

O. FELIX AMORUWA

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

Enterprise infrastructure PM with 12+ years turning bespoke, deal-by-deal platform gaps into supported, priced capabilities at scale — from Intuit's ICE platform (675M+ engagements, 50K TPS, sub-25ms TP99) to Splunk Cloud's Go microservices and Kaiser Permanente's healthcare-regulated Logging-as-a-Service (1.7 TB/day, 200+ enterprise customers). Shipped identity/access controls, usage-based billing, and multi-cloud deployments; credible with security, compliance, and procurement buyers and with the engineers who build it. Hands-on ML infrastructure builder: RL post-training workbench benchmarking TRL, VeRL, OpenRLHF, and NeMo RL. NeurIPS published.

PROFESSIONAL EXPERIENCE

INTUIT | Mountain View, CA

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

- Achieved 275% YoY growth in ICE engagements, scaling to 675M+ in FY23 across QuickBooks, TurboTax, Mint, Mailchimp, and Credit Karma; scaled throughput from 6K to 50K TPS via rSocket migration supporting ~1.5M concurrent connections with sub-25ms TP99 — turning a bespoke per-team integration into a platform-wide, supported capability.
- 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 — converting deal-by-deal exceptions into a priced, sellable product.
- Led Mailchimp GCP-to-AWS migration for MSaaS, delivering Golang service template, MySQL persistence integration, and updated DevPortal documentation — demonstrating multi-cloud deployment ownership across regulated enterprise environments.
- Initiated MSaaS Drift Detection and Resolution program: authored Java JAR library to scan Git repos for configuration drift, partnered with Design on DevPortal UI, and built remediation roadmap using OpenRewrite — establishing compliance posture and governance at platform scale.
- Conducted enterprise-wide Service Language Assessment across 9 languages (Java, Python, Kotlin, Go, TypeScript, Scala, PHP, C++, Groovy), 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; deployed Background-to-Foreground Messaging on iOS/Android with <100ms latency. DeveloperWeek 2022 speaker.

SPLUNK INC. | San Francisco, CA

Senior Product Manager — Search Orchestration & Cloud Services 01/19 – 01/21

- Owned Search Service (Go microservices), Search Catalog (PostgreSQL metadata service), and Splunk Processing Language (SPL/SPL2) for Splunk Cloud Services — built roadmaps, PRDs, and acceptance criteria credible with both Fortune 500 enterprise buyers and the engineers building the platform.
- Led query performance optimization for Fortune 500 beta customer (Assurance), building mirrored Enterprise topology for benchmark testing and achieving up to 10x performance improvements in SCS search — turning a deal-blocking requirement into a platform-wide capability.
- Delivered Scheduler Service end-to-end in ~4 months (June–Oct 2019), enabling scheduled search for first-party applications; demoed at Splunk .conf19.
- Designed repeatable RICE-based prioritization framework for 3 microservice backlogs, balancing internal partner, third-party developer, and Fortune 500 customer requirements. Splunk .conf18 Speaker.

KAISER PERMANENTE | Pleasanton, CA

SOA Technical Product Manager — Enterprise Platform Infrastructure 08/12 – 12/18

- Led development and enterprise rollout of Splunk Logging-as-a-Service in a healthcare-regulated environment (1.7 TB daily volume, 200+ internal enterprise customers) and ITS! Application Monitoring-as-a-Service — delivering compliance posture and observability as a supported, priced platform capability.
- Built enterprise caching capability using Redis and XC10 to address scalability, fault tolerance, and data redundancy — infrastructure product depth in a regulated industry.
- Led demand forecasting and capacity planning for SOA product suite across multiple datacenters, developing IT capacity plans that balanced security, compliance, and operational rigor requirements.

STREAMIO AI | Oakland, CA

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

- Shipped identity/access controls and usage-based billing from scratch: Kinde OAuth 2.0 authentication, Stripe subscriptions with tiered plans, and Electron SafeStorage for secure credential management — demonstrating end-to-end enterprise readiness surface thinking in a 0-to-1 product.
- Built cross-platform deployment pipeline (macOS DMG with code signing/notarization, Linux .deb, iOS) with native ScreenCaptureKit integration via Swift and MCP server exposing screen capture tools to AI coding assistants — multi-platform deployment options ownership.
- Implemented OpenClaw multi-agent orchestration framework with gateway protocol, subagent delegation, and session management across real estate, insurance, health/dental, and financial markets verticals.

FINTELLECT AI | Oakland, CA

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

- Architected multi-provider LLM orchestration (Claude, GPT-4, Gemini) with fallback routing, RAG retrieval pipeline (ChromaDB), structured output validation, and token budget optimization — production model serving infrastructure for a regulated financial services use case.
- Built AI-powered charting and macroeconomic analysis tools for retail investors; executed go-to-market through App Store launch with customer discovery driving iterative product refinement.

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 software; led cross-functional root cause analysis and customer-facing resolution, improving case resolution time by 20% — credibility with Fortune 500 enterprise buyers built from the engineering side.

BANK OF AMERICA MERRILL LYNCH | Charlotte, NC

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

- Developed Monte Carlo simulation model to optimize DMAIC phase distribution across a \$494M portfolio in Global Technology & Operations — quantitative rigor in a regulated financial services environment.

AI/ML RESEARCH & PROJECTS

RL Workbench — Post-Training RL Platform (2026)

- Built 3-phase RL post-training workbench benchmarking TRL, VeRL, OpenRLHF, and NeMo RL (directly mapping to Baseten Loops' frontier RL workloads) — Reward Lab for A/B testing reward functions, Playground for live GRPO/DPO training with SSE metric streaming on MPS/CUDA, and Arena for head-to-head framework benchmarking with GPU Docker passthrough.
- Implemented 12 RL algorithms (PPO, GRPO, DAPO, REINFORCE, REINFORCE++, RLOO, DPO, SimPO, IPO, KTO, ORPO, SPPO) with standardized throughput/memory/convergence benchmarking across frameworks — hands-on ML infrastructure builder credibility.

aeval — AI Model Evaluation Platform (2025–2026)

- Built local-first evaluation platform with 5 core eval types, adversarial safety testing with refusal detection, and CI/CD integration with automated safety gates — compliance posture and governance thinking applied to ML infrastructure. Stack: FastAPI, TimescaleDB, Redis, Next.js, Ollama.
- Statistical rigor: bootstrap confidence intervals, Welch's t-test, Cohen's d effect size, saturation detection, and regression detection — rigorous researcher who lets data drive the call.

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

- Production ML platform in PyTorch with 5 neural architectures (feedforward, GRU, Transformer, ESM-2, multi-task), MLflow experiment tracking, Optuna HPO, and FastAPI serving — 823 automated tests, Docker orchestration (6 containers), 413 to 8B parameters (19M-fold scale increase).
- NeurIPS 2014 accepted paper on artificial neural networks for protein secondary structure prediction.

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

- Repurposed screen capture + HLS streaming + multimodal AI pipeline to score robot model outputs against natural-language rubrics — reducing evaluation cycles from 72 hours to ~4 minutes; zero-integration architecture captures from any visualization tool via screen capture.

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