



APACHE KAFKA

MASTER OF DATA STREAMS DISTRIBUTED STREAMING PLATFORM

Apache Kafka, which originally was developed as a message queuing system within LinkedIn, is a distributed streaming platform with a broad range of functions. It is used for processing and storing all kinds of data streams. Apache Kafka is particularly suitable for the transmission of metrics or system logs. It was subsequently open-sourced in 2011, making its way into a vast amount of modern data-centric IT systems. A very active open-source community is playing its part to add to its popularity. The distributed architecture is optimized for high-fault tolerance and scalability. The risk of a total-system failure is drastically reduced as a result of the way the system has been designed into independent clusters. Messages are therefore available at any time, regardless of the status of third-party systems.

ADVANTAGES OF APACHE KAFKA

- Combination of versatility and simplicity
- Extreme horizontal scaling allows focused usage in modern cluster systems
- Low latency
- High throughput
- Possibility to process real-time streams
- Protection through SSL and authentication

Apache Kafka is capable of processing high volumes of data for applications in connection with big data, microservices and IoT. It can distribute several hundreds of thousands of messages per second independent of the original data format.

AERIS EXPERTISE – APPLICATIONS & ARCHITECTURES

- Data exchange transport in microservice-based IoT applications
- Aggregation of logs and KPIs from an IT system
- Stream processing and machine learning
- Website performance tracking

WHY APACHE KAFKA IS SO VALUABLE TO AERIS' CLIENTS

- Rapid deployment and accurate configuration of Apache Kafka clusters in cloud or on-premise environments
- Effective support of clients for projects, in which monolithic applications are migrated to a microservice architecture based on Apache Kafka
- Swift synchronization of data in distributed systems
- Development of Kafka event-based solutions both as batch or real-time solutions
- Monitoring of live clusters
- Knowledge transfer to the development and operational teams of our clients