Hudhaifa Ahmed

Senior Big Data Architect | Databricks Specialist | Berlin

Professional Overview

I'm a Senior Big Data Architect and Databricks Specialist with over 13 years in the IT industry, more than 7 of those years focused on Big Data. My work centers on designing and building scalable, secure, and cloud-native data platforms that actually move the needle for businesses.

Since 2024, I've been working as a **Resident Solutions Architect (RSA)** through the **Databricks Partner Program (DPP)**. In this role, I collaborate closely with enterprise clients across industries, offering both strategic guidance and hands-on technical leadership. It's a role that sits right at the intersection of engineering and high-level architecture, exactly where I like to be.

Alongside that, I also teach **"Data Engineering with Databricks"** as an accredited trainer (ATP program), in both English and German. It's something I really enjoy, breaking down complex topics and making them actionable for teams of all technical levels.

One of my key strengths is leading **Cloudera-to-Databricks migration projects**. I help organizations modernize their data stacks without disrupting business operations, ensuring a smooth and secure transition.

I've earned **three Databricks certifications** (Data Engineer Associate, Data Engineer Professional, and Generative AI Engineer Associate), along with cloud credentials like **AWS Solutions Architect Professional** and **Azure Solutions Architect Expert**. This multi-cloud background allows me to design solutions that aren't locked into any one platform.

Technically, I have deep hands-on experience across the **Databricks ecosystem**, from Unity Catalog to DBX CLI, REST APIs, and workspace design. I'm also well-versed in the **Hadoop stack** and a strong advocate for clean, scalable **Apache Spark** applications that are built to last.

Key Areas of Expertise

- Deep experience with Databricks implementation and large-scale migration projects
- Designing cloud-native data platforms on **AWS** and **Azure**
- Building and optimizing **Apache Spark** architectures for performance and scale
- Strong focus on data security and governance, especially using **Unity Catalog**
- In-depth knowledge of the **Cloudera/Hadoop** ecosystem
- Practical, team-focused approach with **Scrum Master** and **Product Owner** certifications
- Passionate about helping organizations modernize and future-proof their data platforms

Get in Touch

- ✉ Email: info@huthaifa-abd.com
- 🌐 Website: huthaifa-abd.com

Certifications

- **Databricks Certified Data Engineer Associate**
- **Databricks Certified Data Engineer Professional**
- **Databricks Certified Generative AI Engineer Associate**
- **AWS Certified Solutions Architect – Professional**
- **Microsoft Certified: Azure Solutions Architect Expert**
- **Cloudera Certified Administrator (CCA131)**
- **Certified Professional Scrum Product Owner & Scrum Master**

Technical Expertise

Databricks & Apache Spark Ecosystem

- Deep hands-on experience with the **Databricks platform**, including Workspaces, Clusters, Notebooks, Repos, and Jobs
- Strong working knowledge of **Unity Catalog**, access control, data lineage, and governance best practices
- Skilled in **DBX CLI** and **Databricks REST APIs** for automation, DevOps, and CI/CD pipeline integration
- Designed and optimized large-scale **Spark** architectures on Databricks, with a focus on performance, scalability, and cost-efficiency
- Led multiple **Cloudera-to-Databricks migration** projects, ensuring seamless transitions and minimal disruption
- Integrated **Azure Databricks** into customer-facing products and enterprise environments
- Implemented and managed secure, multi-user Databricks environments with fine-grained access control and workspace isolation

Hadoop, Cloudera & Legacy Data Platforms

- Broad experience with the **Apache Hadoop ecosystem**: Spark, Hive, HBase, Kafka, Oozie, and Impala
- Administered and maintained **Cloudera clusters**, including installation, configuration, and day-to-day operations
- Extensive use of **Cloudera Manager**, **Hue**, and **Ambari** for monitoring and managing cluster resources
- Set up **Hadoop High Availability** (HA) for critical components: NameNode, ResourceManager, Hive Metastore, Impala, and Oozie
- Implemented enterprise-level **security and hardening**: Kerberos integration, TLS encryption, log redaction, and LAN isolation
- Managed operational tasks including replication, backup strategies, hardware replacement, and in-place upgrades
- Built and maintained **ETL workflows** using **Sqoop** and **Oozie**, aligned to enterprise data pipelines

