




THE TALES OF EVENT-DRIVEN TOLD BY A SPRING DEVELOPER

September 2020 / SpringOne



@gamussa | #springone | @confluentinc

PREFACE











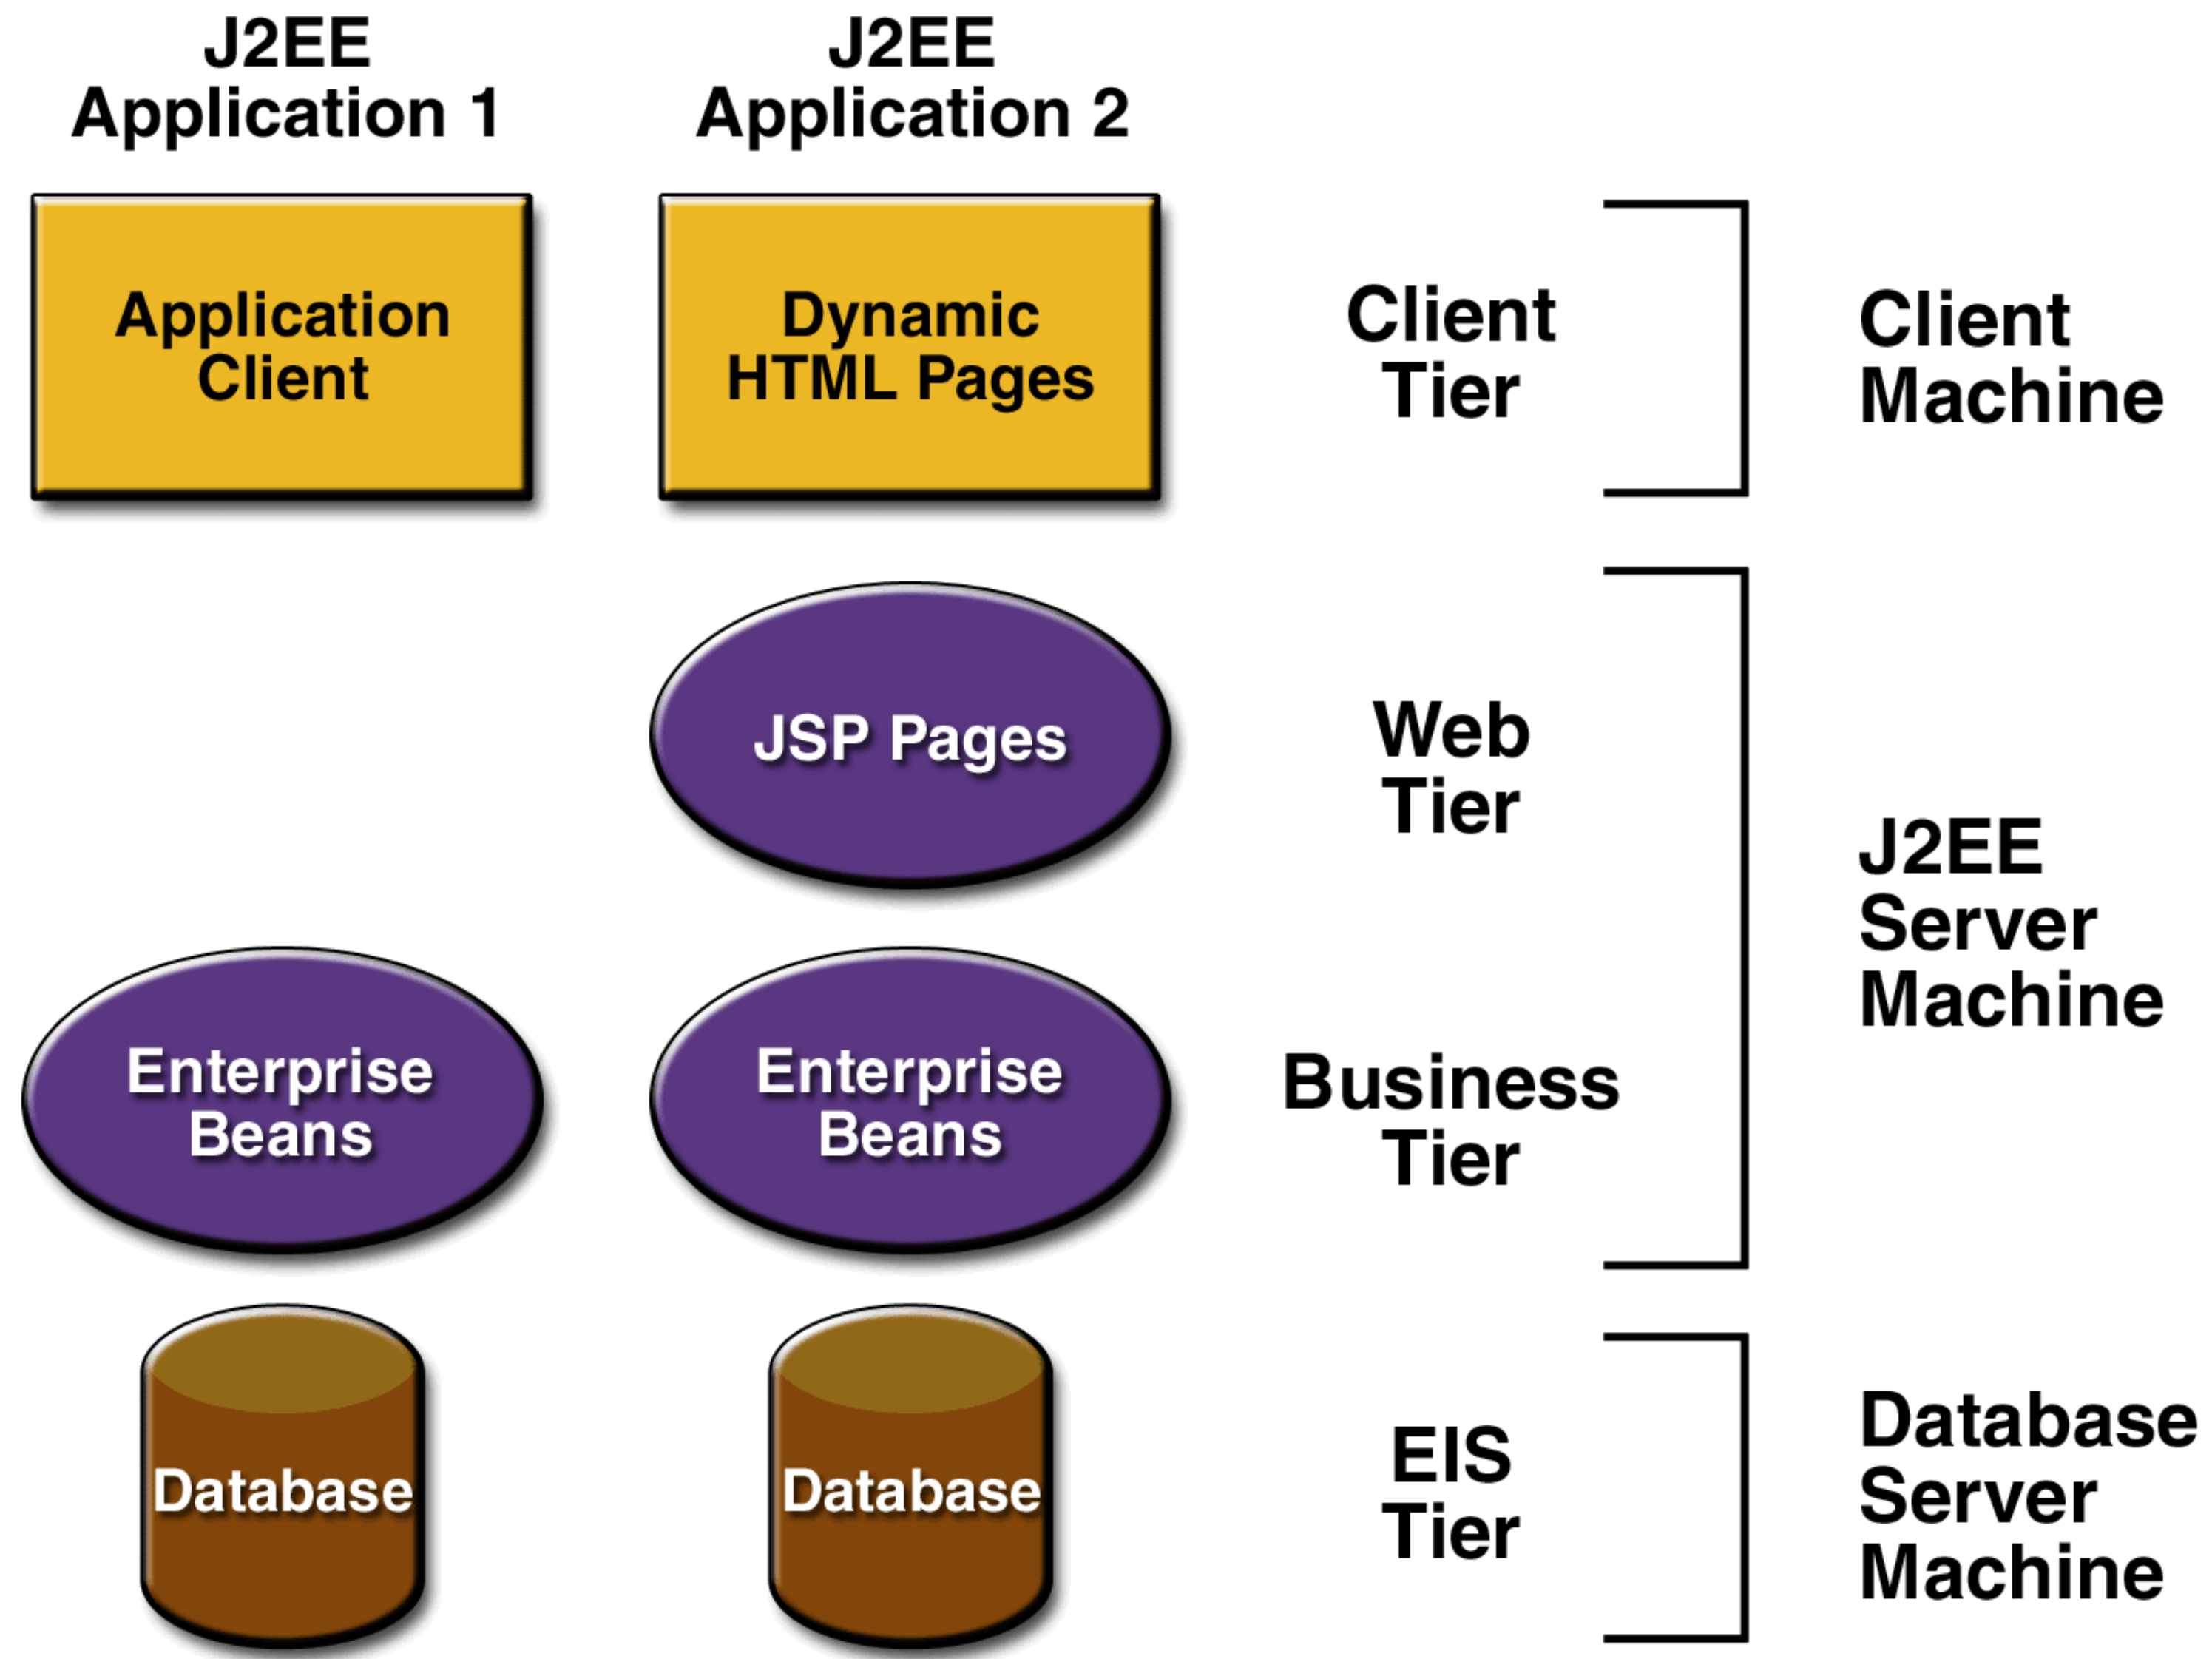





REFACTORING TO MICROSERVICES

IT USED TO BE SO SIMPLE...





