

Muazzam Ali

Software Development Engineer (Backend & AI)

📞 (+91) 7668177718 — ✉ Email — 🌐 Portfolio — 🌐 LinkedIn — 🌐 GitHub — 📄 LeetCode

Professional Summary

Backend Engineer with experience in building high-concurrency systems handling **5x traffic spikes** and autonomous AI agents reducing manual work by **90%**. Proven: **25% latency reduction**, **30% cost optimization**, production-grade REST APIs, resilient caching strategies.

Professional Experience

Loyalty Juggernaut Inc.

April 2025 – Present

Product Engineer

- Reduced API latency by **25%** for 10,000+ daily users by engineering 15+ endpoints with **Django REST Framework** and optimizing complex **PostgreSQL** queries
- Enabled the system to handle **5x traffic spikes** during peak loads by architecting a sliding-window rate limiter and caching strategy using **Redis**, improving response times by **40%**
- Eliminated unauthorized access vectors by implementing strict domain whitelisting and IP-based access control middleware, significantly hardening the API security posture
- Built AI-ready backend foundations: stable request/response schemas, deterministic error contracts, and strict validation to safely support LLM/RAG feature integration.
- Improved reliability for heavy/AI-like workloads using background processing patterns (timeouts, retries, idempotency) and better observability (latency + failure reason logging).

Loyalty Juggernaut Inc.

Jan 2025 – April 2025

Product Engineer Intern

- Increased team sprint velocity by **15%** by delivering 12 production-ready features within 3 months through efficient Agile/Scrum development cycles
- Reduced direct database query load by **30%** by implementing a multi-layer caching mechanism using **Redis**, optimizing cloud infrastructure costs
- Standardized validation, error handling, and reusable service-layer logic across APIs to reduce regressions and simplify future ML/LLM module integration.

Queuify

Feb 2024 – July 2024

Machine Learning Engineer Intern

- Achieved **85%** validation accuracy on predictive models by building end-to-end Machine Learning pipelines using **TensorFlow** and **Scikit-learn**
- Improved model performance by **20%** for core business cases by fine-tuning deep learning architectures using Transfer Learning techniques
- Strengthened ML engineering workflow by formalizing train/eval runs (clean splits, metric tracking, and error analysis) to guide model iteration.
- Packaged inference into a clean interface (single + batch predictions) to support easier backend/product integration.

Technical Skills

Backend Engineering: Python • Django/DRF • FastAPI • Celery • Flask • JWT/RBAC

Database & Caching: PostgreSQL (optimization, indexing) • Redis (patterns, rate limiting)

GenAI: LangChain • LangGraph • OpenAI API • HuggingFace • RAG Pipelines • Ollama • TensorFlow

Tools & Platforms: AWS • Git • GitHub • Postman • Docker

Technical Projects

Distributed Ticket Booking Engine 🔄 – Flask, Redis, PostgreSQL, Docker

- Architected a high-concurrency reservation system using Redis Distributed Locking (Redlock) and Celery to eliminate race conditions and ensure zero double-bookings during traffic spikes.

Distributed Inventory Management API 🔄 — Django REST Framework, Redis, Docker

- Engineered a high-concurrency REST API with **Django** and **PostgreSQL**, utilizing **Redis** caching to maintain sub-100ms latency under high load. Secured sensitive endpoints with Role-Based Access Control (RBAC) via JWT

Education

Maulana Azad National Urdu University, Hyderabad (A Central Government University)

July 2021 – April 2025

Bachelor of Technology in Computer Science

CGPA: 8.54/10.0

Certifications & Achievements

- Solved **230+** **Data Structures & Algorithms** problems on LeetCode with **1,649** contest rating
- Machine Learning Specialization (Stanford) • MIT Computational Thinking