



Mohamed Mokhtar Riahi

CONSULTANT CLOUD & DATA ENGINEER

Munich, Germany · 80809

+491778362685 | mohamedmokhtarriahi1998@gmail.com | [mokhtarriahi](#)

Cloud and data engineer with hands-on experience designing, building, and operating data platforms on AWS. Specialized in ETL/ELT, modern data warehousing, and infrastructure-as-code with Terraform, delivering scalable pipelines and analytics solutions for automotive and home appliances use cases.

Education

Technical University of Munich

DATA ENGINEERING AND ANALYTICS - MASTER OF SCIENCE

Munich, Germany

10/2022 - 06/2024

Technical University of Munich

COMPUTER SCIENCE - BACHELOR OF SCIENCE

Munich, Germany

10/2019 - 09/2022

Skills

Programming	Python, C/C++, Scala, Java, SQL, JavaScript, Kotlin, LaTeX, R, Object-oriented programming
Data Engineering & Big Data	Apache Spark, Apache Hudi, Apache Iceberg, Delta Lake, Apache Kafka, PostgreSQL, MySQL, MongoDB
Cloud & DevOps	AWS, Databricks, Terraform, CI/CD (GitHub Actions, Bitbucket, Bamboo), Docker, Kubernetes, Linux, Git
Data Science & ML	Machine learning, Applied regression, Data visualization (dplyr), Deep learning, PyTorch
Web	FastAPI, NodeJS, ReactJs, HTML5, CSS
Tooling & Practices	Bash scripting, GitHub workflow, Ansible, Slurm, Unit testing
Languages	English (fluent), French (advanced), German (advanced), Arabic (native)
Industry	Automotive, Home appliances

Certificates

Amazon Web Services

AWS CERTIFIED DATA ENGINEER – ASSOCIATE

Munich, Germany

12/2024

Amazon Web Services

AWS CERTIFIED SOLUTIONS ARCHITECT – ASSOCIATE

Munich, Germany

11/2024

Amazon Web Services

AWS CERTIFIED CLOUD PRACTITIONER

Munich, Germany

09/2022

Databricks

DATABRICKS CERTIFIED DATA ENGINEER ASSOCIATE

Munich, Germany

12/2025

Confluent

CONFLUENT CERTIFIED DEVELOPER FOR APACHE KAFKA

Munich, Germany

09/2025

Confluent

CONFLUENT CERTIFIED ADMINISTRATOR FOR APACHE KAFKA

Munich, Germany

09/2025

Scrum.org

PROFESSIONAL SCRUM MASTER I (PSM I)

Munich, Germany

03/2024

Professional Experience

Data Reply (Consultancy)

DATA ENGINEER

Munich, Germany

07/2024 - Present

Client: Confidential Premium Automotive Manufacturer

- Manage and optimize cloud infrastructure on AWS with a strong focus on automation, scalability, and cost efficiency using Terraform.
- Develop and maintain dbt models and integrate them with GitHub Actions to build automated deployment pipelines.
- Design end-to-end data flows where data is ingested into Amazon S3, processed using AWS Glue and Step Functions, and exposed to stakeholders and users via Amazon QuickSight dashboards.
- Leverage Python, SQL, and AWS (Glue, Step Functions, Lambda, Athena) throughout the data and analytics stack.
- Integrate with Azure APIs and UIs to ensure seamless collaboration between AWS and Azure environments and implement an email alerting system for spend thresholds.
- Implementation of a cloud efficiency score from scratch merging all data sources from all cloud accounts of the company to provide scores and cloud recommendations in a dashboard to improve cloud efficiency per cloud account and department within the company.
- Implementation of a daily report provided via Teams to the team that incorporates all of the errors that happened during the day from all resources of the project. This pipeline was implemented in AWS using Lambda, SQS Queue, and Step Functions
- **Technologies used:** AWS, Python, CI/CD, Github, ECS, Step Functions, Lambda, DBT, SQL, QuickSight, Glue, Athena

Internal Projects

- Implemented demos using industrial data platforms (HighByte for data modeling and integration, Litmus for edge connectivity and device management) to showcase end-to-end industrial analytics solutions.
- Deployed these solutions on AWS (EC2, S3, Glue, Athena, QuickSight, Grafana) to demonstrate scalable, cloud-based industrial data pipelines.
- Contributed to white papers and RFPs within the consultancy, supporting solution design and technical positioning for clients.
- Delivered internal and client-facing presentations within the consultancy on data platforms, cloud architectures, and best practices.
- **Technologies used:** HighByte, Litmus, AWS (EC2, S3, Glue, Athena, QuickSight, Grafana, Lambda, Step Functions), Python, Terraform