NORMAL APPLICATIONS

(I.E., MONOLITHS)

***MONOLITHS
ARE HARD TO THINK ABOUT***

***MONOLITHS ARE HARD TO
CHANGE***



RE-INTEGRATION





***THERE ARE NO GOOD WAYS
TO INTEGRATE
MICROSERVICES***



***THERE ARE NO-GOOD WAYS
TO INTEGRATE
MICROSERVICES***

FILESYSTEM



DATABASE

MURIATIC ACID 20°
UN 1789

WARNING-CONTENTS DANGEROUS-CAUSES SEVERE BURNS


1. When Handling: Wear Goggles or Face Shield.
2. Do Not Get in Eyes or on Skin or Clothing. Do Not Take Internally.
3. In Case of Accidental Contact: Immediately Flush Thoroughly with Water.
4. When Diluting: DO NOT ADD WATER TO CONTENTS.
5. Add contents SLOWLY to Water.
6. Before Moving Containers: Be Sure Caps are Sealed & Tight.
7. Loosen Capsules Carefully.
8. Keep Out of Sun and Away from Heat.
9. Never Use Pressure to Empty.
10. In Case of Spillage: Flush Thoroughly with Water.
11. Completely Drain Containers Before Recycling.
12. DANGER: Vapor Extremely Irritating. Avoid Breathing of Vapor.

VAPOR HARMFUL

Milport Chemical Company
P.O. BOX 1000, NEW YORK, N.Y. 10001

Nope. Don't like that.

INTEGRATING MICROSERVICES THROUGH THE DATABASE




- "I have a database and I know how to use it."
- Eventually causes services to co-mingle.
- Violation the «bounded context»
- Great to use **inside** a service boundary!
- Terrible for **sharing data** or **negotiating change**.


RPC



INTEGRATING MICROSERVICES VIA RPC



- Avoids problems of database integration
- Feels natural
- Aligns with the request/response paradigm
- Problem: cascading failures
- Question: how do you **debug** this system?
- Answer: *you build a l*



DemoApplicationTests > contextLoads() PASSED

TransactionTransformerTest > shouldHaveEnoughFunds() PASSED

TransactionTransformerTest > shouldHaveInsufficientFunds() PASSED

TransactionTransformerTest > shouldStoreTransaction() PASSED

KStreamConfigTest > balanceShouldBe300() PASSED

KStreamConfigTest > shouldCreateSuccessfulTransaction() PASSED

KStreamConfigTest > testDriverShouldNotBeNull() PASSED

KStreamConfigTest > shouldHaveInsufficientFunds() PASSED

2020-09-01 23:27:30 org.apache.kafka.clients.consumer.KafkaConsumer [INFO] (KafkaConsumer.java:1082) - [Consumer

2020-09-01 23:27:30 o.a.kafka.streams.processor.internals.StreamThread [INFO] (StreamThread.java:221) - stream-

2020-09-01 23:27:30 o.a.kafka.streams.processor.internals.StreamThread [INFO] (StreamThread.java:1130) - stream-

2020-09-01 23:27:30 org.apache.kafka.clients.producer.KafkaProducer [INFO] (KafkaProducer.java:1182) - [Produce

2020-09-01 23:27:30 o.a.kafka.streams.processor.internals.StreamThread [INFO] (StreamThread.java:221) - stream-

2020-09-01 23:27:30 o.a.kafka.streams.processor.internals.StreamThread [INFO] (StreamThread.java:1130) - stream-

2020-09-01 23:27:30 o.a.kafka.clients.admin.internals.AdminMetadataManager [INFO] (AdminMetadataManager.java:235) -

org.apache.kafka.common.errors.TimeoutException: Call(callName=fetchMetadata, deadlineMs=159901) timed out waiting to send the call.

2020-09-01 23:27:30 org.apache.kafka.streams.KafkaStreams [INFO] (KafkaStreams.java:207) - stream-client [trans

2020-09-01 23:27:30 org.apache.kafka.streams.KafkaStreams [INFO] (KafkaStreams.java:989) - stream-client [trans

2020-09-01 23:27:30 org.apache.kafka.streams.KafkaStreams [INFO] (KafkaStreams.java:989) - stream-client [trans


2020-09-01 23:27:30 org.apache.kafka.streams.KafkaStreams [INFO] (KafkaStreams.java:989) - stream-client [trans

2020-09-01 23:27:30 org.apache.kafka.streams.KafkaStreams [INFO] (KafkaStreams.java:989) - stream-client [trans

2020-09-01 23:27:30 org.apache.kafka.streams.KafkaStreams [INFO] (KafkaStreams.java:989) - stream-client [trans

2020-09-01 23:27:30 org.apache.kafka.streams.KafkaStreams [INFO] (KafkaStreams.java:989) - stream-client [trans

LOG4J TYPE OF LOG?



EVENTS?


WHAT'S AN EVENT?

***AN OCCURRENCE, ESPECIALLY ONE OF SOME
IMPORTANCE.***

A COMBINATION OF:




NOTIFICATION



STATE TRANSFER



***ALSO,
EVENTS
ARE IMMUTABLE.***



CORE ABSTRACTION


- DB - table[⊗]
- Hadoop -
file
- Kafka - ?

