

O. FELIX AMORUWA

famoruwa@berkeley.edu | 909-731-9011 | felixamoruwa.info

Staff Product Manager with 12+ years delivering data infrastructure and platform products at enterprise scale — from owning Splunk's search and data pipeline services to scaling Intuit's data platform to 675M+ engagements across QuickBooks, TurboTax, and Mailchimp. Proven record translating ambiguous technical customer feedback into bold product vision, leading cross-functional teams of engineers and designers, and executing 0-to-1 data products end-to-end. Deep fluency with data/operations personas, API integrations, and cloud data infrastructure; NeurIPS published researcher.

PROFESSIONAL EXPERIENCE

INTUIT | Mountain View, CA

Staff Product Manager — Data Platform & Developer Infrastructure 06/21 – 09/24

- Scaled Intuit's core data platform (ICE) to 675M+ engagements in FY23 across QuickBooks, TurboTax, Mint, Mailchimp, and Credit Karma — 275% YoY growth — by owning product roadmap, cross-functional execution, and platform strategy across 20+ mobile apps and 30+ product SKUs.
- Led Mailchimp GCP-to-AWS data infrastructure migration (MSaaS), delivering Golang service template, MySQL persistence integration, and updated developer documentation to meet production deadline — directly analogous to cross-cloud data pipeline work.
- Built Asterias, a declarative asset lifecycle management platform with GraphQL API, enabling data-driven visibility into service dependencies across the enterprise.
- Worked closely with telemetry and usage data (SQL, BigQuery) to distill developer pain points into prioritized product direction; conducted enterprise-wide Service Language Assessment across 9 languages, presenting findings to CTO.
- Delivered ICE Self-Service platform (DevPortal, GitOps config, ICE Playground), reducing developer onboarding from 2–3 weeks to minutes in pre-prod and <24 hours for production, while mitigating \$1M+ in projected opex growth.
- Scaled platform throughput from 6K to 50K TPS via rSocket migration supporting ~1.5M concurrent connections with sub-25ms TP99; implemented ICE Presence in async chat, generating \$480K/month in additional invoicing.
- Initiated MSaaS Drift Detection and Resolution program: authored Java JAR library to scan Git repos for configuration drift and built remediation roadmap using OpenRewrite. DeveloperWeek 2022 speaker.

SPLUNK INC. | San Francisco, CA

Senior Product Manager — Data Pipeline & Search Infrastructure 01/19 – 01/21

- Owned Search Service (Go microservices), Search Catalog (PostgreSQL metadata service), and Splunk Processing Language (SPL/SPL2) — core data pipeline and query infrastructure products for Splunk Cloud Services (SCS); built product roadmaps, PRDs, and acceptance criteria.
- Led query performance optimization initiative for enterprise beta customer (Assurance), building mirrored topology for benchmark testing and achieving up to 10x performance improvements in SCS search — turning ambiguous customer pain into measurable product wins.
- Delivered Scheduler Service end-to-end in ~4 months (June–Oct 2019), enabling scheduled data pipeline capabilities for first-party applications; demoed at Splunk .conf19.
- Designed repeatable RICE-based prioritization framework for 3 microservice backlogs, balancing Fortune 500 customer requirements, internal partners, and third-party developer needs. Splunk .conf18 Speaker.

STREAMIO AI | Oakland, CA

Founder & CEO — streamio.ai (macOS, Linux, iOS & Web) 09/24 – Present

- Led 0-to-1 product strategy, customer discovery, and go-to-market execution for a multi-platform AI product integrating real-time data from Redfin/Zillow APIs, Claude (MCP SDK), and ElevenLabs — demonstrating end-to-end ownership across ambiguous, fast-moving product territory.
- Implemented OpenClaw multi-agent orchestration framework with gateway protocol and subagent delegation, enabling coordinated data workflows across real estate, insurance, health/dental, and financial markets — a reverse-data-flow architecture routing enriched outputs back to operational contexts.
- Built end-to-end auth and payments pipeline (Kinde OAuth 2.0, Stripe subscriptions) and shipped cross-platform desktop builds (macOS DMG, Linux .deb) with native OS integrations.

