Michael Humphrey

Raleigh, NC 27606 | michael.humphrey.biz@gmail.com | (919) 446-5346 |
linkedin.com/in/michael-humphrey-a95343100

PROFESSIONAL SUMMARY

Cloud-native software engineer and technical leader with 9+ years of experience building the platform-level infrastructure behind enterprise SaaS products. Consistently operates at the intersection of deep AWS architecture and measurable business outcomes — reducing commercial bank onboarding from weeks to hours, engineering a supply chain event system designed for ~100 million events per year, and building automated pipelines that empower non-engineers to ship to production without any engineering involvement. Brings a rare cross-domain background spanning fintech, food safety, and computational biology, including co-invention of a US patent in machine learning-based molecular classification. Targeting Senior, Staff, and Principal Engineer roles at companies building serious infrastructure in fintech, healthcare technology, and enterprise SaaS.

SKILLS

Languages & Frameworks: TypeScript, Node.js, Go, Python, React, Java, C#, SQL, HTML, CSS

AWS: API Gateway, Lambda, DynamoDB, Step Functions, SQS, SNS, EventBridge, S3, KMS, IAM, Secrets Manager, ECS (Fargate), QuickSight, CodePipeline, CloudWatch, CDK

Salesforce Platform: Apex, Queueables, Platform Events, Flows, Managed Packages, sObjects

Data & Streaming: Confluent Kafka, Databricks, Snowflake, MongoDB, PostgreSQL (Aurora), Redis

Infrastructure & DevOps: Docker, GitHub Actions, Azure DevOps, Azure Container Registry

AI/ML: scikit-learn, pandas, Qwen2-VL, RunPod, FFMPEG, OpenCV

Observability: Datadog

Other: Kong Konnect, Informatica IICS, Auth0, JWT, go-swagger

Tools: Git, Jira, Confluence, VS Code, IntelliJ, Cursor, Claude Code

PROFESSIONAL EXPERIENCE

nCino | Senior Software Engineer | December 2021 – Present

- Served as one of three primary architects and engineering lead across two cross-functional teams (6 engineers total; 1 senior and 2 mid-level per team) to design and deliver nCino's International Commercial Onboarding & Loan Origination platform — reducing commercial bank customer onboarding from weeks to hours. Set technical direction, managed delivery timelines, interfaced with product and executive stakeholders, authored work breakdown structures, and remained deeply hands-on in code throughout. Currently live in production with two UK financial institutions including Yorkshire Building Society, with active international expansion underway via a Dun & Bradstreet integration. Incorporated FullCircl's KYB and eKYC solutions (a nCino acquisition) to meet international regulatory compliance requirements.

  Technologies: AWS API Gateway, Lambda, DynamoDB, PostgreSQL (Aurora), Step Functions, Secrets Manager, IAM, React, TypeScript, Salesforce Platform

- Contributed as API architect on nCino's greenfield Omnichannel loan origination platform — a unified solution consolidating previously fragmented products across consumer, small business, commercial, and mortgage banking into a single cohesive system. Designed and implemented RESTful API endpoints for headless deposit account opening for small business banking clients. The platform now serves approximately 10% of nCino's client base.

  Technologies: Node.js, TypeScript, AWS CDK, IAM, KMS, API Gateway, Lambda, DynamoDB, Salesforce APIs, Ruby on Rails, React

- As team lead working alongside one engineer and a team of AWS Solutions Architects, architected nCino's Benchmarking & Analytics platform — replacing Domo — by building a multi-tenant QuickSight embedding infrastructure that surfaces loan performance, portfolio risk, and origination lifecycle dashboards directly within nCino's Salesforce Managed Package. Engineered row-level security and tenant isolation across pooled Databricks data sources, ensuring complete data segregation

between financial institutions sharing a single dashboard asset.

