Vybavnag Kandasamy

510-361-8806 | vybavnag2003@gmail.com | linkedin.com/in/vybavnag | US Citizen | Open To Relocation

Skills

Languages: C, C++, Python, Java, R, TypeScript, JavaScript, Golang.
Frameworks: React, Node.js, Flask, FastAPI, JUCE.
Developer Tools: Git, AWS, Docker, Firebase, Google Cloud Platform, MySQL, Perforce.
Libraries: Pandas, NumPy, Matplotlib, TensorFlow, PyTorch.
Certifications: AWS Certified Cloud Practitioner, NVIDIA Fundamentals of Deep Learning.

Experience

Software Engineer I **Aug 2025 – Present**
UKG *Sunrise, FL*

- Instrumented all AI pillar microservices with OpenTelemetry, enabling end-to-end distributed tracing reducing **mean time to debug production incidents by 40%.**
- Optimized DLP pipeline and OTEL Collector in Go with Google Cloud DLP table format, **cutting latency by 74%** and eliminating span loss.
- Built a Go trace processor to extract agent metrics (input/output, intent, tool calls) from OTEL spans into BigQuery, cutting data scientist **query turnaround by 53%**.
- Architected a Google Pub/Sub feedback loop for real-time customer feedback ingestion, enabling analysts to act on insights without manual data pulls.
- Integrated Grafana and Arize across AI microservices, centralizing observability for service monitoring and model evaluation.

Software Engineer Capstone **Jan 2025 – June 2025**
Universal Audio (Professional audio hardware and software manufacturer) *Santa Cruz, CA*

- Researched and developed a 4-input, 4-output audio device using a microcontroller with real-time digital signal processing (DSP) capabilities.
- Designed a system on a dual-processor microcontroller to offload DSP and console workloads.
- Engineered DSP algorithms using JUCE/C++ and optimized them using techniques like Decimation and FFT, reducing **processing latency by 25%**.
- Used DMA optimizations to tune Linux kernel parameters on the microcontroller reducing audio latency **below 5 ms**.

Software Engineer Intern **Sept 2024 – Dec 2024**
TextBuddy *Santa Cruz, CA*

- Built an SMS-based application using Java and Spring enabling LLM access for users without internet connectivity.
- Integrated LiteLLM API to unify access across multiple LLMs, **reducing model-switching latency by 56%**.
- **Shortened setup time by 70%** by containerizing with Docker and automating deployment on GCP.

Projects

Gateway | *AWS, React/Vite, Docker, Kubernetes, GitHub Actions* **Jan 2026 – Present**

- Built an AI-native learning platform where users design distributed systems on a visual canvas with components mapped to real cloud services.
- Developed AI chatbots that review user architectures in real-time, suggest improvements, and surface bottlenecks in system designs.
- Enabled one-click Terraform export and deployment, letting users go from visual design to live cloud infrastructure.
- Deployed on AWS EKS with Helm-managed microservices and CI/CD via GitHub Actions.

Education

University of California, Santa Cruz **June 2025**
Bachelor of Science in Computer Science