LOG

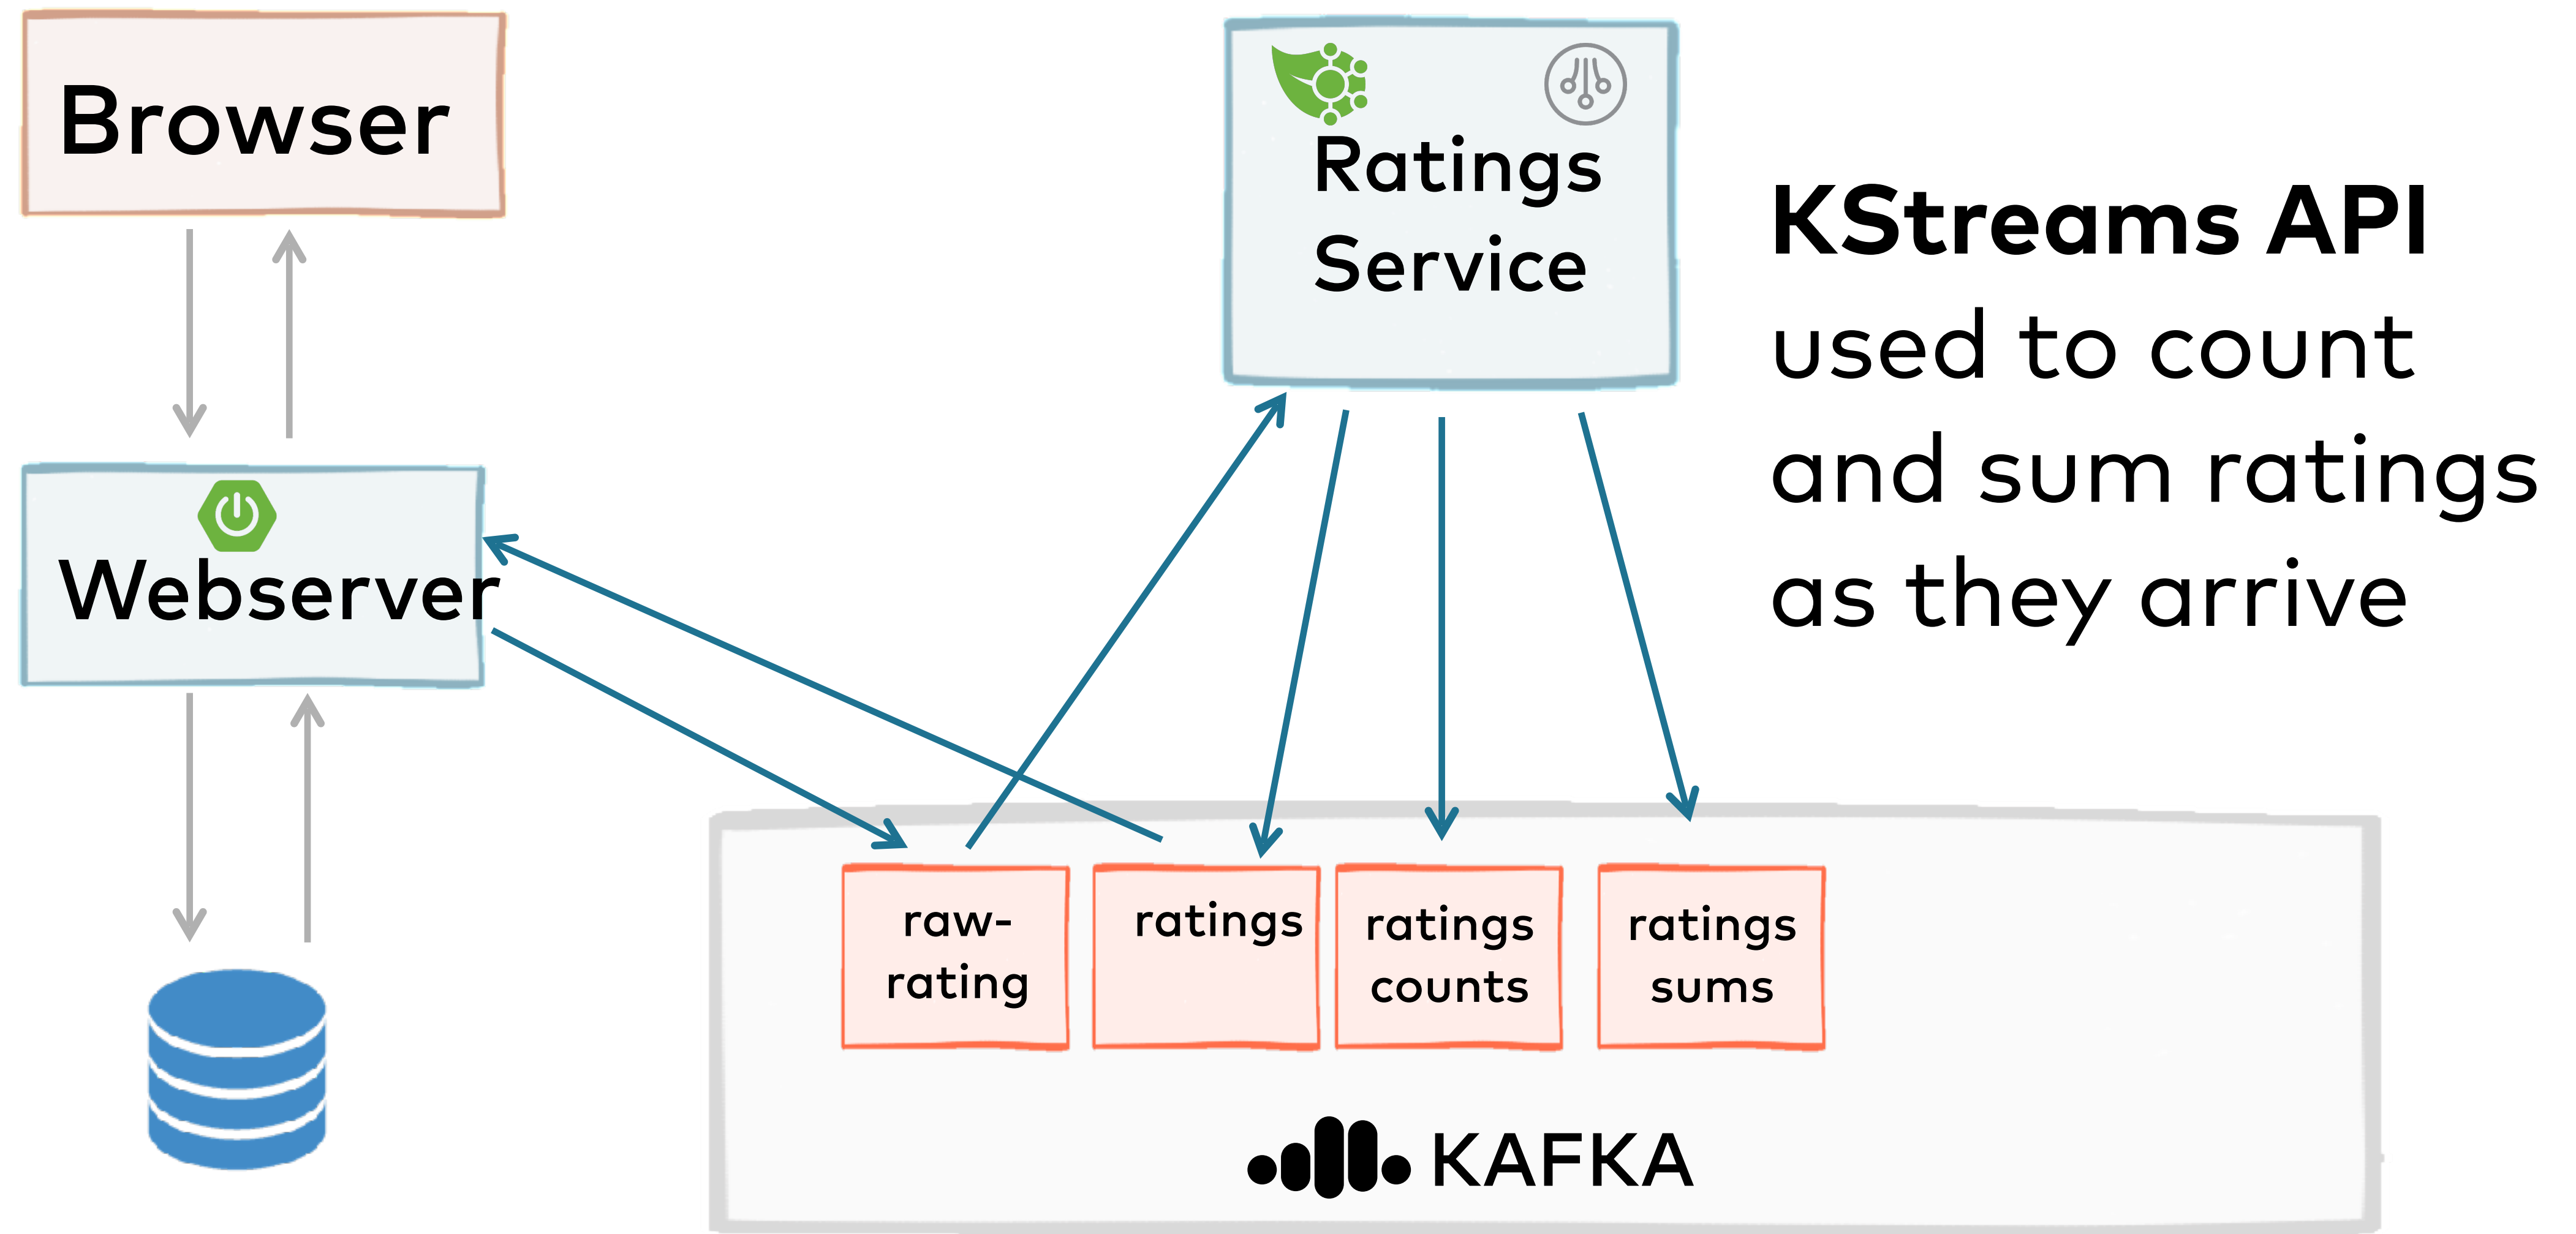