Technologies: Node.js, TypeScript, AWS CDK, IAM, KMS, API Gateway, Lambda, DynamoDB, QuickSight, Databricks, Salesforce

- As sole engineer, designed and implemented a fully automated GitOps-style CI/CD pipeline enabling business analysts to promote QuickSight dashboards from development to production with zero engineering involvement — reducing deployment from a multi-hour manual process to approximately 5 minutes. The pipeline automatically validates Row-Level Security guardrails at promotion time, rejects non-compliant dashboards with owner notifications, versions all dashboard assets in source control, and publishes availability records to a database consumed by the Salesforce Managed Package.

  Technologies: Node.js, TypeScript, AWS CDK, EventBridge, Lambda, QuickSight, S3, CodePipeline, GitHub Octokit, GitHub Actions, GitHub Apps

- Architected the core REST API suite for MTRNS, nCino's proprietary banking core integration mapping framework, enabling engineers to centrally create, version, and reuse XPath/XQuery data transformation rules across all financial institution tenants. Replaced fully bespoke per-client implementations maintained by a dedicated professional services team with a reusable, multi-tenant mapping repository that standardizes core banking integrations (e.g., Jack Henry jXchange) while preserving per-FI configurability.

  Technologies: Node.js, TypeScript, AWS CDK, IAM, KMS, API Gateway, Lambda, DynamoDB, S3

- Contributed to automating deployment of nCino's banking core integration framework across all tenant environments, transitioning from unmaintainable per-financial-institution implementations delivered manually by a professional services team to a centralized, repeatable automated process. Integrated the MTRNS mapping framework with Informatica IICS to enable a single unified set of integration workflows to serve all tenants securely via GitHub Actions.

  Technologies: GitHub Actions, Informatica IICS

- Served as lead engineer on a team of 3 building nCino's native Salesforce API orchestration framework — an Airflow-inspired workflow engine built entirely on Salesforce platform technologies. Engineered to overcome fundamental Salesforce governor limit constraints, including the platform prohibition on performing DML operations across callout transaction boundaries, the framework uses Queueable Apex, custom Business Activity sObjects for job lifecycle tracking, Platform Events, and Flows to orchestrate complex multi-step API workflows across transaction boundaries within a single unified execution model. A general-purpose framework now powering all nCino integrations, including Alloy Fraud Detection and KYC compliance workflows.

  Technologies: Salesforce Apex, Queueables, Platform Events, Flows, sObjects

FoodLogiQ | Software Engineer | June 2019 – December 2021

- As sole engineer, built a CLI tool and custom YAML parser — extending the open-source go-swagger library — to auto-generate Swagger documentation for hundreds of RESTful APIs at build time, with configurable public/private endpoint visibility to prevent internal APIs from being exposed in the developer portal. This greenfield initiative enabled developer partners, including Whole Foods, to self-serve integrations for the first time and directly supported the launch of FoodLogiQ's formal partner integrations program.

  Technologies: Go, go-swagger

- Drove a one-month initiative that elevated FoodLogiQ's integration test coverage from effectively 0% to over 90% across hundreds of API endpoints, leveraging swagger-test-templates to auto-generate test scaffolding from the Swagger documentation built in the prior initiative. Caught real production bugs during the process, significantly accelerated deployment confidence across the engineering organization, and established a testing pattern that became a permanent fixture in the team's API development workflow.

  Technologies: Node.js, mocha, chai, swagger-test-templates

- As sole engineer, designed and implemented a fault-tolerant batch product upload system enabling clients to sync millions of products into FoodLogiQ's Manage & Monitor platform — a capability that did not previously exist. Individual product API calls averaged ~1.5 seconds, making large catalog syncs effectively impossible without parallelized infrastructure. Engineered SNS/SQS fanout for parallel processing with per-record retry logic, dead-letter queues, and partial batch recovery, ensuring individual product failures never abort a full upload. Enabled Whole Foods to perform full catalog syncs from Glye into FoodLogiQ for the first time.

