

O. FELIX AMORUWA

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

Technical Product Leader with 15+ years building infrastructure platforms and orchestration systems at scale — scaling distributed services to 675M+ engagements and 50K TPS at Intuit, owning search orchestration microservices at Splunk, and architecting agentic AI workflow frameworks in production today. Proven track record defining SLOs, driving failure attribution to automated resolution, and translating complex system state into operator UX that engineers can act on under pressure. NeurIPS published researcher; hands-on GPU infrastructure experience benchmarking RL frameworks across CUDA/MPS with Docker passthrough.

PROFESSIONAL EXPERIENCE

INTUIT | Mountain View, CA

Staff Product Manager — Platform Infrastructure & Workflow Orchestration 06/21 – 09/24

- Scaled distributed platform infrastructure to 675M+ engagements in FY23 across QuickBooks, TurboTax, Mint, Mailchimp, and Credit Karma; drove throughput from 6K to 50K TPS via rSocket migration supporting ~1.5M concurrent connections at sub-25ms TP99 — owning fleet availability SLOs end-to-end.
- Delivered ICE Self-Service platform (DevPortal, GitOps config, ICE Playground), reducing operator onboarding from 2–3 weeks to minutes in pre-prod and <24 hours for production — directly mitigating \$1M+ in projected opex growth through workflow automation.
- Initiated MSaaS Drift Detection and Resolution program: authored Java JAR library to scan Git repos for configuration drift, built remediation roadmap using OpenRewrite, and partnered with Design on operator-facing DevPortal UI — establishing audit trails and automated repair workflows at enterprise scale.
- Built Asterias, a declarative asset lifecycle management platform with GraphQL API, enabling structured failure attribution and lifecycle tracking across 20+ mobile apps and 30+ product SKUs using SQL and BigQuery telemetry.
- Led Mailchimp GCP-to-AWS migration for MSaaS under production deadline, delivering Golang service template, MySQL persistence integration, and updated operator documentation — managing blast radius across multi-cloud topology.
- Conducted enterprise-wide Service Language Assessment across 9 languages, analyzing usage data and developer feedback to inform strategic investment decisions presented to CTO; implemented ICE Presence in async chat generating \$480K/month in additional invoicing.

SPLUNK INC. | San Francisco, CA

Senior Product Manager — Search Orchestration & Distributed Systems 01/19 – 01/21

- Owned Search Service (Go microservices), Search Catalog (PostgreSQL metadata service), and Splunk Processing Language (SPL/SPL2) — defined product roadmaps, SLOs, PRDs, and acceptance criteria for Splunk Cloud Services (SCS) spanning distributed orchestration infrastructure.
- Led query performance optimization initiative for enterprise beta customer, building mirrored topology for benchmark testing and achieving up to 10x performance improvements — demonstrating reliability engineering instincts and blast radius management in production search infrastructure.
- Delivered Scheduler Service end-to-end in ~4 months, enabling scheduled workflow execution for first-party applications; designed RICE-based prioritization framework balancing Fortune 500 SLA requirements across 3 microservice backlogs. Splunk .conf18/.conf19 Speaker.

KAISER PERMANENTE | Pleasanton, CA

SOA Technical Product Manager — Infrastructure Reliability 08/12 – 12/18

- Led enterprise rollout of Splunk Logging-as-a-Service (1.7 TB daily volume, 200+ internal enterprise customers) and ITSI Application Monitoring-as-a-Service — owning availability SLOs and operator experience for mission-critical healthcare infrastructure.
- Architected caching capability using Redis and XC10 across the enterprise to address fault tolerance, data redundancy, and scalability challenges; led datacenter capacity planning across multiple sites to maintain fleet availability under demand growth.

STREAMIO AI | Oakland, CA

Founder & CEO — AI Orchestration Platform (macOS, Linux, iOS & Web) 09/24 – Present

- Architected OpenClaw multi-agent orchestration framework with gateway protocol, subagent delegation, profile management, and session switching — building agentic AI workflow software with human-in-the-loop intervention points and coordinated agent execution across multiple domains.
- Shipped MCP SDK integration exposing screen capture tools to AI coding assistants; built end-to-end auth and payments pipeline (Kinde OAuth 2.0, Stripe, Electron SafeStorage) with real-time HLS livestreaming pipeline supporting up to 9 concurrent sources at 30fps.

FINTELLECT AI | Oakland, CA

Founder & CEO — AI Financial Platform (iOS/Android/Web) 09/24 – Present

- Architected RAG retrieval pipeline with ChromaDB vector store, multi-provider LLM orchestration (Claude, GPT-4, Gemini) with fallback routing and structured output validation — building reliability-first agentic AI workflows with automated confidence thresholds and token budget optimization.

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 software; led cross-functional root cause analysis and failure attribution workflows, improving case resolution time 20% across the Support Organization — early foundation in operational consequence and blast radius management.

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

RL Workbench — GPU Infrastructure & Post-Training RL Platform (2026)

- Built multi-framework RL benchmarking platform with GPU Docker passthrough, enabling head-to-head comparison of TRL, VeRL, OpenRLHF, and NeMo RL — standardized throughput, memory, and convergence benchmarking across CUDA and Apple Silicon (MPS) with live SSE metric streaming; directly analogous to AI factory fleet performance instrumentation.

- Implemented 12 RL algorithms (PPO, GRPO, DAPO, REINFORCE, REINFORCE++, RLOO, DPO, SimPO, IPO, KTO, ORPO, SPPO) with algorithm-specific metric profiles, cross-tab workflow lineage tracking, and 3-phase pipeline: Reward Lab → Playground → Arena.

aeval — AI Model Evaluation Platform with Automated Safety Gates (2025–2026)

- Built local-first evaluation platform with CI/CD integration, regression detection, and automated safety gates — implementing automation confidence thresholds and human-in-the-loop intervention criteria analogous to break-fix automation guardrails. Stack: FastAPI orchestrator, TimescaleDB, Redis job queue, Next.js dashboard, Ollama.
- Statistical rigor: bootstrap confidence intervals, Welch's t-test, Cohen's d effect size, saturation detection across 5 eval types (factuality, reasoning, instruction-following, safety, code generation).

AutoEval — Automated Visual Evaluation for Robot Model Training (2025)

- Built zero-integration automated evaluation pipeline reducing evaluation cycles from 72 hours to ~4 minutes; multimodal AI (Claude/GPT-4V) generates structured PASS/FAIL reports with confidence scores and edge-case flagging — a detection-to-resolution pipeline directly analogous to break-fix automation with audit trails.

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

- Production ML platform in PyTorch with 5 neural architectures, MLflow experiment tracking, Optuna HPO, FastAPI serving — 823 automated tests, Docker orchestration (6 containers); scales from 413 to 8B parameters. NeurIPS 2014 accepted paper.

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

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

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