

YUSUF ÖZKAN

Software Engineer | AI & Backend Development

Istanbul, Türkiye • +90 532 571 0558 • contact@ozkanyusuf.com

linkedin.com/in/ozkanyusuf • www.ozkanyusuf.com

PROFESSIONAL SUMMARY

Software Engineer with 3+ years of experience building scalable backend systems and production-grade Generative AI platforms. Specialized in agentic AI architectures, hybrid Long-Term / Short-Term Memory systems, RAG pipelines, real-time streaming infrastructure, and Human-in-the-Loop workflows for high-stakes domains. Proven track record delivering AI-powered platforms in fintech, banking, and e-commerce using Java/Spring Boot and Python. Strong cross-functional contributor with experience in DevSecOps, observability (OpenTelemetry, Kafka), and cloud-native deployments, and active engagement with technical communities through publications and knowledge sharing.

WORK EXPERIENCE

Software Engineer • Softtech (İşbank Technology Subsidiary)

Sep 2022 – Present

Istanbul, Türkiye • Data & AI Applications Team

Data & AI Applications Team — Current

- Architected an enterprise agentic AI platform delivering production-grade Generative AI assistants for banking, with focus on context-aware, humanistic conversations Worked on **İsCep Maxi** Assistants and migration to GenAI.
- Designed a hybrid Long-Term / Short-Term Memory (LTM/STM) architecture combining Milvus (vector recall), Redis (session state), and MSSQL (durable storage), orchestrated by a custom MemoryRouter layer
- Built real-time streaming infrastructure on Server-Sent Events (SSE) + Redis Pub/Sub for low-latency token delivery and step-level event propagation across distributed agent nodes, Human-in-the-Loop (HITL) workflows — approval gates, tool-call confirmation, and intervention checkpoints — to support high-stakes banking operations
- Engineered RAG and multi-agent systems using LangChain, LangGraph, Azure OpenAI function calling, and pgvector-backed hybrid retrieval; integrated Ollama for self-hosted model evaluation
- Built end-to-end observability via OpenTelemetry and a Kafka-based three-signal pipeline (logs, metrics, traces); developed real-time voice agent capabilities with Pipecat, LiveKit, and Deepgram

Data Analytics Team — Previous

- Developed backend services for an internal Data Visualization & Analytics platform using Java 11/17 and Spring Boot, Integrated Hazelcast for distributed caching; led migration from Liquibase to Flyway for database versioning, Built an AI-powered conversational analytics service leveraging GPT-3.5/4 and PaLM models for natural-language queries on dashboards, Increased unit test coverage from 35% to 70% using JUnit 5 and Mockito
- **Stack:** Docker, Jenkins, Bitbucket, Elasticsearch, RabbitMQ, MinIO, Spring Boot, Agile/Scrum

Software Engineer (Part-Time) • Sensey

Feb 2024 – Jul 2025

Delft, Netherlands (Remote) • Backend, AI & Architecture • 24 hrs/week

- Led backend and DevOps development for an e-commerce platform built on Java Spring Boot with Azure-based infrastructure, Architected and led the Payment Microservice end-to-end, integrating Adyen as the payment service provider
- Deployed services using Kubernetes and Azure DevOps pipelines; contributed to CI/CD automation and cloud scalability improvements.
- Developed an internal company chatbot and shaped the AI integration roadmap for future product features

Software Engineer Intern • Softtech

Aug 2022 – Oct 2022

Istanbul, Türkiye

- Led backend development of an Android app for personal location tracking aimed at travelers
- **Stack:** Java, Spring Boot, JWT, Hibernate, PostgreSQL

Backend Developer Intern • Red and Grey

Apr 2021 – Oct 2021

Istanbul, Türkiye

- Developed backend services for web scraping and REST APIs using Python and Node.js

Software Engineer Intern • Becure Global

Jul 2020 – Sep 2020

Kocaeli, Türkiye

- Developed a real-time posture analysis system using NVIDIA Jetson Nano, OpenCV, and OpenPose
- Project featured on the NVIDIA Developer Forum with public video demonstration

EDUCATION

B.Sc. Business Administration • Istanbul Medeniyet University

Sep 2020 – Present

Istanbul, Türkiye • Double Major

B.Sc. Computer Engineering • Istanbul Medeniyet University

Sep 2019 – Jun 2023

Istanbul, Türkiye

- **GPA:** 3.44 / 4.00
- **Thesis:** Library Reservation System Web Application for Medeniyet University Library
- **Leadership:** Co-Founder, AI Team Lead & Core Member — Google Developer Student Club; organized courses, competitions, and live trainings on YouTube

Erasmus Exchange — Computer Science • Lodz University of Technology

Dec 2021 – Jul 2022

Lodz, Poland

PUBLICATIONS

Tasan, M., Ozkan, Y., Ozgur, A.O. (2024). *Data Plateau: A Unified Analytics Platform with Intuitive Interfaces for Real-Time and ML-Driven Insights*. Orclever Proceedings of Research and Development, 4(1), 73–89. [DOI: 10.56038/oprd.v4i1.457](https://doi.org/10.56038/oprd.v4i1.457)

CERTIFICATIONS

- **Microsoft Certified: Azure Fundamentals (AZ-900)** — Microsoft, May 2024

TECHNICAL SKILLS

- **Languages:** Java (8/11/17), Python, SQL / T-SQL, JavaScript, C
- **Backend:** Spring Boot, Spring Security, Spring Cloud, Spring Data, Spring AOP, FastAPI, REST APIs, Microservices, SSE
- **AI / ML & Voice:** LangChain, LangGraph, RAG, Vector DBs (Milvus, Chroma, pgvector), Azure OpenAI, OpenAI Function Calling, Multi-Agent Systems, HITL Workflows, LTM/STM Memory Architectures, Pipecat, LiveKit, Deepgram, Cartesia TTS
- **Databases & Caching:** PostgreSQL, MySQL, Oracle, MS SQL Server, Redis
- **Messaging, Streaming & Observability:** Kafka, RabbitMQ, Redis Pub/Sub, OpenTelemetry, Elasticsearch, Kibana
- **DevOps & Cloud:** Docker, Kubernetes (OpenShift), Jenkins, Azure DevOps, Azure (AZ-900), CI/CD, MinIO S3
- **Other:** Hazelcast, Liquibase, Flyway, Git, Agile / Scrum

LANGUAGES

Turkish — Native • **English** — B2 (Independent User)