  Technologies: Go, AWS Lambda, SNS, SQS, DynamoDB, MongoDB

- As sole engineer, designed and implemented FoodLogiQ's webhook notification system, eliminating the need for customer polling — some clients polling as frequently as every 30 seconds — by delivering near-real-time supply chain event notifications across the entire customer base. Built with at-least-once delivery guarantees, automatic retry on failure, and dead-letter queue handling for all platform event types. Adopted by all customers upon release.

  Technologies: Go, MongoDB, Redis, AWS Lambda, SNS, SQS, S3, Athena

- Implemented automated role mapping for FoodLogiQ's Single Sign-On feature, eliminating the manual process of administrators assigning platform roles to every newly onboarded user. Roles and groups from users' identity providers — including Azure AD and SAML/OIDC providers via Auth0 — are now automatically translated into FoodLogiQ platform roles at onboarding time, meeting enterprise access control requirements across all customer organizations.

  Technologies: Go, AngularJS, MongoDB, JWT, Auth0, Auth0 Actions, JavaScript

- Served as architect and lead engineer on a team of 3 in a 4-month ground-up rebuild of FoodLogiQ's supply-chain traceability platform, replacing a non-scalable in-memory Java async queue with a production-grade event streaming architecture. The new system ingests MongoDB change streams into Confluent Kafka, processes farm-to-shelf shipping and receiving events (farm → distribution center → grocery store) through Kafka consumers, and loads transformed data into Snowflake for analytics. Designed to support ~100 million supply chain events per year. Live in production today.

  Technologies: Go, Confluent Kafka, Cucumber/GoDog, MongoDB, Snowflake, SnowSQL

Metabolon | Software Engineer | January 2017 – June 2019

- Sole author of a Python rules-based classification module serving as the first-pass stage in Metabolon's automated metabolomics analysis pipeline. Applies decision-tree classification to HPLC-MS/MS retention time data to identify 30–40 small molecules per study with near-100% accuracy, routing clear-cut identifications away from the computationally expensive multilinear ridge regression stage and reducing analysis time from up to a week per study to near-fully-automated. Contributed to the patented machine learning workflow for automated molecular identification (US11387087B2).

  Technologies: Python, scikit-learn, pandas

- Extended Metabolon's legacy data curation application with an Expected Retention Time (RT) visualization derived from the multilinear ridge regression model at the core of the patented analytical workflow. The feature displays expected, actual, and reference RTs side-by-side, transforming scientists' manual curation experience and significantly accelerating the study review process.

  Technologies: C#, WinForms, Oracle PL/SQL

- Built CI/CD pipelines for the Data Science team's AI/ML projects using Azure DevOps, replacing a manual script-execution workflow with reproducible, versioned deployments across all data science Python packages.

  Technologies: Azure DevOps, Azure Container Registry, Docker, Twine

CONTRACT WORK

Software Engineer, Contractor | September 2024 – January 2025

- Sole engineer on an automated video and image content labeling system for a confidential client, integrating a fine-tuned multimodal AI model (Qwen2-VL) with a FastAPI service to enable customers to submit media files for automated semantic tagging. The system labels a 10-minute video in 1–2 minutes, generating searchable metadata (scene descriptions, subjects, actions) to power content search and discovery. Deployed FastAPI on AWS ECS Fargate with GPU inference hosted on RunPod, and asynchronous job processing via SQS. Shipped to production and remains a core feature of the client's studio product.

  Technologies: Node.js, TypeScript, AWS ECS (Fargate), AWS SQS, FastAPI, Python, Kong Konnect, RunPod, Qwen2-VL, FFMPEG, OpenCV

EDUCATION

North Carolina State University

B.S. Biochemistry | 2015 | Summa Cum Laude

PATENT

Method for analyzing small molecule components of a complex mixture, and associated apparatus and computer program product

US Patent No. US11387087B2 | Co-inventor

Awarded for a novel machine learning approach to automated metabolite classification in complex biological samples, combining rules-based molecular detection with multilinear ridge regression for retention time prediction.