

Muhammad Zulqarnain

Senior Software & Applied AI Engineer

zulqarnain.mailbox@gmail.com · (+92) 331-7431748 · Lahore, Pakistan · LinkedIn · GitHub

SUMMARY

Senior Software Engineer with nearly 10 years of experience building and scaling backend systems in Python, Django, and FastAPI. Maintained infrastructure across 140+ microservices for a platform serving 40M+ learners, and built data pipelines and enrichment systems covering 500k+ companies and 2M+ founders. Hands-on experience building agentic AI systems, LLM-powered workflows, retrieval-augmented generation, and a production MCP server exposing platform data to Claude, using OpenAI and Claude APIs from design through production.

EXPERIENCE

Senior Software & AI Engineer *April 2026 – Current*

Independent Consulting · Remote

Stack: Python · Django · FastAPI · PostgreSQL · AWS · OpenAI · Claude · MCP · RAG · Agentic AI Workflows

- Independent consultant vetted through Toptal, working with clients on Python backend systems, applied AI and LLM integration, and inheriting and hardening existing codebases.
- Built an open-source RAG system end to end: structure-aware chunking with header context, embeddings in PostgreSQL with pgvector, two-stage retrieval with a cross-encoder reranker, grounded answers with inline citations, and an evaluation harness measuring hit rate, MRR, answer faithfulness and abstention.

Senior Data & Software Engineer *May 2022 – Mar 2026*

Basis Set Ventures

Stack: Python · Django · FastAPI · PostgreSQL · pgvector · AWS · Docker · Apache Airflow · Argo Workflows · Next.js · MCP · Agentic AI Workflows

- Built and maintained the firm's core investment platform (Pascal) as one of three engineers, owning the backend and data layer across 500K+ companies and 2M+ founder profiles.
- Proposed and led the migration of Pascal from Django templates to FastAPI and Next.js, owning the architectural decision and its implementation.
- Solely designed and built an end-to-end founder enrichment pipeline that ran unattended in production for years, combining web scraping, third-party tools and verification steps, with step-level state tracking so partial failures resumed rather than silently leaving records incomplete.
- Designed and maintained data collection pipelines processing thousands of companies daily, covering company profiles, GitHub activity, founder backgrounds and social presence across GitHub, Product Hunt, YCombinator, Discord and Twitter.
- Built an MCP server exposing Pascal's investment data to Claude, designing the tool surface over 500K+ companies and 2M+ founder profiles within the platform's permission model.
- Built LLM-powered tools for investors, including an automated competitor analysis tool that turned market research taking an analyst a day into structured comparison reports, plus extraction and categorization workflows feeding the platform.
- Built a financial bookkeeping application from scratch on FastAPI and React, including migrating years of financial records out of Excel into PostgreSQL without data loss.
- Built a portfolio tracking application with an API layer tying investor activity back to companies, and replaced a complex form-based search with natural full-text search.
- Migrated pipeline orchestration from Airflow to Argo Workflows on Kubernetes, rebuilding the deployment path alongside it.

Senior Software Engineer *Oct 2021 – Apr 2022*

Stack Builders · (Client: FindKeepLove)

Stack: Python · Flask · ReactJS · MySQL · Celery · Redis · Docker · CircleCI

- Built a recurring payment and subscription billing service from scratch for a platform where fashion and lifestyle brands run sweepstakes campaigns, covering billing cycles, failed and cancelled payments, retry logic that avoided duplicate charges, and alerting.
- Worked across a Flask backend and React frontend on feature development and high-priority production issues, with background processing on Celery and Redis.

Senior Software Engineer *Jan 2018 – Nov 2021*

Arbisoft · (Clients: **edX**, **Wanderu**)

Stack: Python · Django · Scrapy · ReactJS · PostgreSQL · MongoDB · AWS · Kubernetes · Jenkins · GoCD

- Owned the infrastructure layer across 140+ microservices on the edX platform, serving 40M+ learners worldwide, working directly with an edX Engineering Manager on dependency upgrades, CI/CD, deprecations and security patching where a bad deploy was visible immediately.
- Built, maintained and debugged 50+ ground, ferry and rail travel provider integrations for Wanderu, including FlixBus, Trenitalia and RailEasy, covering thousands of stations plus significant geolocation and maps API work.
- Owned end-to-end integration for Deutsche Bahn, Trenitalia, RailEasy, GoFerry and several other providers across search, booking and credit card payment flows, including the idempotency work that ensured a retried request never produced a double booking or double charge.

Software Engineer *Nov 2016 – Jan 2018*

QC Technologies · (Client: **Vital Interaction**)

Stack: Python · SQLAlchemy · Django · MySQL · jQuery · Bootstrap

- Built features for Vital Interaction, an automated patient communication platform serving healthcare organizations across the US, working with protected health information under HIPAA.
- Integrated hundreds of US medical facilities into the appointment automation portal using two-way SMS, email and IVR workflows.
- Shipped web application features and backend services supporting mobile apps, with involvement in planning, architecture and direct client interaction.

CORE SKILLS & TECHNOLOGIES

Languages: Python, JavaScript, TypeScript

Frameworks: Django, FastAPI, Flask, ReactJS, NextJS, Scrapy

AI / LLM: OpenAI API, Claude API, MCP, Claude Skills, RAG, pgvector, Prompt Engineering, AI pipeline design

Data / ETL: Apache Airflow, Argo Workflows, Celery, Redis, Selenium, PostgreSQL, MongoDB, MySQL, ArangoDB

Cloud / Infra: AWS (EC2, S3, RDS, Lambda, ECS, EKS), Docker, Kubernetes, CI/CD (Jenkins, GoCD, Github Actions)

Other: Alembic, **Pydantic**, **pytest**, RESTful & GraphQL APIs, SQLAlchemy, System Design, Agile/Scrum

PROJECTS

FastAPI-Docs-RAG · <https://github.com/mzulqarnain1/FastAPI-Docs-RAG>

Retrieval-augmented question answering built from scratch: structure-aware chunking, embeddings in PostgreSQL with pgvector, two-stage retrieval with cross-encoder reranking, and grounded answers with citations plus refusal on out-of-corpus questions. Includes an evaluation harness measuring retrieval quality (hit@5, MRR), answer faithfulness, and abstention separately.

EDUCATION & CERTIFICATIONS

- **BS Software Engineering**, Punjab University College of Information Technology (PUCIT), Lahore
- **AWS Cloud Developer** Nanodegree, Udacity
- **React.js** Nanodegree, Udacity