FINTELLECT AI | Oakland, CA

Founder & CEO — fintellectai.io (iOS/Android/Web) 09/24 – Present

- Architected RAG retrieval pipeline with ChromaDB vector store and multi-provider LLM orchestration (Claude, GPT-4, Gemini) with fallback routing — a multi-source data integration layer delivering context-aware insights to retail investor operations personas.
- Led customer discovery, iterated product based on trader feedback, and executed go-to-market through App Store launch.

KAISER PERMANENTE | Pleasanton, CA

SOA Technical Product Manager 08/12 – 12/18

- Led development and enterprise rollout of Splunk Logging-as-a-Service (1.7 TB daily volume, 200+ internal enterprise customers) — a large-scale data ingestion and observability platform serving data/operations personas across the enterprise.
- Built caching capability using Redis and XC10 to address scalability, fault tolerance, and data redundancy; led demand forecasting and capacity planning for SOA product suite across multiple datacenters.

IBM | Pittsburgh, PA / Foster City, CA / RTP, NC

Software Engineer — Business Intelligence Products 07/05 – 03/10

- Resolved multiple \$1–5M high-severity customer escalations for enterprise BI and data infrastructure products; led cross-functional root cause analysis, improving case resolution time by 20% across the Support Organization.

BANK OF AMERICA MERRILL LYNCH | Charlotte, NC

Tech MBA Summer Associate — Global Corporate & Investment Banking T&O 06/11 – 08/11

- Developed enhanced estimation model using Monte Carlo simulation to optimize DMAIC phase distribution across a \$494M portfolio in Global Technology & Operations.

AI/ML RESEARCH & PROJECTS

aeval — AI Model Evaluation Platform (2025–2026)

- Built local-first evaluation platform with 5 core eval types, adversarial safety testing, and CI/CD regression detection. Stack: FastAPI orchestrator, TimescaleDB, Redis job queue, Next.js dashboard, Ollama — demonstrating end-to-end data pipeline design from ingestion to structured reporting.

RL Workbench — Post-Training RL Platform (2026)

- Built 3-phase RL post-training workbench benchmarking GRPO/DPO across TRL, VeRL, OpenRLHF, and NeMo RL; implemented 12 RL algorithms with standardized throughput/memory/convergence benchmarking across frameworks.

BRAIN — Protein Structure Prediction ML Platform (UC Berkeley 2004; rewritten 2026)

- NeurIPS 2014 accepted paper; production ML platform in PyTorch with 5 neural architectures, MLflow experiment tracking, Optuna HPO, and FastAPI serving — 823 automated tests, Docker orchestration (6 containers).

EDUCATION

CARNEGIE MELLON UNIVERSITY, SILICON VALLEY

Master of Science — Software Management Mountain View, CA (08/16)

CARNEGIE MELLON, TEPPER SCHOOL OF BUSINESS

Master of Business Administration — Concentrations: Finance, Quantitative Analysis, Economics Pittsburgh, PA (03/12)

UC BERKELEY, COLLEGE OF ENGINEERING & HAAS SCHOOL OF BUSINESS

Bachelor of Science — Computational Engineering Science Berkeley, CA (12/05)

Certificate — Corporate Finance for Financial Engineering

TEACHING EXPERIENCE

DE ANZA COLLEGE | Cupertino, CA

Adjunct Faculty — Computer Information Systems Department 04/18 – Present

- Java Programming (CIS 35A/B, 36A/B); Introduction to PL/SQL (CIS 64C/G); Introduction to Data Analytics (CIS 44F); Fundamentals of Large Scale Cloud Computing — AWS, GCP (CIS 64E); Ethical Hacking (CIS 102); Digital Forensics (CIS 104); Managing Cloud Projects (CIS 95F).

ADDITIONAL INFORMATION

- FAA Commercial Drone License; FAA Private Pilot License (intended CFI)
- O'Reilly & OER Publications: Data Analytics and Managing Cloud Projects (forthcoming Summer 2026)
- First Tee Assistant Golf Coach — San Francisco Chapter (2012–Present), Pittsburgh Chapter (2011–2012)
- Interests: Classical & Jazz Saxophone (Alto/Tenor/Baritone), Triathlon, Golf, Intramural Basketball & Soccer