

O. FELIX AMORUWA

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

Platform-focused Technical Product Manager with 12+ years owning platform-as-a-product at enterprise scale — from scaling Intuit's developer platform to 675M+ engagements and 50K TPS, to building production agentic integration infrastructure with MCP servers and multi-agent orchestration today. Deep fluency in distributed microservices, runtime execution, observability, and hybrid/on-premises deployment models. Proven track record shipping developer SDKs, self-service DevPortals, and SLO-driven platform investments at Global 500 companies. NeurIPS published; Carnegie Mellon MS & MBA.

PROFESSIONAL EXPERIENCE

INTUIT | Mountain View, CA

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

- Achieved 275% YoY growth in platform engagements, scaling to 675M+ in FY23 across QuickBooks, TurboTax, Mint, Mailchimp, and Credit Karma; scaled runtime throughput from 6K to 50K TPS via rSocket migration supporting ~1.5M concurrent connections with sub-25ms TP99 — a textbook data-plane performance and reliability win.
- Delivered ICE Self-Service platform (DevPortal, GitOps config, ICE Playground) embodying platform-as-a-product philosophy — reducing developer onboarding from 2–3 weeks to minutes in pre-prod and <24 hours for production, while mitigating \$1M+ in projected opex growth.
- Extended Java and Python SDK Starter Kits with scaffolding templates, build configurations (Gradle/Maven), testing frameworks, and CI/CD integration — empowering developers to go from zero to production-ready microservice in minutes; authored developer and operator experience strategy across ~20 mobile apps and 30+ product SKUs.
- Initiated MSaaS Drift Detection and Resolution program: wrote Java JAR library to scan Git repos for configuration drift, partnered with Design on DevPortal UI, and built remediation roadmap using OpenRewrite — establishing safe rollout mechanics and lifecycle governance across the platform.
- Conducted enterprise-wide Service Language Assessment across 9 languages (Java, Python, Kotlin, Go, TypeScript, Scala, PHP, C++, Groovy), analyzing telemetry and usage data (SQL, BigQuery) to inform strategic investment decisions presented to CTO; built Asterias, a declarative asset lifecycle management platform with GraphQL API.
- Led Mailchimp GCP-to-AWS migration for MSaaS, delivering Golang service template, MySQL persistence integration, and updated DevPortal documentation; implemented ICE Presence in async chat generating \$480K/month in additional invoicing.

STREAMIO AI | Oakland, CA

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

- Built OpenClaw multi-agent orchestration framework with gateway protocol, subagent delegation, profile management, and session switching — production agentic integration infrastructure enabling coordinated AI agent workflows across real estate, insurance, health/dental, and financial markets.
- Integrated Claude MCP SDK to expose screen capture tools to AI coding assistants via MCP servers — hands-on experience building the infrastructure layer on which AI agents and MCP servers run reliably in production.
- Engineered real-time HLS livestreaming pipeline with multi-stream canvas compositing (up to 9 concurrent sources, 30fps), FFmpeg transcoding, WebSocket communication layer, and local HTTP playback server — demonstrating data-plane execution and runtime resilience across hybrid desktop/cloud deployment.
- Shipped cross-platform desktop builds (macOS DMG with code signing/notarization, Linux .deb) with native ScreenCaptureKit integration via Swift, embedded terminal (xterm.js/node-pty), and end-to-end auth/payments pipeline (Kinde OAuth 2.0, Stripe subscriptions, Electron SafeStorage).
- Led 0-to-1 product strategy, AI development, and go-to-market execution including customer discovery interviews and iterative product refinement.

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, structured output validation, and token budget optimization — production infrastructure supporting generative integration workloads.
- Built domain-specific conversational AI agents embedded across mobile and web, each scoped to distinct financial focal points; executed go-to-market through App Store launch following customer discovery.

SPLUNK INC. | San Francisco, CA

Senior Product Manager — Search Orchestration 01/19 – 01/21

- Owned Search Service (Go microservices), Search Catalog (PostgreSQL metadata service), and Splunk Processing Language (SPL/SPL2) — built product roadmaps, PRDs, and acceptance criteria for Splunk Cloud Services (SCS), operating as a platform-as-a-product owner across three interdependent runtime services.
- Delivered Scheduler Service end-to-end in ~4 months (June–Oct 2019), enabling scheduled search capabilities for first-party applications; demoed at Splunk .conf19.
- Led query performance optimization initiative for beta enterprise customer (Assurance), building mirrored topology for benchmark testing and achieving up to 10x runtime performance improvements in SCS search.
- 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 08/12 – 12/18

- Led development and enterprise rollout of Splunk Logging-as-a-Service (1.7 TB daily volume, 200+ internal enterprise customers) and ITS! Application Monitoring-as-a-Service — owning the observability and telemetry platform for a regulated, hybrid on-premises enterprise.
- Built caching capability using Redis and XC10 across the enterprise to address scalability, fault tolerance, and data redundancy challenges; led demand forecasting and capacity planning 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 software; led cross-functional root cause analysis and customer-facing resolution, 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 FastAPI orchestrator, TimescaleDB, Redis job queue, Next.js dashboard, and Ollama — an observability and SLO-adjacent infrastructure layer for AI workloads with CI/CD regression detection and automated safety gates.
- 5 core eval types (factuality, reasoning, instruction-following, safety, code generation), adversarial safety testing with refusal detection, and statistical rigor: bootstrap confidence intervals, Welch’s t-test, Cohen’s d effect size, saturation detection.

RL Workbench — Post-Training RL Platform (2026)

- Built 3-phase RL post-training workbench: Reward Lab, Playground (real TRL-powered GRPO/DPO training with live SSE metric streaming on MPS/CUDA), and Arena for head-to-head framework benchmarking (TRL, VeRL, OpenRLHF, NeMo RL) with GPU passthrough in Docker containers.
- Implemented 12 RL algorithms (PPO, GRPO, DAPO, DPO, SimPO, IPO, KTO, ORPO, and more) with standardized throughput/memory/convergence benchmarking — production infrastructure for agentic and generative AI workloads.

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, and Docker orchestration (6 containers); 823 automated tests. NeurIPS 2014 accepted paper. Original 2004 system: hand-coded neural network in C++ with custom BPTT.

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

- Zero-integration architecture captures from any visualization tool via screen capture; multimodal AI (Claude/GPT-4V) performs spatial reasoning on prediction frames, generating structured PASS/FAIL reports — reducing evaluation cycles from 72 hours to ~4 minutes.

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/64G — 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