SUNNY ANIL KHANORKAR

McLean, VA | 571-347-0496 | sunnykhanorkar22@gmail.com | LinkedIn | GitHub

Summary

Software Engineer with 4+ years of backend and full-stack experience in healthcare technology, specializing in **Python microservices**, **REST APIs**, and **agentic AI** systems. Proven track record building high-throughput clinical data pipelines, **ML**-integrated workflows, and cloud-native infrastructure on **AWS** and **GCP**. Currently architecting **GenAI**/agentic solutions at GE HealthCare; prior roles at Fortis Healthcare and HealthPliz delivered measurable reductions in system latency, operational costs, and manual processing overhead.

Work Experience

GE HealthCare **Jan 2025 – Present**

Software Engineer – AI/Backend (Agentic AI)

- Architected a multi-agent **agentic AI** platform using **LangGraph** and **AutoGen** to automate radiology report triage, reducing radiologist review time by 38% across 3 hospital integrations.
- Built and deployed **Python**-based **microservices** on **AWS EKS** serving 500K+ monthly inference requests with p99 latency under 180ms; **Kubernetes** HPA scaled pods automatically during peak diagnostic loads.
- Designed a **RAG pipeline** over 2M+ clinical documents using **Elasticsearch** + **OpenAI** embeddings, cutting unstructured data query time from 14s to 1.8s.
- Engineered **LLM** orchestration layer with **LangChain** tool-calling to interface with 6 downstream **FHIR**-compliant **EHR** APIs, eliminating 4 manual integration touchpoints.
- Implemented real-time event streaming via **Apache Kafka** to feed agentic decision loops with live patient vitals, enabling sub-second context updates for on-call alert agents.
- Built monitoring dashboards in **Datadog** and **Grafana** tracking agent task completion rates, token costs, and error traces; reduced mean time-to-detection for agent failures from 22 min to 4 min.
- Deployed **infrastructure-as-code** (**Terraform**) for all AI services across **AWS** dev/staging/prod environments, reducing environment provisioning time from 3 days to 45 minutes.
- Collaborated with clinical AI researchers to evaluate 4 **LLM** candidates (GPT-4o, Claude 3, Llama 3, Mistral) on **HIPAA**-compliant benchmarks; selected final model achieving 91% diagnostic note accuracy.

Fortis Healthcare **Apr 2021 – Jul 2023**

Software Development Engineer II – Backend

- Owned backend architecture for Fortis's patient portal serving 1.2M+ registered users; migrated monolith to 11 **Python FastAPI microservices** on **AWS**, reducing deployment cycle from 3 weeks to 2 days.
- Designed and implemented an **event-driven** appointment scheduling engine using **AWS SQS** + **Celery**, processing 80K+ bookings/day with zero message loss across 28 hospitals.
- Built a real-time clinical data pipeline (**Python** + **Apache Kafka** + **Spark Streaming**) ingesting lab results from 6 diagnostic partners, reducing report delivery latency from 40 min to under 3 min.
- Integrated a fraud detection **ML** model (**scikit-learn** gradient boosting) into the billing API, flagging 94% of duplicate insurance claims and saving ~$1.2M annually in erroneous reimbursements.
- Led **PostgreSQL** schema redesign for patient records with partitioning and composite indexing, cutting average query execution time by 62% on a 500M-row table.
- Containerized all backend services with **Docker** and deployed to **AWS ECS**; wrote **CI/CD** pipelines in **GitHub Actions** cutting manual release effort by 70%.
- Built a **HIPAA**-compliant audit-logging service using **DynamoDB Streams** and **AWS Lambda**, capturing every PHI access event with 100% replay fidelity for compliance audits.
- Mentored 3 junior engineers on code review standards, system design patterns, and test coverage targets; team unit test coverage rose from 41% to 87% over 6 months.

HealthPliz **Feb 2020 – Mar 2021**

Associate Software Developer – Backend & Full Stack

- Developed 35+ **REST API** endpoints in **Python** (**Django REST Framework**) for patient onboarding, teleconsultation booking, and pharmacy ordering; APIs served 15K+ daily active users at launch.

- Built full-stack telemedicine module (**Django** + **React**) integrating **WebRTC**-based video consultations with automated appointment reminders via **Twilio** SMS/email, reducing no-show rate by 27%.
- Designed **MySQL** schema with normalized tables and stored procedures for prescription and drug inventory management; query response time averaged under 80ms at 10K concurrent sessions.
- Deployed **Django** application to **AWS EC2** behind an **Application Load Balancer** with auto-scaling groups; sustained 99.7% uptime during a 4x traffic spike during COVID-19 lockdown.
- Implemented **Redis** caching layer for doctor profile and availability endpoints, reducing database read load by 55% and cutting API response time from 420ms to 90ms.
- Wrote **Dockerfiles** and **docker-compose** configurations for local and staging environments, eliminating "works on my machine" defects and reducing onboarding time for new developers by 3 days.
- Integrated **Razorpay** and **Stripe** payment gateways with webhook-based reconciliation, processing INR 8M+ monthly transactions with <0.1% payment failure rate.
- Collaborated directly with 2 product managers to translate clinical workflow requirements into technical specs, delivering all 3 MVP modules 1 week ahead of schedule.

Technical Skills

Languages: Python, Java, JavaScript, TypeScript, SQL, Bash

Frameworks & Libraries: FastAPI, Flask, Django, Node.js, Express.js, React, LangChain, LangGraph, AutoGen, Celery, gRPC, GraphQL, REST APIs, Microservices, Event-Driven Architecture, WebSockets

AI / ML: Agentic AI Systems, RAG Pipelines, Advanced Retrieval (Hybrid Search, Reranking), LLM Fine-Tuning (LoRA, PEFT), Prompt Engineering, OpenAI APIs, Function Calling, Structured Output (JSON Schema), LlamaIndex, Hugging Face Transformers, Vector Databases (FAISS, Pinecone, pgvector), Embedding Models, Model Evaluation (RAGAS, LangSmith), Hallucination Detection, Guardrails, scikit-learn, TensorFlow, PyTorch

Cloud & DevOps: AWS (EC2, S3, Lambda, SQS, RDS, EKS), GCP (BigQuery, Pub/Sub, Cloud Run), Docker, Kubernetes, Terraform, CI/CD (GitHub Actions, Jenkins), Serverless Architecture

Databases & Big Data: PostgreSQL, MySQL, MongoDB, Redis, Elasticsearch, Apache Kafka, Apache Spark, Hive

Tools: Git, JIRA, Postman, Datadog, Grafana, PagerDuty, OpenTelemetry, Prometheus, SonarQube

Education

Virginia Polytechnic Institute and State University (Virginia Tech) — Alexandria, VA
Master of Engineering in Computer Engineering — Aug 2023 – May 2025

Nagpur University — Nagpur, India
Bachelor of Engineering in Computer Science — Jun 2018 – Jul 2022