Professional Experience

Ultra Tendency, Germany (Aug 2017 – Present)

Role: Berlin Territory Manager & Senior Lead Big Data Developer

Responsibilities:

- Led the development, implementation, and design of enterprise data solutions, including architecture reviews, product management, pre-sales support, and technical leadership.
- Applied emerging best practices in Enterprise Data Management using Hadoop and software-based platforms.
- Ensured high standards in delivering Big Data Management and Integration solutions.
- Designed and executed strategies to significantly boost the performance of both existing and new Big Data solutions.
- Developed and implemented risk mitigation strategies.
- Programmed middleware using Java, Scala, and Python.
- Built and maintained a Hadoop-based ecosystem for enterprise-wide analysis of structured, semi-structured, and unstructured data.
- Wrote scalable, high-quality application code.

Achievements:

- Consistently maintained high customer satisfaction across all client engagements.
- Successfully delivered multiple large-scale migrations and end-to-end deployments of new data systems.
- Implemented a streamlined, event-driven, near real-time engineering stack.
- Resolved numerous data quality issues in projects that had a direct impact on revenue.
- Designed, implemented, and deployed production-grade, patentable machine learning technologies at scale.
- Fostered and sustained an engineering culture rooted in Agile principles.
- Significantly improved database access and caching performance.
- Led the restructuring of a matrix organization and launched new operations in a previously unserved region.

Career Progression:

- **Berlin Territory Manager & Senior Lead Big Data Developer** – Jul 2021
- **Regional Manager & Senior Lead Big Data Developer** – Jun 2018
- **Lead Big Data Developer** – Jan 2018
- **Big Data Developer** – Aug 2017

Earthlink Telecommunications (July 2011 – Aug 2017)

Role: Research and Development Manager

Responsibilities:

- Led software development, solution design, architecture reviews, market analysis, risk assessments, stakeholder management, recruitment, coaching, technical leadership, reporting, documentation, and pre-sales support.
- Revamped R&D programs to align with organizational goals and explore new product opportunities.
- Collaborated with senior management and internal teams to interpret and act on R&D outcomes.
- Recruited, trained, and managed Research and Development staff.

- Promoted a culture of open-source adoption by migrating critical solutions to open-source platforms such as OpenStack and Ceph.

Achievements:

- Boosted employee engagement by over 150% within one year by fostering a competitive and motivating work environment.
- Enhanced project management workflows through the integration of CI/CD tools like Rundeck, Redmine, and GitLab.
- Modernized the company's VAS services, resulting in a 50% increase in user engagement.
- Led third-party solution development, including client coordination and product promotion to maximize satisfaction.
- Improved hiring efficiency by 80% through the adoption of specialized recruitment tools and refined interview processes.
- Reduced software deployment time by over 16 hours per deployment using automation and integration tools.
- Cut hardware resource usage significantly by implementing cloud technologies, resulting in better performance and lower costs.
- Increased solution uptime by introducing monitoring tools, establishing service thresholds, and providing detailed maintenance documentation.

Career Progression:

- **Research and Development Manager** – Feb 2015
- **Lead Developer** – May 2014
- **Senior Developer** – Jan 2013

Projects

> Due to NDA constraints, specific project names have been anonymized.

Ultra Tendency, Germany

> For representative case studies and client success stories from my company, please visit: https://www.ultratendency.com/en/case-studies/.

