

Brian Hyun

[LinkedIn](#) | [Site](#) | 626-622-8188 | hyunbrian97@gmail.com

EXPERIENCE

Stackline

Software Engineer II, Team Lead

Seattle, WA

Feb 2026 – Present

Ad Intelligence and Audience Platform Systems

- Led product and technical execution of Stackline's **Ad Manager** Platform as team lead for 4 engineers. Drove roadmap sequencing, resolved product ambiguity, bootstrapped and managed the **Jira** board, and owned end-to-end feature execution from design through production across **50+** brands and **100M+** daily ad records using **Next.js**, **TypeScript**, **Shadcn UI**, and **AWS**. (**2026 Achiever's Award** — **3 of 233** selected)
- Defined the foundational data model for Ad Manager. Structured the canonical advertising hierarchy across Campaigns, Ad Groups, Ads, Products, and Targets spanning Search and Demand-Side Platform (DSP) domains; domain-driven modeling decisions were fundamental to the platform's ability to function and scale across multiple retailers.
- Drove platform-wide AWS infrastructure cost optimization across all Ad Manager systems. Reduced monthly cloud spend by **50%** (**\$35K** → **\$17K**) by eliminating redundant cross-region data transfers and unnecessary EBS volume snapshots.
- Owned production observability end-to-end. Instrumented the platform with **Sentry** and **Mixpanel**, built centralized **Slack alerting**, and led incident response and cross-team debugging across a 4-person team with no dedicated DevOps function.
- Scaffolded and managed **Apache Doris OLAP** clusters across dev and prod environments. Designed canonical schemas, async materialized views, indexes, and partitions; monitored cluster health via **Prometheus** and **Grafana**; reduced P95 query times to sub-second across **100M+** daily metrics.
- Built an **MCP server** for **Zenith**, an internal **agentic AI assistant** for conversational campaign analytics and optimization. Exposed Stackline's ad database as agent-callable tools (query execution, schema introspection) over **100M+** daily ad records, powered by **Claude Sonnet** via **AWS Bedrock** and **Amazon Strands**.

Software Engineer I

Mar 2025 – Feb 2026

Ad Intelligence and Audience Platform Systems

- Led end-to-end development of Stackline's **Ad Manager** Platform from zero. Shipped MVP in **3 months**, owning everything from UI to data pipeline across **Amazon Ads**, **Walmart Connect**, and **Instacart Ads**; processed **100M+** ad records daily for **50+** brands using **Next.js**, **TypeScript**, **Shadcn UI**, and **AWS**.
- Architected the AWS ingestion framework (**Kinesis**, **ECS Fargate**, **DynamoDB**, **S3**) powering Ad Manager's real-time and batch data workloads. Delivered automated scaling and sub-hour data latency underpinning the analytics UI.
- Scaffolded **CI/CD** from scratch using **GitLab CI** and **AWS CDK**, automating deployments for Ad Manager and establishing the release infrastructure adopted across the engineering team.
- Spearheaded **Audience Syndication** in ShopperOS, enabling audience exports and real-time status tracking across **Amazon Ads**, **Google Ads**, **Meta Ads**, and **TikTok Ads** via secure **OAuth 2.0 PKCE** flow.
- Engineered a user segmentation pipeline compressing **5,900-dimensional** preference vectors for **900K+** users via custom **autoencoder** + **MiniBatchKMeans**; co-developed a **polynomial contextual bandit** engine to recommend optimal campaign actions based on real-time context.

Pangea Laboratory

Software Engineer

Tustin, CA

Jun 2021 – Mar 2025

Laboratory Information Management System (LIMS)

- Designed and built a full-stack **React** + **FastAPI** LIMS platform as the sole engineer, displacing a **~\$300K+/year** third-party vendor solution with an in-house build, owning the entire stack from UI to deployment and powering end-to-end accessioning, processing, and reporting pipelines.
- Built responsive, role-based dashboards in **React** and **Tailwind CSS** for **10+** lab technicians and **5+** clinical scientists; partnered with scientists and operations staff to streamline data-entry UX, reducing submission errors by **25%** and cutting report turnaround by **10+ hours/week**.
- Implemented **Auth0** authentication and **RBAC** across three organizations; automated deployments with **Docker**, **GitHub Actions**, and **AWS Amplify/Elastic Beanstalk**, reducing release time by **40%**.

EDUCATION

University of Southern California

B.A. in Cognitive Science, Summa Cum Laude

Los Angeles, CA

TECHNICAL SKILLS

Languages & Frameworks: TypeScript, JavaScript, React, Node.js, Next.js, Python, FastAPI, Tailwind CSS

Cloud & Infrastructure: AWS (Cognito, ECS, Fargate, Kinesis, Step Functions, SQS, EventBridge, DynamoDB, S3, CDK), Docker, Nginx, Unicorn, CI/CD (GitLab CI, GitHub Actions)

Data & Systems: PostgreSQL, Apache Doris (OLAP), MongoDB, REST & GraphQL API design, Nx Monorepo Management

Observability & Tooling: Prometheus, Grafana, Sentry, Mixpanel, Auth0, Figma

AI & Agents: Claude Code, Cursor, AWS Bedrock, Amazon Strands, Model Context Protocol (MCP server development, tool use)