

Ken King

Software Engineer | Frontend & Full-Stack Product · Applied AI | TypeScript, React, AWS

San Francisco Bay Area | kenkingtech@gmail.com | linkedin.com/in/kking935 | kenk.ing

Frontend-leaning full-stack engineer from the Amazon Bedrock console: led the Intelligent Prompt Routing frontend for re:Invent 2024, built the Agent Evaluations experience, and cut UI release lead time from 2+ weeks to hours. Then founding engineer at MASLOW Health, shipping a patient-facing clinical AI web app with EHR integration.

RELEVANT EXPERIENCE

Founding Engineer — MASLOW Health · Healthcare AI startup

Nov 2025 – Present

- **Product ownership:** Built and deployed the patient-facing React/TypeScript clinical AI web app as sole full-time engineer: voice and form intake, draft persistence and recovery, responsive layouts; iterated onboarding from patient and clinician feedback; managed interns and contractors.
- **Clinical AI workflows:** Built voice intake (AWS Transcribe) feeding guardrailed Bedrock workflows on Step Functions and Lambda; delivered PDF reports into eClinicalWorks via FHIR (server-to-server OAuth) for clinician approval.
- **Offline deployment:** Ported the app to Rust/Tauri with bundled Qwen/Whisper for pediatric clinical research.

Amazon Web Services · Santa Clara, CA

Software Development Engineer — Amazon Bedrock

Jul 2024 – Nov 2025

- **Product UI:** Led Intelligent Prompt Routing's frontend deployment for re:Invent 2024; built React/TypeScript features for Bedrock Flows, Prompt Optimization and Chat Playground.
- **Agent evaluation:** Built Bedrock Agent Evals UI for benchmarks, automated metrics, LLM-as-judge, human review and traces; deployed internally and production-launch ready before redirection to AgentCore.
- **Shared-platform release engineering:** Separated the team's console codebase and release pipeline from a shared monolith, cutting UI release lead time from 2+ weeks to hours; drove staged, feature-flagged rollouts across 12+ regions and GovCloud.
- **Cross-team ownership:** Go-to engineer for console release coordination across internal and partner teams; removed a years-old Weblab dependency from a legacy service outside the team's ownership.
- **On-call:** Triageed 50+ Sev-2 incidents in a single Bedrock rotation, debugging backend and cross-service failures.

Software Development Engineer Intern (2x) — SageMaker Ground Truth Plus

Summers 2022 & 2023

- **LLM evaluation (2023):** Built a full-stack LLM evaluation prototype: user-selected Hugging Face models and benchmarks, hosted SageMaker inference, automated scoring, expert review, infrastructure as code.

Graduate Research — Virginia Tech · Human-AI interaction · 5 peer-reviewed papers

Jan 2022 – Jul 2024

- **Research:** First-author HCII 2024 paper (SafeSpace, a dementia-care voice assistant); built a clinical-text benchmark from three datasets and compared fine-tuned ClinicalBERT with few-shot Llama 3.

Head of Product Development — Vector Rideshare

Feb 2020 – Sep 2021

- **Product development:** Promoted from intern to technical lead; built React/TypeScript, Node.js and GraphQL workflow apps; led a seven-person intern team.

TECHNICAL SKILLS

Frontend: TypeScript, React, Next.js, TanStack Query, Zustand, React Hook Form, Zod, Tailwind, shared component libraries, feature flags

Backend & Cloud: Python, Node.js, FastAPI, GraphQL, REST, Rust, SQL; Lambda, Step Functions, Cognito, CDK, CloudWatch; DynamoDB, PostgreSQL/pgvector

AI & ML: Bedrock, SageMaker, Transcribe, Polly, PyTorch, Hugging Face, MCP, tool calling, RAG, LangGraph, evaluations, Langfuse, guardrails, human-in-the-loop

Testing & Tooling: Vitest, Jest, React Testing Library, Playwright, Cypress, GitHub Actions, Docker; Claude Code, Codex

EDUCATION

M.S. Computer Science, Virginia Tech

Jul 2024

GPA 3.95 · Thesis-based research in AI benchmarking and model evaluation · Graduate Teaching Assistant

B.S. Computer Science, Virginia Tech

May 2021

GPA 3.71 · Magna Cum Laude · NCAA Division I football · Undergraduate research · Completed in 3 years