A RICHER

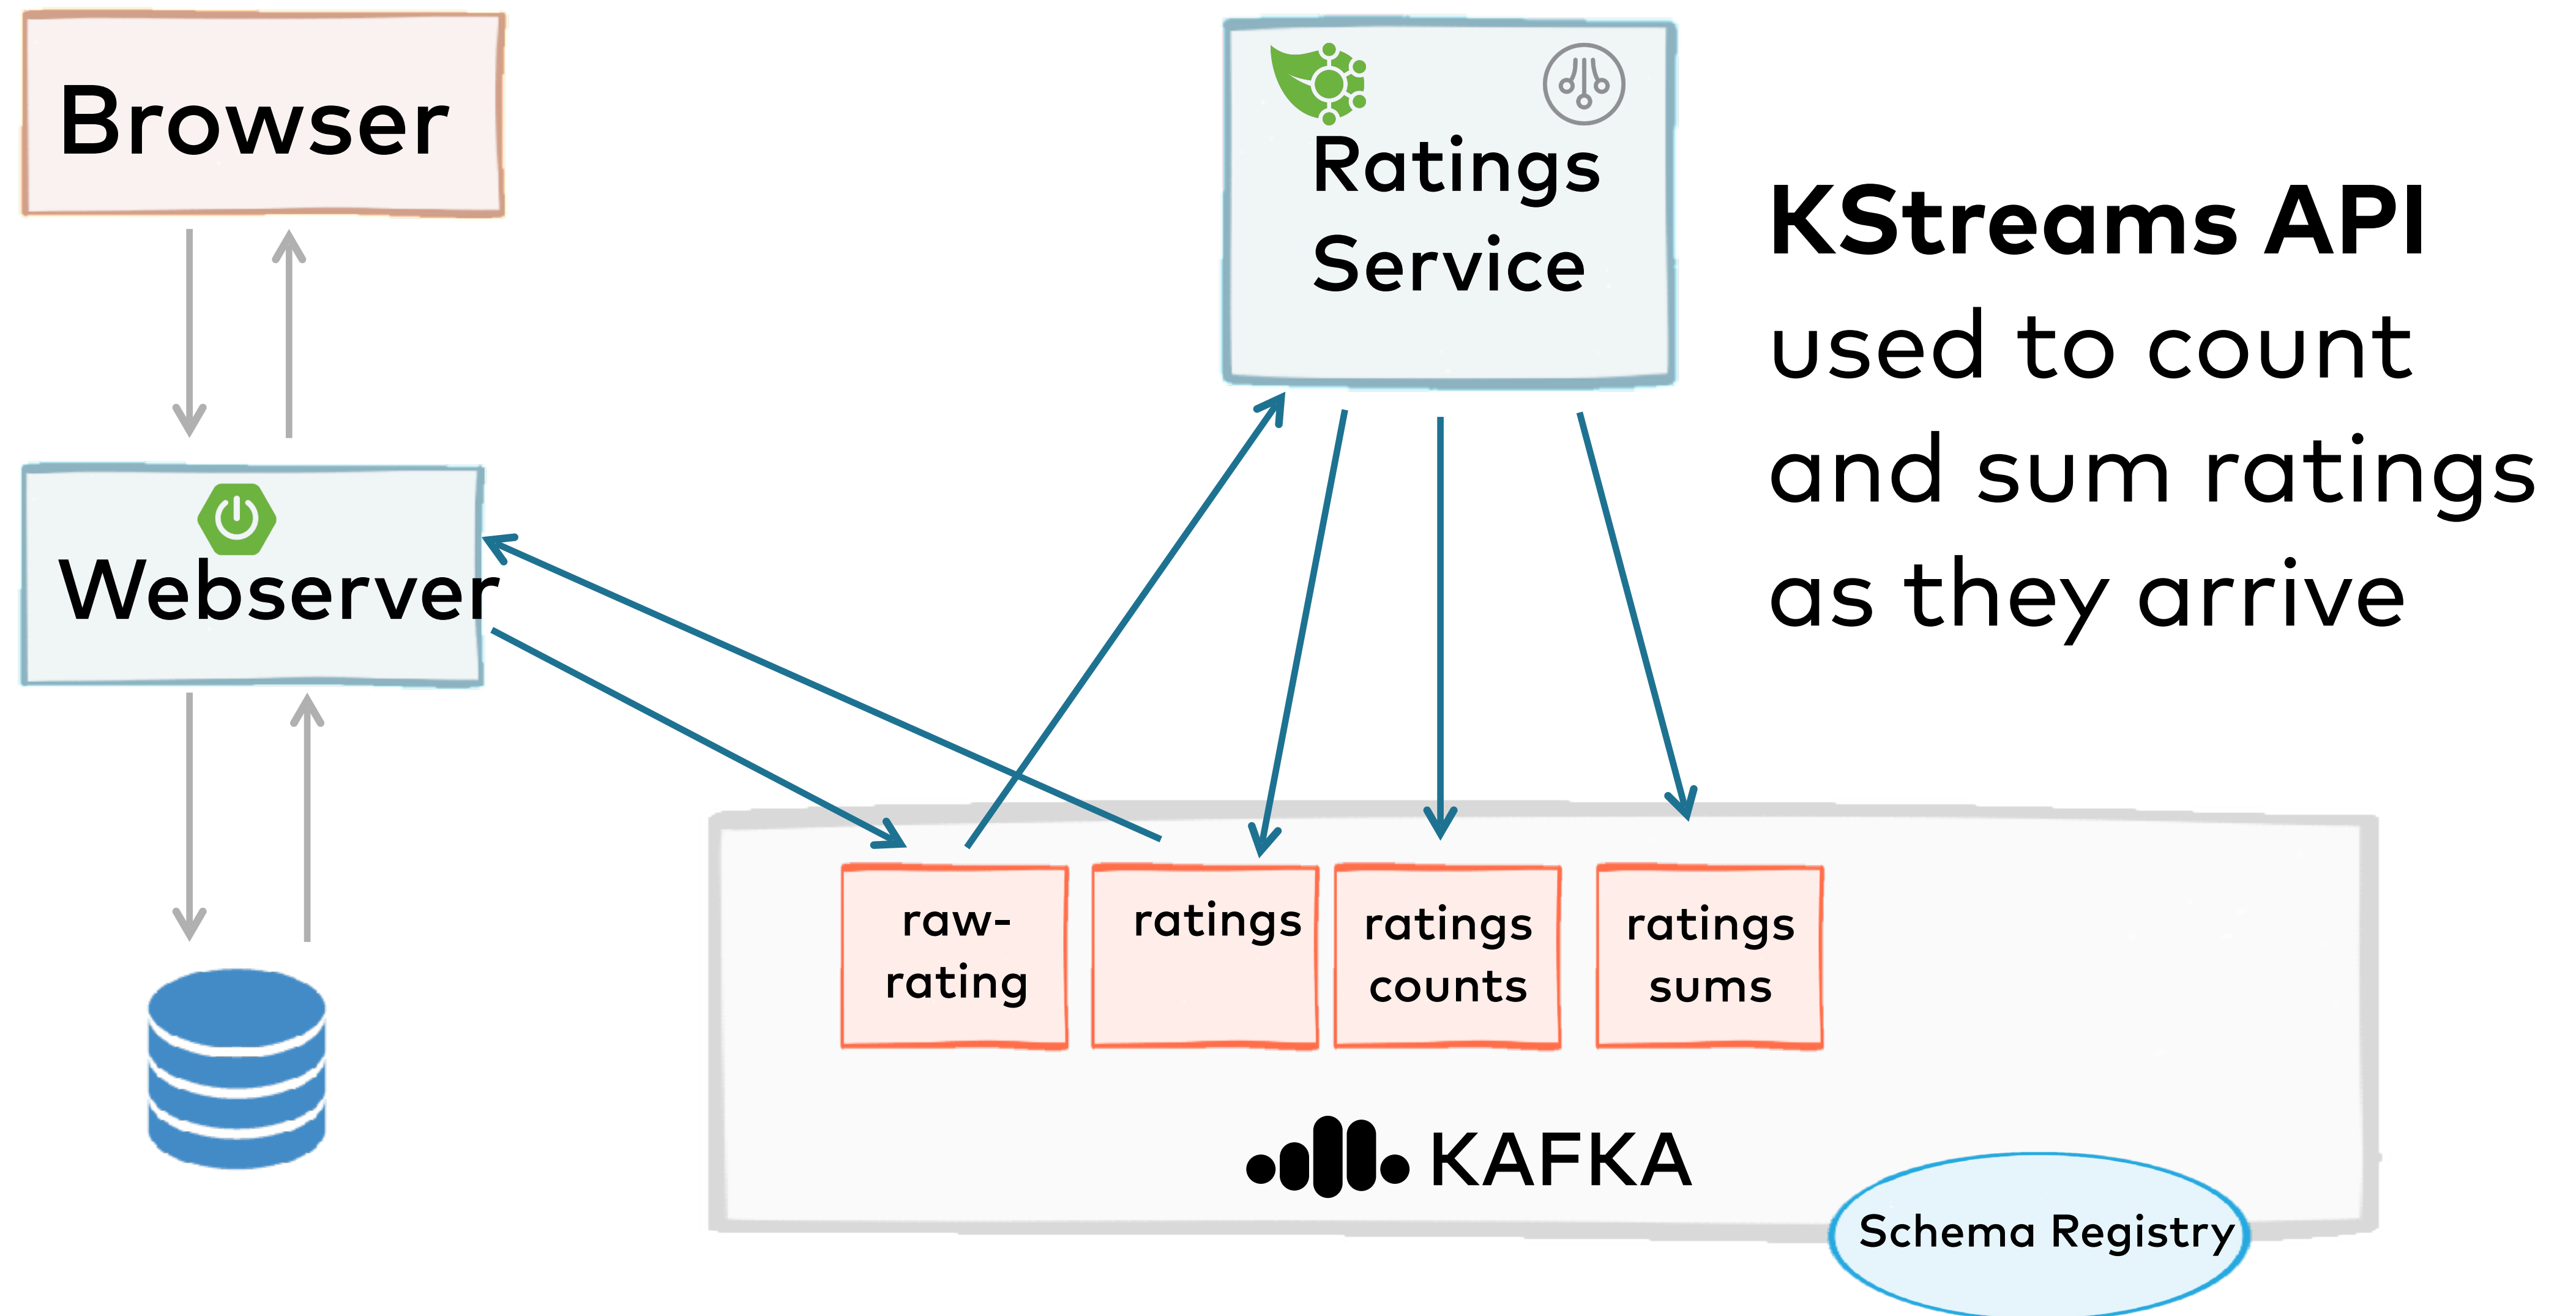

