

YASH AMOL KANKAL

+1 602-565-5915 | yashamolkankal@outlook.com | linkedin.com/in/yashkankal | github.com/yash-kankal

EDUCATION

Master of Science in Information Technology

Arizona State University, United States

Relevant Coursework: Distributed Systems, Algorithms & Data Structures, Database Systems, Cloud Computing, Computer Networks

Aug 2023 – May 2025

GPA: 3.6/4.0

Bachelor of Engineering in Information Technology

Prof. Ram Meghe Institute of Technology and Research, India

Relevant Coursework: Data Structures & Algorithms, Computer Networks, Database Management, Software Engineering

Jul 2019 – Jun 2022

GPA: 3.6/4.0

TECHNICAL SKILLS

Languages	: Java, Python, C++, TypeScript, JavaScript, Kotlin, Swift, SQL
Backend & APIs	: Node.js, Express.js, FastAPI, REST APIs, WebSockets, OAuth 2.0, JWT
Frontend & Mobile	: React.js, Next.js, Jetpack Compose, SwiftUI, MVVM, Retrofit, Coroutines
Databases	: PostgreSQL, Redis, MongoDB, MySQL, Elasticsearch, AWS RDS
Cloud & DevOps	: AWS (Lambda, S3, Kinesis, RDS, ECS, EC2, App Runner, CloudFormation), Docker, Kubernetes, GitHub Actions, CI/CD
Foundations	: Data Structures & Algorithms, Distributed Systems, System Design, Machine Learning Systems, Testing & Debugging, Code Reviews

EXPERIENCE

DriverAI, LLC, Software Engineer, Tempe, AZ / Remote, USA

Jan 2026 – Present

- Architected a distributed real-time backend system using WebSockets and Redis to synchronize state across mobile clients under variable network conditions, reducing perceived latency by 40%.
- Engineered scalable backend APIs using FastAPI and PostgreSQL with schema validation, rate limiting, and centralized exception handling, reducing API response latency by 25%.
- Built a mobile-to-ML inference pipeline capturing live camera frames via CameraX and transmitting them to backend ML services for real-time product recognition, enabling automated visual identification across Android and iOS clients.
- Automated AWS infrastructure provisioning (RDS, S3, Lambda, App Runner) via GitHub Actions CI/CD pipelines, eliminating manual deployment steps and reducing provisioning overhead by 40%.
- Applied algorithmic reasoning to optimize in-memory state propagation and real-time update flows across distributed mobile and web clients, improving feature delivery velocity by 30%.
- Developed unit and integration tests for backend APIs and real-time session workflows, improving system reliability and reducing post-release regressions.

Pangian, Full Stack Developer Intern, Remote, USA

Jul 2025 – Aug 2025

- Designed scalable backend APIs using Node.js, Express.js, and TypeScript with PostgreSQL and Prisma ORM; applied query optimization and Redis caching to reduce API response latency by 30% under high-concurrency workloads.
- Automated CI/CD pipelines using AWS CodePipeline, GitHub Actions, and CloudFormation infrastructure-as-code, eliminating manual deployment steps and reducing release cycle time across multiple environments.
- Developed React.js and Next.js frontend features integrating secure REST API contracts to improve maintainability across shared application workflows.

AppWelt Pvt. Ltd, Full Stack Developer, Nagpur, India

Oct 2021 – Jun 2023

- Engineered and scaled 12+ production microservices using Node.js, Express.js, and TypeScript; deployed on AWS EC2 and ECS with Docker-based CI/CD pipelines, improving release velocity by 15%.
- Improved database throughput by 40% under peak load through PostgreSQL query optimization, index restructuring, and MongoDB aggregation pipeline tuning for backend-heavy workloads.
- Developed React.js and Next.js frontend components integrated with backend APIs, reducing page load time by 35% and increasing session duration by 20% through Redis caching and Elasticsearch indexing.

PROJECTS

Video Sharing Platform | TypeScript, Node.js, Express.js, PostgreSQL, Redis, Elasticsearch, AWS S3, Kinesis

- Engineered a scalable video streaming platform supporting 100+ concurrent uploads and distributed media storage using AWS S3 and PostgreSQL.
- Built event-driven analytics pipelines using AWS Kinesis for real-time processing of streaming metrics; applied Redis caching and Elasticsearch indexing to enable sub-second content discovery.
- Designed RESTful backend services for authentication, upload, playback, and search with algorithmic indexing strategies to optimize content delivery at scale.

Real-Time Mobile ML Inference Pipeline | Kotlin, CameraX, FastAPI, WebSockets, Jetpack Compose, PostgreSQL

- Engineered a mobile-to-backend ML inference pipeline capturing live camera frames via CameraX and transmitting them to backend services for real-time model inference and product recognition.
- Implemented WebSocket-based streaming to deliver low-latency predictions and asynchronous UI updates to mobile clients under variable network conditions.
- Designed backend processing workflows for scalable image ingestion, inference handling, and response propagation across distributed clients.