signals.

- Identified freight cost patterns, rate discrepancies, and carrier-switching opportunities through data warehouse analysis; supported \$50K+ in quarterly savings through analytics-backed cost-reduction and carrier-performance recommendations.

Senior Operations Analyst | Hapag-Lloyd | Atlanta, GA | 02/2019 – 11/2019

- Managed key international freight accounts and escalations for ocean export operations, coordinating with carriers, customs brokers, and freight forwarders on AES filings, exception resolution, and on-time service delivery across export lanes.

- Built Power BI dashboards and Salesforce/CRM reports tracking export volume, port performance, and exception rates by customer and lane; advised account teams on incoterms, routing options, service constraints, and cost considerations.

KEY DELIVERABLES & PLATFORM PROJECTS

acceptance-harness — Public reference implementation demonstrating deterministic document-ingestion patterns: embedded in-memory test database, canonical test fixtures with deterministic input/output pairs, ACCEPT/QUARANTINE/SYSTEM routing validation with structured failure reasons, duplicate-detection and ambiguity-hold logic, and full GitLab CI configuration. Reviewers can inspect the Spring Cloud Stream test-binder setup, fixture design, replay-safe manifest structure, and acceptance criteria documentation without needing a deployed infrastructure stack. Running the Maven wrapper executes the fixture suite and produces a JUnit report confirming deterministic ACCEPT, QUARANTINE, SYSTEM-failure, and duplicate routing states against `ACCEPTANCE_CRITERIA.md` without a deployed stack.

airgap-messaging-ingest — Public reference implementation demonstrating idempotent XMPP-to-Kafka-to-PostgreSQL ingestion: schema validation at the consumer boundary, idempotent message-processing with deduplication and replay handling, and correlation lineage for distributed trace queries across pipeline stages. Includes Grafana dashboard configurations, Prometheus metric definitions, Loki log-aggregation setup, Docker Compose for local deployment, and operations runbooks covering startup, failure-recovery, and local air-gapped testing procedures without external service dependencies. Running `docker compose up` followed by Maven `verify` submits sample payloads, triggers deduplication and schema-validation paths, and produces JUnit and JaCoCo coverage evidence aligned with the Grafana and Loki configuration in the repository.

Both repositories are personal reference implementations demonstrating production-oriented engineering patterns. No confidential employer code is included.

CERTIFICATIONS

- Confluent Data Streaming Engineer — Confluent | Sep 2025 – Sep 2027
- Azure Databricks Platform Architect — Databricks | Nov 2025 – Nov 2026
- Microsoft Certified: Azure Data Fundamentals | Nov 2025
- Microsoft Certified: Azure AI Fundamentals | Oct 2025