Data Insights GmbH (Consultancy)

WORKING STUDENT: DATA ENGINEER

Munich, Germany

10/2021 - 06/2024

Client: Confidential Multinational Home-appliances Manufacturer

- Managed AWS accounts and underlying infrastructure for data platforms used in production.
- Reduced technical debt by fixing code smells and improving reliability in Scala, Java, and Python services.
- Built and maintained data pipelines and CI/CD workflows using Bitbucket, Bamboo, and Gradle for multiple applications.
- Implemented generic Gradle build automation and parameterized templates to standardize CI/CD across applications.
- Improved performance of the production EMR cluster by orchestrating parallel execution via AWS Step Functions.
- **Technologies used:** AWS, Java, Scala, Gradle, Python, CI/CD, Bitbucket, Bamboo, EMR, Step Functions

Internal Projects

- Compared Apache Iceberg, Apache Hudi, and Delta Lake and developed a unified API to run common operations across all three table formats.
- Deployed the solution on AWS EMR and performed benchmarking using Spark History Server, summarizing results in an internal blog post.
- Supported AWS accounts migration using Terraform and helped design secure, automated infrastructure.
- Developed an internal utilization forecast application to estimate consultant capacity, expenses, and profit/loss scenarios.
- Extended an internal framework to integrate Apache Iceberg, Apache Hudi, and Delta Lake and added new features based on project needs.
- Rebuilt the utilization forecast solution as a full web application to replace a Power Apps prototype and improve usability and maintainability.
- **Technologies used:** Scala, Java, MySQL, JavaScript, EJS, CSS, HTML, NodeJS, Express, Azure AD, Apache Kafka, Apache Spark, Apache Hudi, Apache Iceberg, Delta Lake, AWS, Databricks, Git, SBT, AWS CLI, Terraform, Microsoft Power Apps, Power Automate, Excel

Personal Projects

2D Pixel Art Urban Life Simulator (Unity)

Munich, Germany

SOLO GAME DEVELOPER

07/2025 - Present

- Developing a 2D pixel art open-world urban life simulator in Unity as a solo project, covering design, programming, and content creation.
- Planning to publish the game on Steam for PC in 2026, with the potential to expand to additional platforms.
- **Technologies used:** Unity, C#, 2D game development, Pixel art, Git

Efficient Parallel Spooling Sort

Munich, Germany

MASTER THESIS

10/2023 - 05/2024

- Implement an external sort algorithm using multithreading as efficient as possible using a limited amount of memory.
- **Technologies used:** C++, cmake, multithreading

Development of a Distributed Build Platform for Isabelle

Munich, Germany

BACHELOR THESIS

05/2022 - 08/2022

- The deployment of a distributed system that builds the images produced by isabelle. Scheduling the building process between multiple nodes on the cloud and using multithreading to reduce the total build time from 8 hours to 1 hour.
- **Technologies used:** Isabelle, LRZ Cloud, Parallelism, Multi-threading, Scala, Slurm, NFS, Ansible, Linux

Chatbot development and using LLMs - FSN Captial

Munich, Germany

DATA INNOVATION LAB

10/2023 - 02/2024

- The implementation of the frontend and backend of a chatbot using LLM-models API calls and the evaluation of the results of the chatbot using machine learning models.
- **Technologies used:** Python, FastAPI, ReactJs, Bootstrap, MongoDB, Machine learning, Pandas, Docker, Docker-compose, Random Forests, Decision Trees, Langchain

Risk Lab Web Application for Allianz Global Investors

Munich, Germany

BACKEND DEVELOPER

05/2023 - 07/2023

- The development of the backend of a web application using FastAPI.
- Implementation of a CI/CD pipeline in Github. And usage of Docker.
- **Technologies used:** FastAPI, Python, Postman, MongoDB, Docker, Kubernetes, CI/CD, Github actions

Machine learning in Medical Imaging - Sim2Real OCT Style Transfer

Munich, Germany

DEVELOPER

05/2023 - 09/2023

- Implement a Cycle-GAN deep learning model to transfer synthetic OCT images into real ones and vice-versa.
- **Technologies used:** Python, PyTorch, Deep Learning

Android App Development

Munich, Germany

DEVELOPER

11/2022 - 12/2022

- Implemented an Android app using kotlin to learn app development environment.
- **Technologies used:** Kotlin, Android.