

AHMET FURKAN DEMIR

Data and Platform Engineer

Istanbul | +90 5525422058 | 1ahmetfurkandemir@gmail.com | https://ahmetfurkandemir.com/

Summary

Data and Platform Engineer passionate about modern data, cloud, and infrastructure technologies. Designs and manages real-time data pipelines with Apache Flink, Spark, and Kafka, writes advanced queries using Trino and Druid, and optimizes ETL pipelines with Apache Airflow.

Builds, manages, and scales containerized systems using Kubernetes, Docker, and GitLab CI/CD. Applies a DevOps mindset to ensure secure, high-performing CI/CD processes, automated infrastructure, and reliable distributed systems. Develops microservices and data processing solutions using Python and Java.

Experience

Nesine.com 2023 August - Present
Data and Platform Engineer
https://www.nesine.com/

(Company Info: Nesine is a leading online betting and gaming platform that offers a wide range of sports betting and lottery services.)

Data Engineer

- Managed Hadoop clusters, performing performance optimization, resource monitoring, maintenance, and scaling operations such as adding/removing nodes.
- Executed performance improvements, maintenance, and updates on HDFS.
- Led the migration from Ambari to Cloudera Data Platform (CDP), successfully completing migration and validation.
- Played an active role in the transition from Cloudera Data Platform (CDP) to Kubernetes, setting up Kubernetes clusters and deploying applications to production while ensuring reliability and stability.
- Implemented Iceberg and S3 (MinIO and Pure) integrations and configured Hive Metastore, Iceberg REST, and Nessie connections, increasing data read performance from the lakehouse by ~45%.
- Managed data transfers from MSSQL and other sources to S3, overseeing 10+ ETL pipelines.
- Built reports and dashboards using Apache Superset and Imply Pivot and presented them to various teams.
- Developed batch ETL pipelines with Apache Spark, actively using PySpark to read data from various sources and sink it to S3; also conducted Apache Spark and Scala testing.
- Implemented streaming-to-S3 sinks using Apache Druid, structuring streaming data for efficient writing to S3; made 10+ Kafka topics accessible to around 25 users.
- Designed and orchestrated ETL DAGs in Apache Airflow, dockerizing Python projects to run as Kubernetes Pod Operators, eliminating SSH-based execution and achieving ~30% faster performance.
- Architected Flink ETL codebase, developing 12 Flink jobs that process streaming data and deliver results to S3/Kafka for high-volume data workloads.
- Conducted CDC (Change Data Capture) experiments with Apache Flink, processing real-time database changes and testing delivery to Kafka/S3 with a focus on data consistency and performance.

DevOps

- Installed and managed Harbor as a local Docker registry, including image security scanning, cross-repo backups, LDAP integration, and full lifecycle management.
- Set up and maintained GitLab Runners, designing .gitlab-ci.yml pipelines for all DE projects using Kaniko to build and push Docker images to Harbor.
- Installed and managed ArgoCD-based CI/CD pipelines, ensuring continuous deployment, version control, and application lifecycle management; deployed around 10 applications to Kubernetes clusters using ArgoCD.
- Administered Nexus Repository, designing GitLab CI/CD pipelines for Java projects and automating deployments of compiled artifacts to Nexus.
- Installed, maintained, and upgraded Kubernetes clusters for the DE team, handling deployments, scaling, performance tuning, and security.
- Utilized Docker and containerd for containerized application development, deployment, and lifecycle management.
- Designed end-to-end CI/CD pipelines for Python/Kubernetes services, including test automation, image builds, registry push, deployments, and rollbacks.
- Implemented monitoring with Grafana, Prometheus, and Loki, creating dashboards, alerts, and centralized log management for proactive incident handling; visualized all Kubernetes logs and logs from 20 different services.
- Experimented with Kubeflow, building ML pipelines, testing training/deployment workflows, and integrating with Kubernetes for workflow orchestration and versioning.
- Integrated and monitored NVIDIA GPU nodes in Kubernetes, optimizing workloads and ensuring GPU availability for AI/ML tasks.
- Authored Helm Charts for internal and external tools, standardizing Kubernetes deployments with version control and automation; created and designed Helm charts from scratch for 4 different applications within Nesine.
- Implemented Kubeseal (SealedSecrets) for secure secret management in Kubernetes, enabling developers to store sensitive credentials safely in Git.

Platform Engineer

- Installed and maintained essential DE tools across Prod and Stage environments, ensuring security, performance, and best practices.
- Integrated LDAP/AD authentication for 15 applications, managing roles and permissions for 8 different teams, and implementing user- and team-based access control and audit logging.
- Deployed and managed Apache Kafka, Flink, and Druid clusters, handling cluster setup, security, scaling, data ingestion, job management, monitoring, and performing HA configurations and failure testing.
- Set up and managed Trino clusters, optimizing for performance and security; the clusters are actively used by 55 users across 6 different teams at Nesine.com and Atyarisi.com.
- Deployed MinIO (S3) on Kubernetes as a scalable storage infrastructure, managing buckets, access policies, replication, and monitoring to ensure high availability and performance for data pipelines and applications.
- Installed and configured Hive Metastore, Iceberg REST, and Nessie for metadata management.
- Established LDAP/AD-based SSH access within Kubernetes for secure development environments.
- Deployed and maintained Kafka UI, Superset, Pivot, and CloudBeaver on Kubernetes, with CloudBeaver enabling ~60 users to actively access S3 (Iceberg) and other databases.
- Performed installations of PostgreSQL, MySQL, and MSSQL for metadata and application support.

