

ID CANDIDATE	ROLE	AVAILABILITY
15367	AI Backend Engineer	30 days

PROFILE RESUME

AI Backend Engineer with over 7 years of experience in data science, machine learning, and backend development. Strong expertise in designing and optimizing backend architectures for AI applications, with a focus on scalability, performance, and real-time data processing. Adept at integrating AI models into production environments, developing high-performance APIs, and deploying cloud-native AI solutions. Passionate about leveraging Python, Go, and Java.

TECHNICAL SKILLS

Python, Go, Java, Bash, SQL, TensorFlow, PyTorch, ONNX, TensorFlow Serving, Triton Inference Server, LLM, AWS, GCP, Azure, Kubernetes, Docker, Terraform, CI/CD Pipelines, Serverless, Apache Spark, Dask, Ray, Kafka, FastAPI, Flask, Django, Spring Boot, GraphQL, RESTful APIs, Prometheus, Grafana, ELK Stack, OpenTelemetry

EXPERIENCE

Artificial Intelligence

07/2022 – present

Senior AI Backend Engineer

- Architected and developed scalable backend services for AI-powered applications, ensuring high availability and low-latency responses.
- Built efficient API gateways and microservices using FastAPI, GraphQL, and gRPC to serve machine learning models.
- Implemented and optimized AI model inference services using TensorFlow Serving, Triton Inference Server, and ONNX Runtime.
- Led the development of MLOps pipelines for model deployment, monitoring, and automated retraining with Kubeflow and MLflow.
- Designed distributed data processing architectures with Apache Kafka, Apache Beam, and Flink for real-time AI applications.
- Integrated cutting-edge LLMs (Large Language Models) into production environments, leveraging Hugging Face Transformers and LangChain.
- Ensured cloud-native deployments using AWS SageMaker, GCP Vertex AI, and Azure ML.
- Introduced observability frameworks with Prometheus, Grafana, and OpenTelemetry, reducing system downtime by 50%.

Python | FastAPI | GraphQL | gRPC | TensorFlow | Triton | ONNX | Hugging Face | LangChain | Apache Kafka | Apache Beam | Flink | AWS | GCP | Azure | Prometheus | OpenTelemetry

Data Analytics

01/2020 – 06/2022

AI Backend Engineer

- Developed and deployed machine learning model services using Flask, FastAPI, and TensorFlow Serving.
- Built data pipelines for AI applications using Apache Spark, Dask, and Feature Stores (Feast).
- Implemented model versioning, retraining, and monitoring with MLflow, Prometheus, and Grafana.

- Deployed NLP-based AI models for automated text analysis and sentiment classification.
- Collaborated with data scientists to transition models from research to production.

Python | Flask | FastAPI | TensorFlow Serving | Apache Spark | MLflow | Prometheus | Grafana

Cloud Computing

03/2017 – 12/2019

Data Engineer (ML Focus)

- Built ETL pipelines for AI datasets using Apache Airflow and Spark.
- Developed real-time streaming solutions for ML applications using Kafka and Apache Flink.
- Implemented distributed model training on cloud infrastructure with Ray and Kubernetes.
- Optimized query performance for AI-driven analytics using PostgreSQL and Cassandra.

Python | Apache Airflow | Spark | Kafka | Kubernetes | PostgreSQL | Cassandra

Data Analytics

07/2015 – 02/2017

Machine Learning Engineer

- Developed machine learning models for predictive analytics using scikit-learn and TensorFlow.
- Built feature engineering pipelines for AI applications, improving model performance by 25%.
- Deployed initial AI models as microservices using Flask and Docker.
- Assisted in A/B testing and performance evaluation of ML-powered applications.

Python | TensorFlow | Flask | Docker

Information Technology

07/2014 – 06/2015

Junior Software Engineer

- Developed and maintained backend services using Python and Flask, supporting data processing workflows.
- Wrote and optimized database queries in PostgreSQL, improving data retrieval efficiency.
- Created unit and integration tests using PyTest, ensuring code reliability and maintainability.
- Gained hands-on experience with Linux and shell scripting for server management and automation.

Python | Flask | PostgreSQL

LANGUAGES

Polish: Native

English: Fluent

COURSES AND CERTIFICATION

Google Professional Machine Learning Engineer	2023
AWS Certified Machine Learning - Specialty	2022
Certified Kubernetes Administrator (CKA)	2021
Advanced AI Systems and Cloud Computing - Coursera	2020
Machine Learning Engineering for Production (MLOps) – DeepLearning.AI	2019

EDUCATION

Warsaw University of Technology	Computer Science	Master of Science	2013 – 2015
Warsaw University of Technology	Computer Science	Bachelor of Science	2010 – 2013

I consent to the processing of my personal data and other data contained in my application and CV for the purposes of the recruitment process and of future recruitment processes conducted by ITFS sp. z o.o. Al. Zwycięstwa 96/98, 81-451 Gdynia. All information about the principles of personal data processing and your rights and obligations can be found in the Privacy Policy of Data Protection: <https://itfs.com/privacy-policy/>