ANAGHA DEVENDRA PATIL

Raleigh, NC | anaghadevendra.patil@gmail.com | Portfolio | LinkedIn | GitHub

WORK EXPERIENCE

Software Engineer, INSAIT Solutions, Remote, US FEB 2026 – PRESENT

- Engineered an AI healthcare management platform using Django, React, PostgreSQL, and AWS, expanding the system across 10+ modules and 4+ user roles covering patient consultations, clinical documentation, prescriptions, and billing.
- Developed GenAI clinical documentation workflows by orchestrating LLM prompts, structured outputs, clinician review, and approval processes, reducing manual documentation workload by 50%.
- Built secure multilingual audio and document processing pipelines using encrypted S3 storage, AWS Transcribe, and AWS Translate, enabling transcription, translation, clinical report processing, and traceable file processing states.
- Partnered with clinicians and stakeholders to translate operational requirements into production APIs and user workflows, conduct User Acceptance Testing, and strengthen traceability through validation, approval tracking, and audit trails.

Software Engineer, Triton Sensors, Remote, US MAY 2025 – FEB 2026

- Led development of Triton Insights, a GenAI analytics platform using AWS Bedrock, Django, and InfluxDB to generate actionable device insights, reducing data retrieval latency by 50% and enabling self-service analytics.
- Designed a real-time LLM chatbot using WebSockets, AWS Bedrock, Django, Redis caching, and asynchronous request handling, reducing average response latency by 30%.
- Optimized GenAI inference through batching, caching, and asynchronous execution across AWS services, cutting processing costs by 20% while maintaining sub-second response times.
- Implemented an alert notification logging pipeline using AWS Lambda, InfluxDB, and secure site-scoped Django APIs to query alert trigger history at 100K+ events per day.

Software Engineer, FlexGen, Durham, NC JUL 2024 – MAY 2025

- Owned an asynchronous Ansible and Python automation framework within a Flask SaaS application, enabling multi-host deployments, live log updates, and per-host status controls while reducing execution time from 2 hours to 1.2 hours.
- Developed REST APIs and configuration management workflows supporting 50+ sites, eliminating manual GitHub downloads and improving deployment reliability.
- Automated Pytest validation across three integration branches and CI workflows, eliminating 10+ hours per week of manual sanity checks and preventing regressions before deployment.

Software Engineer Intern, FlexGen, Durham, NC MAY 2023 – AUG 2023

- Automated server deployment and Grafana and InfluxDB configuration across 30+ servers using Ansible, cutting deployment time by 40% and reducing manual setup effort.
- Enhanced GitHub Actions validation pipelines using Pylint and automated tests while parallelizing AWS build workflows, reducing build time by 50% and improving release reliability.

TECHNICAL SKILLS

- **Programming Languages:** Python, Go, Java, JavaScript, TypeScript, SQL.
- **Backend and AI:** Django, Flask, FastAPI, REST APIs, WebSockets, Server-Sent Events, LangGraph, AWS Bedrock.
- **Frontend:** React, Next.js.
- **Cloud and DevOps:** AWS (S3, Lambda, Transcribe, Translate), Docker, Ansible, GitHub Actions, Linux.
- **Data, Testing, and Observability:** PostgreSQL, MySQL, MongoDB, InfluxDB, Redis, Pytest, JUnit, Grafana.

PROJECT EXPERIENCE

Trade Order Management System JUN 2026 - PRESENT

- Built a Java Spring Boot trade order management backend with REST APIs for portfolios, holdings, orders, and transactions, supporting validated BUY and SELL workflows with PostgreSQL persistence and Docker based local setup.
- Implemented order validation, execution, cancellation, transaction history, filtering, structured exception handling, and 19 JUnit tests to enforce state transitions and maintain consistent cash balance, holdings, and order records.

AI Incident Commander MAY 2026 – PRESENT

- Built an agentic incident response platform using FastAPI, Next.js, LangGraph, and Docker to correlate 10+ log patterns, classify incident severity, stream live analysis, and generate structured remediation plans.
- Designed a multi-agent workflow with state-based orchestration, asynchronous APIs, and structured outputs to coordinate log analysis, impact assessment, and remediation planning.

EDUCATION

Master of Computer Science, North Carolina State University, Raleigh, NC, USA AUG 2022 - MAY 2024

Bachelor of Engineering in Computer Engineering, University of Mumbai, India AUG 2018 - JUN 2022