DS/AI

- Built and deployed FastAPI services to expose Data Science models, designing APIs capable of handling millions of users and successfully testing them for production deployment.
- Contributed to the development of crawler detection systems for Nesine.com and Atyarisi.com.
- Participated in brute-force attack detection projects for Nesine.com and Atyarisi.com.
- Supported LLM and RAG initiatives, assisting with architecture, experimentation, and deployment.
- Dockerized DS team projects and automated deployment on Kubernetes via CI/CD pipelines.

TEXINSIGHT 2023 Jan - 2023 August
AI and Backend Engineer
https://www.texinsight.com/

(Company Info: AI-based photo and video analyzing mobile apps and a platform to increase the productivity of companies.)

- Deployed AI projects end-to-end on AWS, managing infrastructure, model hosting, and scalable production deployment.
- Developed APIs and web applications using Python frameworks including Flask, FastAPI, gRPC and Django, delivering scalable and production-ready solutions.
- Developed computer vision applications using PyTorch and TensorFlow, including object detection systems and real-time human presence/liveness detection.
- Managed domain configurations and operations using Cloudflare APIs, automating DNS and security settings.
- Enabled AI model deployment on mobile devices, optimizing models for performance, memory, and inference speed.
- Deployed projects on AWS using EC2, S3, RDS, DynamoDB, and SageMaker, designing and managing end-to-end CI/CD pipelines for scalable and reliable deployments.

İkonik Technology and Coredinat UK 2022 January - 2022 September
Software and AI Engineer
https://coredinat.com/

(Company Info: Company that produces CRM software and personnel tracking applications.)

- Developed an automatic business card information extraction system using Python Flask API, Google Tesseract OCR, and Regex, enabling high-accuracy data entry into the system.
- Built a work and task distribution application with CRM automation using Python Flask, HTML, CSS, JavaScript, and PostgreSQL, improving internal task management efficiency.
- Implemented customer integration workflows using Java Spring, Confluent, Microsoft SQL Server, and Kafka, streamlining data exchange and communication processes.
- Developed and deployed AI models for computer vision and NLP tasks using TensorFlow, PyTorch, and OpenCV, including emotion detection and real-time image/video analysis.
- Implemented natural language processing applications for sentiment analysis and sentence polarity determination, leveraging AI techniques to extract insights from text data.

MOVE ON 2021 June - 2021 August
AI Engineer Intern
https://moveon.ai/

(Company Info: Company working on Agricultural Artificial Intelligence - Autonomy - Mobility.)

- Developed AI applications using TensorFlow Object Detection API and YOLO architecture, including a Python-Flask web platform to deploy and distribute models, and an application to detect and censor unwanted or immoral images.
- Presented Flask application deployment on AWS EC2, demonstrating cloud deployment strategies and scalable infrastructure best practices.
- Conducted a workshop on AI for Jetson Nano for fellow interns and published tutorials/articles on pose estimation and AI applications, sharing knowledge with the wider community.

Education 2019 - 2023
Necmettin Erbakan University
Computer Science Engineering
3.46 GPA
https://www.erbakan.edu.tr/en/homepage
Lessons

Volunteering October 2020 - June 2023
NVIDIA
DLI University Ambassador, Jetson AI Specialist and Ambassador
https://www.nvidia.com/en-us/training/educator-programs/university-ambassador-program/
Participated in the Jetson AI Specialist program, submitting projects and code written on Jetson Nano to NVIDIA, and earned the Jetson AI Specialist certification.
Acted as an Ambassador to teach and mentor others on Jetson Nano, AI, Deep Learning, and Machine Learning, conducting workshops and sharing resources with the community.

Link1, Link2, Link3

Amazon Web Services 2021 October - 2023 January
AWS Community Builder
https://builder.aws.com/connect/community/community-builders
Participated as an AWS Community Builder, sharing knowledge on cloud computing, best practices, and modern data engineering techniques through workshops, presentations, and community contributions.
Link

HUAWEI 2022 April - 2022 October
Huawei Student Developers Campus Ambassador
https://developer.huawei.com/consumer/en/programs/hsd/
Served as a Huawei Student Developers Campus Ambassador, promoting Huawei developer programs, organizing workshops, and facilitating student engagement in technology and software development initiatives.
Link

Kodluyoruz 2021 May - 2021 October
Teaching Assistant and Mentorship
https://www.kodluyoruz.org/
(Company Info: Kodluyoruz is a nonprofit working to make Turkey a talent hub in tech by creating the right pathways for youth.)
Served as a Teaching Assistant for the Istanbul Data Science Program for Social Good, supporting instructors and assisting students with course-related questions.
Mentored students in Data Science and Python, guiding them through coursework and graduation projects. Link

Interests

Skills

Kubernetes, Docker, containerd, ArgoCD, Gitlab Runner, Apache Flink, Apache Spark, Java, Maven, Nexus, Iceberg, DevOps, CI/CD, Harbor, S3, Python, FastAPI, Flask, Django, Airflow, Kafka, Apache Superset, Druid, Pivot, PostgreSQL, Microsoft SQL Server, Change Data Capture (CDC), Trino, Helm, KubeFlow, Grafana, Prometheus, Loki, Rancher, Ranger, Hive, MinIO, Nifi, Scala, HDFS, Hadoop, Cloudera Data Platfrom(CDP), Ambari, TensorFlow, PyTorch, PyQT5, WebSockets, Google Remote Procedure Call (gRPC), AWS, Amazon S3, Serverless, EC2, RDS, Docker Swarm, CloudFlare, BeautifulSoup, Selenium, Spring

Certifications

All My Certificates
https://ahmetfurkandemir.github.io/AhmetFurkanDEMIR/Course.html

Profiles

1dfurkan

AhmetFurkanDEMIR

ahmetfurkandemir

ahmetfurkandemr