Client: Austrian Insurance Company

Role: Development Lead

Tasks:

- Development Lead of Databricks migration workstream
- Planning and design cloud architecture for new Databricks implementation
- Develop Security concept and authentication methods
- Hands on migration of an application from on prem to Databricks on Azure

Client: Banking Sector

Role: Full Stack Lead Developer (Product Owner / Architect)

Tasks:

- Requirement Engineering & Solution Documentation
- Scalable RESTful Micro service design and integration
- Planning & Supervising solutions deployments
- Planning and management of the test resources
- Assessing the progress and effectiveness of the test effort
- SCRUM & Product Life Cycle Management
- Define user task flows and interaction models
- Showcase design ideas through sitemaps, storyboards, and process flows
- Implement the latest design innovations to ensure that the product is in accordance with the cutting-edge technology
- Develop standards and procedures to determine product quality and release readiness

Client: German Insurance Company

Role: Senior DevOps

Tasks:

- Maintenance of an existing Cloudera cluster landscape (HDFS, Spark, Kafka, Hive)
- Planning and execution of performance enhancements

Client: German Telecommunications Company

Role: Lead developer in a Cloud-based Big Data project on Microsoft Azure

Tasks:

- Micro service development with Spring Boot and Kubernetes
- Cloudera on Azure, Databricks on Azure
- Introduction and maintenance of a KPI system to monitor and improve the code quality based on SonarQube

Client: German Automobile Company

Role: Cloudera Expert

Tasks:

- Construction, installation and automation of a Cloudera cluster
- Kerberizing of cluster services and TLS for Web Frontends
- Automatic cluster deployment using Puppet
- Cluster Administration
- HBase, Spark, Kafka, Impala

Client: Service Provider

Role: Big Data Developer

Tasks:

- Construction, installation, and automation of a Cloudera cluster
- ETL Workflow under deployment of Sqoop und Oozie
- Cluster Administration

Client: German Financial Institution

Role: Big Data Developer

Tasks:

- Construction and hardening of a Hortonworks HDP Cluster
- Connection of HDP Cluster to SAP HANA
- Automatic cluster deployment using Ansible
- Cluster administration

Client: Service Provider

Role: Big Data Developer

Tasks:

- Construction and installation of a Cloudera cluster
- ETL from MS SQL Server using Talend

Client: Telecommunication Provider

Role: Big Data Developer

Tasks:

- Software Development (Enterprise Web / Mobile Applications)
    - Spring Boot Framework, Java, Maven, SBT
- Solution design (Distributed Microservices, Scalable SQL Databases)
    - HBase, Oracle, DB2, MongoDB, MySql, MSSQL
- Software Architecture review
    - Microservices, MVC, Entity, Openstack, Foreman, VSphere, KVM
    - Windows VDI, CEPH, Ansible, Chef, Puppet
- Market analysis, Risk assessment & stakeholder management
- Recruiting, Coaching, Technical leadership
- Reporting, Documentation

Earthlink Telecommunications

Additional role and task details available upon request.

- **On-Premise Cloud Solutions [OpenStack] (2014)** – Cloud Architect
- **LMS Manifest System (2013)** – Lead Developer
- **Zain DOS System (2012)** – Senior Developer
- **Multiplayer Gaming System (2011)** – Senior Developer

- **Human Resource Management System – HR (2013)** – Full Stack Developer
- **TV Streaming Service Web Application (2012)** – Full Stack Developer
- **Content Management System (2013)** – Full Stack Developer
- **eGovernment Centralized Archiving Solution (2013)** – Systems Architect

Education

- Bachelor of Science in Computer Science, University of Tikrit (2007 - 2011)
  - 3D Now [Augmented Reality Project] (2011) - BC Graduation Project

Languages

- Arabic (native)
- English (C2)
- German (C1)

Interests

- Travel to share knowledge on conferences, meetups, podcasts, workshops, mentorship, etc.
- Solving complex IoT / Big Data problems
- Figuring out best & worst practices (by asking "Why?"), and publishing them at Professional Blog and Game Development Blog
- Cloud Computing, Hadoop, Big Data Analytics, Internet of Things, Data Warehouse, Business Intelligence, Cross Platform Enterprise Software Development, Distributed Storage Solutions, Process Automation, CMS/Ecommerce Systems, System Administration, Database Management
- Exploring the world of AI and LLM technologies and utlizing them into real life usecases.