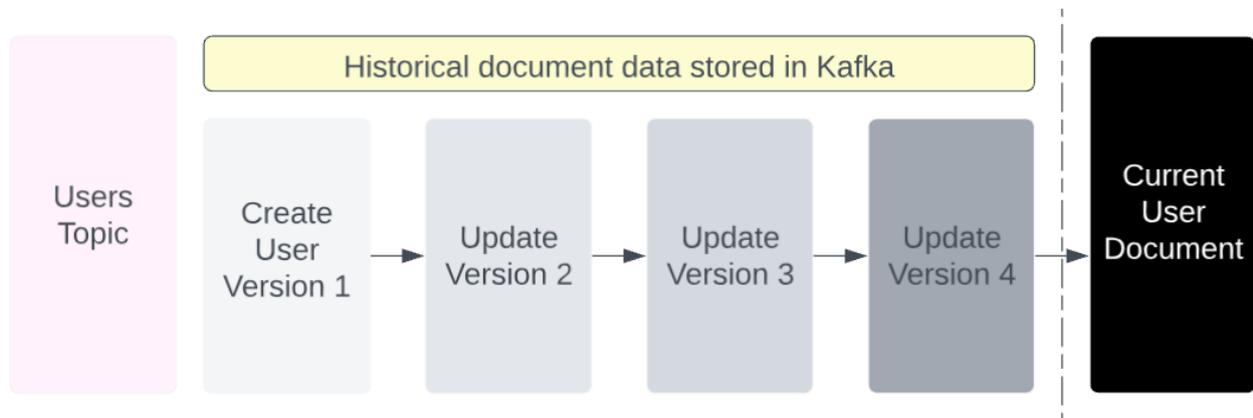


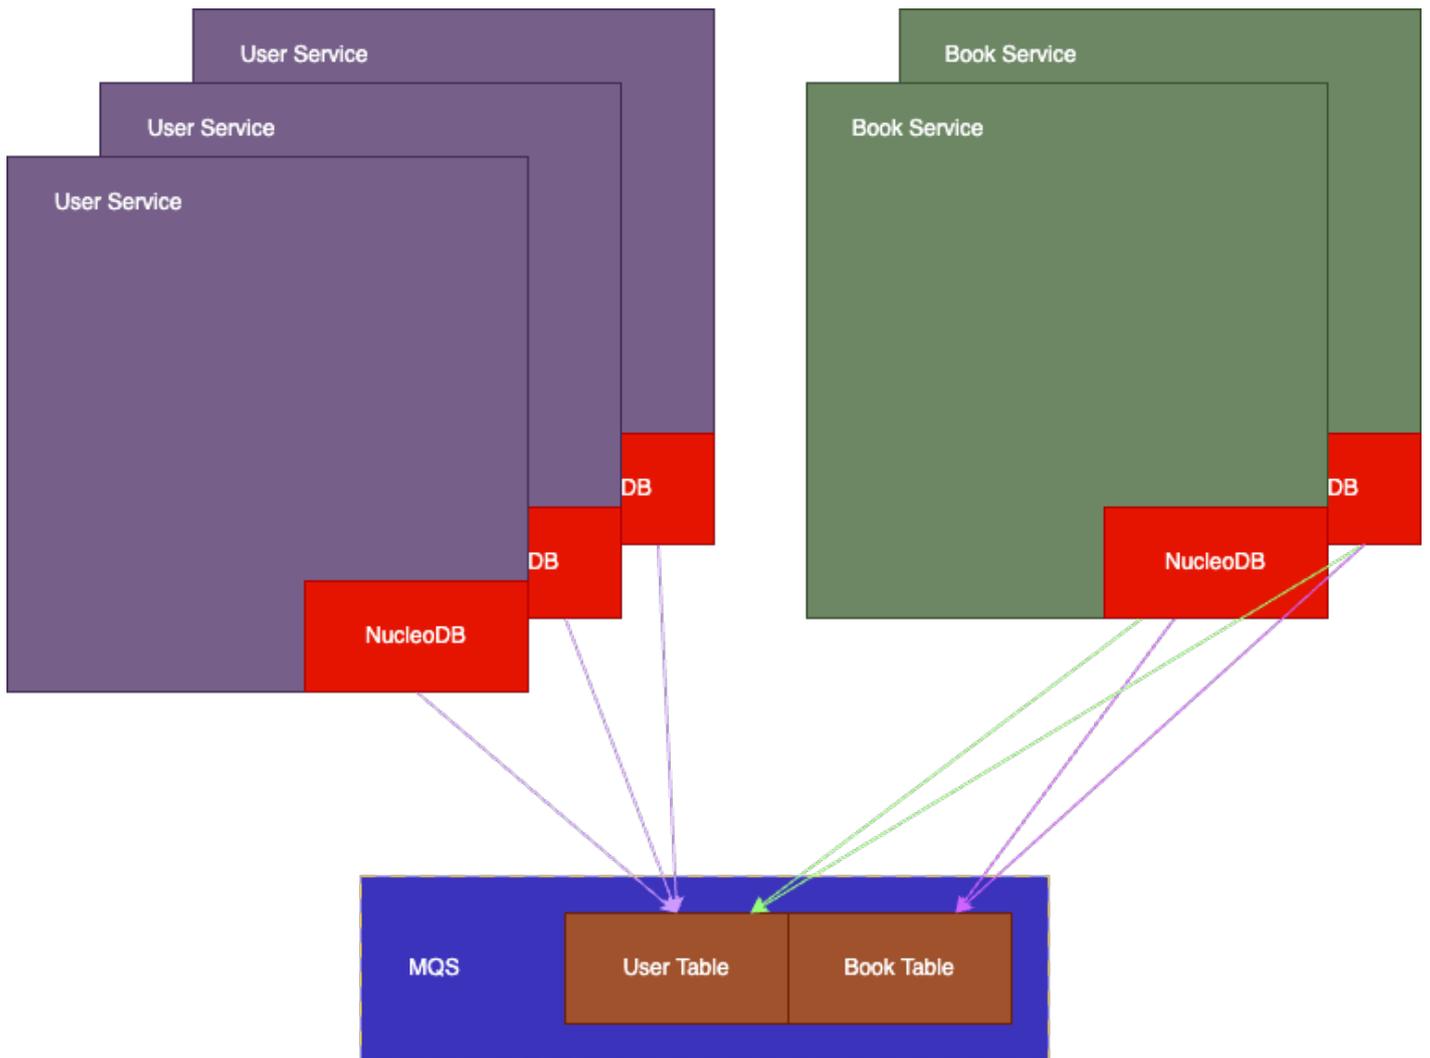
Description



NucleoDB is an in-memory, embedded database system designed to provide high-speed data management and processing. Its in-memory architecture ensures rapid access and manipulation of data, making it ideal for applications where performance and quick response times are critical. Being embedded, NucleoDB seamlessly integrates into various applications, offering a streamlined and efficient way to handle data within the application's own environment. This design approach not only simplifies the application architecture but also enhances overall system performance by reducing the need for external data calls.

One of the key features of NucleoDB is its interaction with Kafka, a popular distributed streaming platform. NucleoDB connects directly to a Kafka cluster, with each table within NucleoDB corresponding to a separate Kafka topic. This setup facilitates efficient data streaming and synchronization between NucleoDB and Kafka, enabling real-time data processing and analytics. Additionally, NucleoDB includes a specialized 'Connections' table, reserved for linking two entries from different tables together, along with associated metadata. This feature allows for the creation of complex relationships and data structures within the database, enhancing its capability to handle diverse and intricate data models. The integration with Kafka, combined with its in-memory and embedded qualities, positions NucleoDB as a powerful tool for modern, data-driven applications.

Architecture



Revision #4

Created 23 December 2023 18:53:03 by Nathaniel

Updated 23 December 2023 22:22:00 by Nathaniel