MICROSERVICES APPLICATION



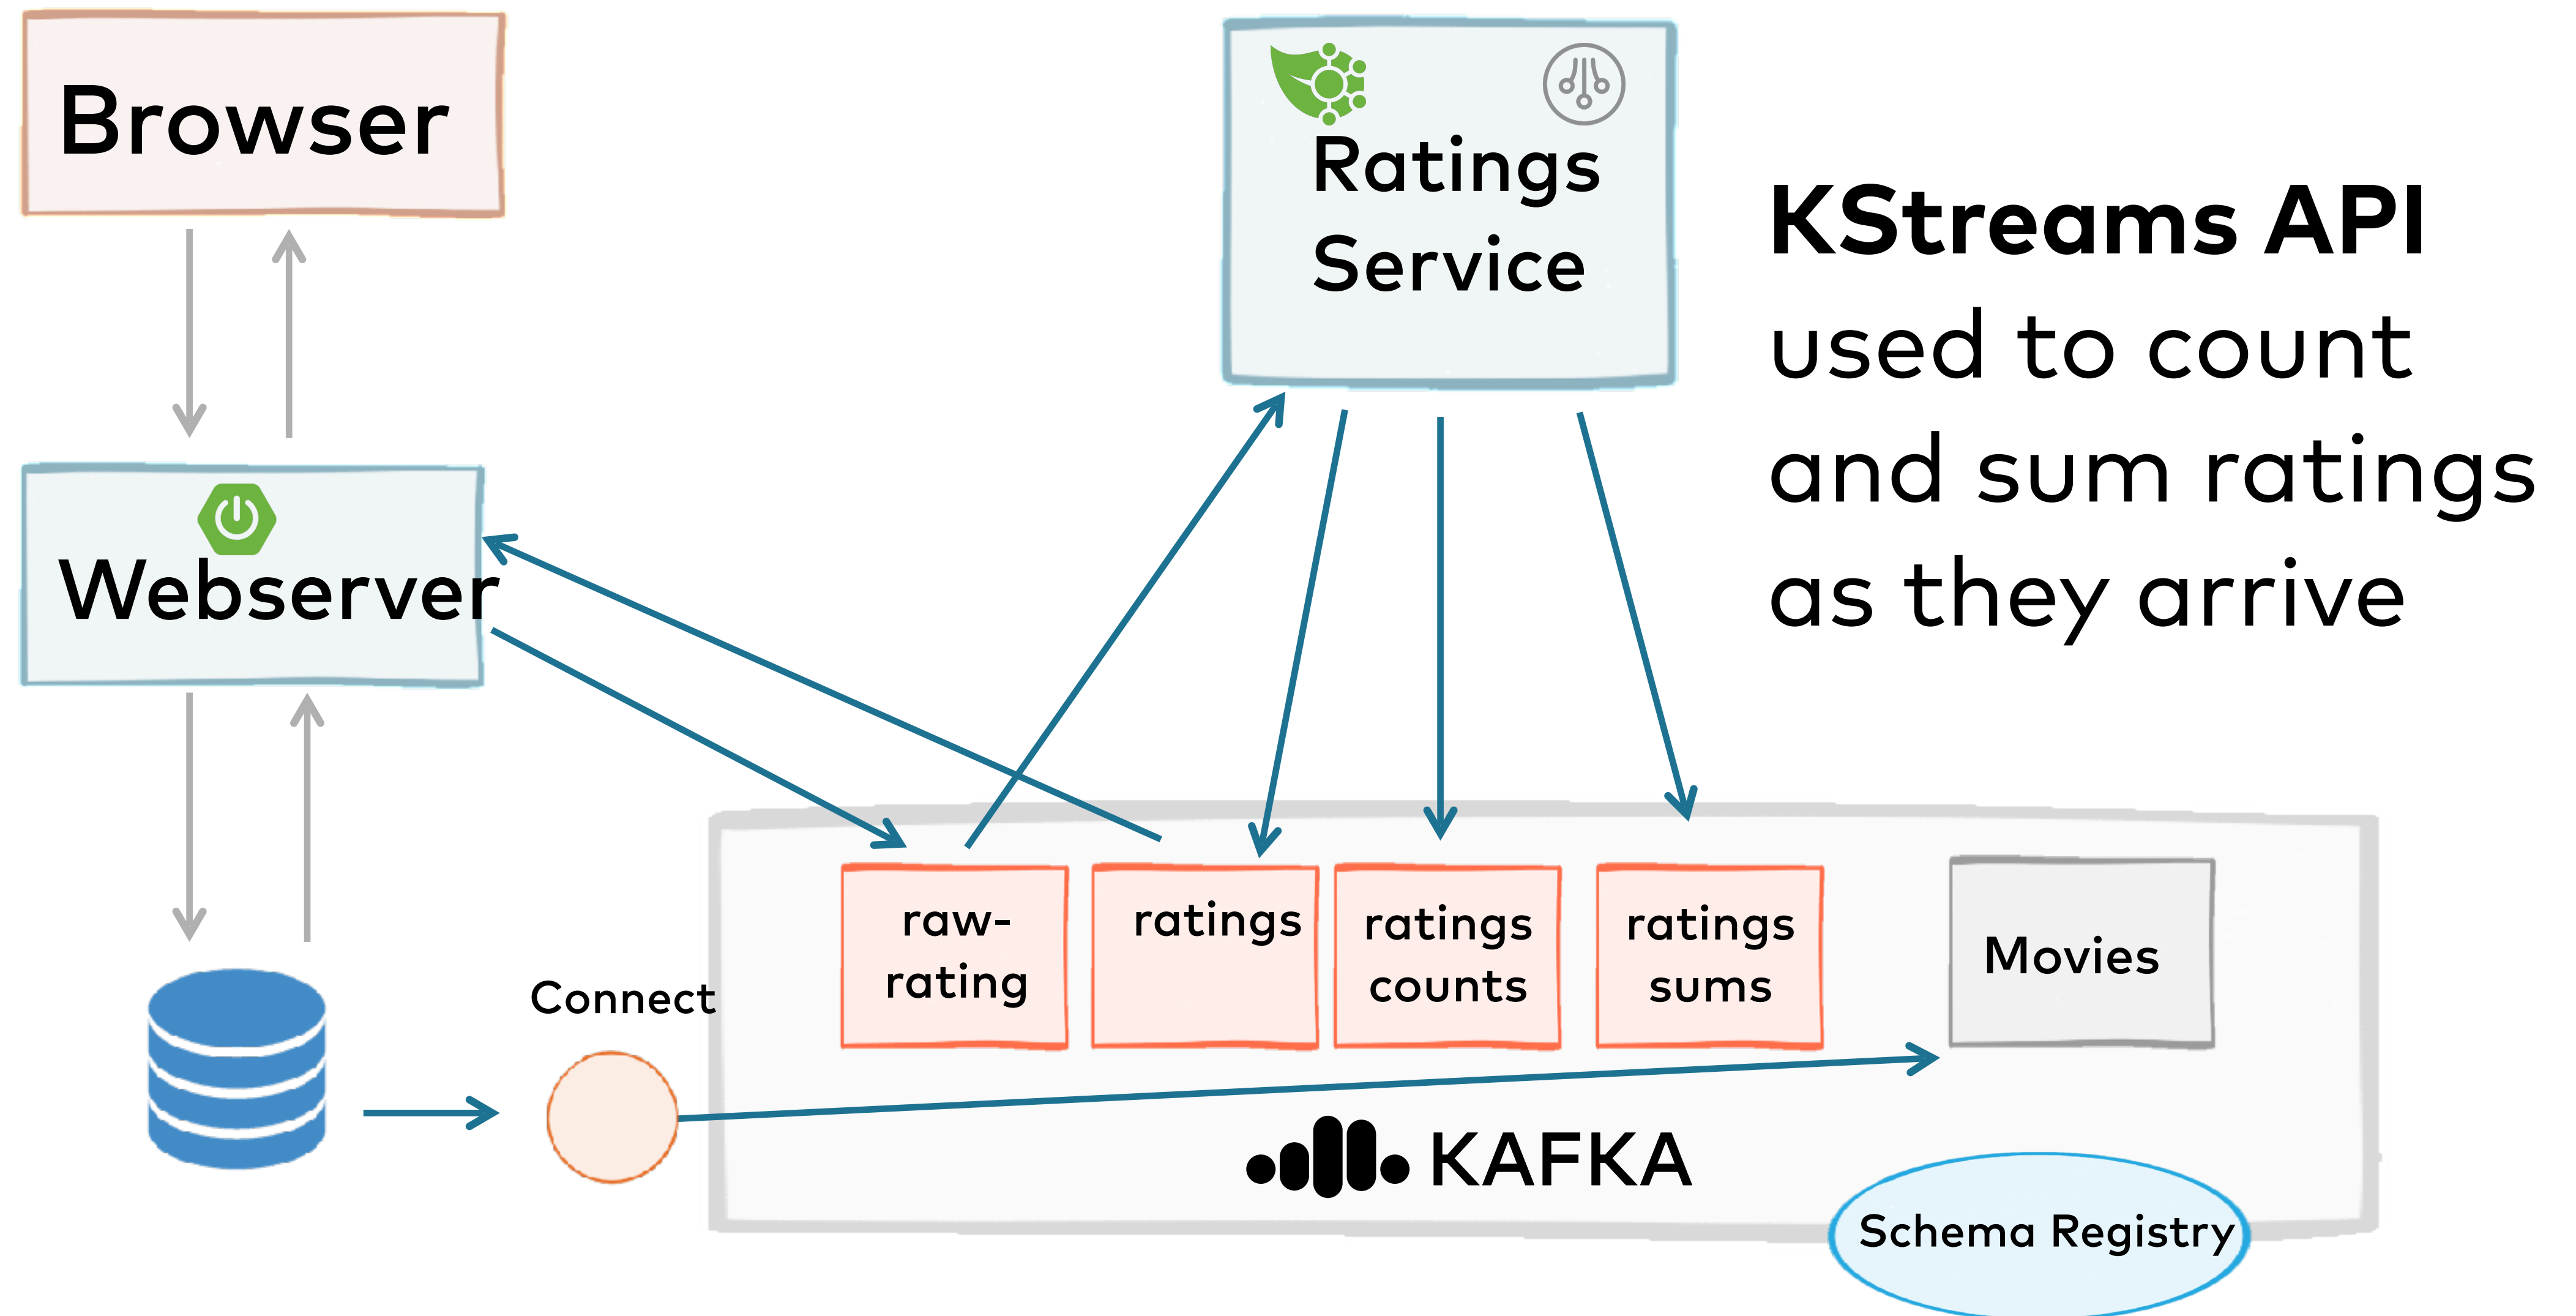

RATINGS SERVICE OPERATING ON THE EVENT STREAM



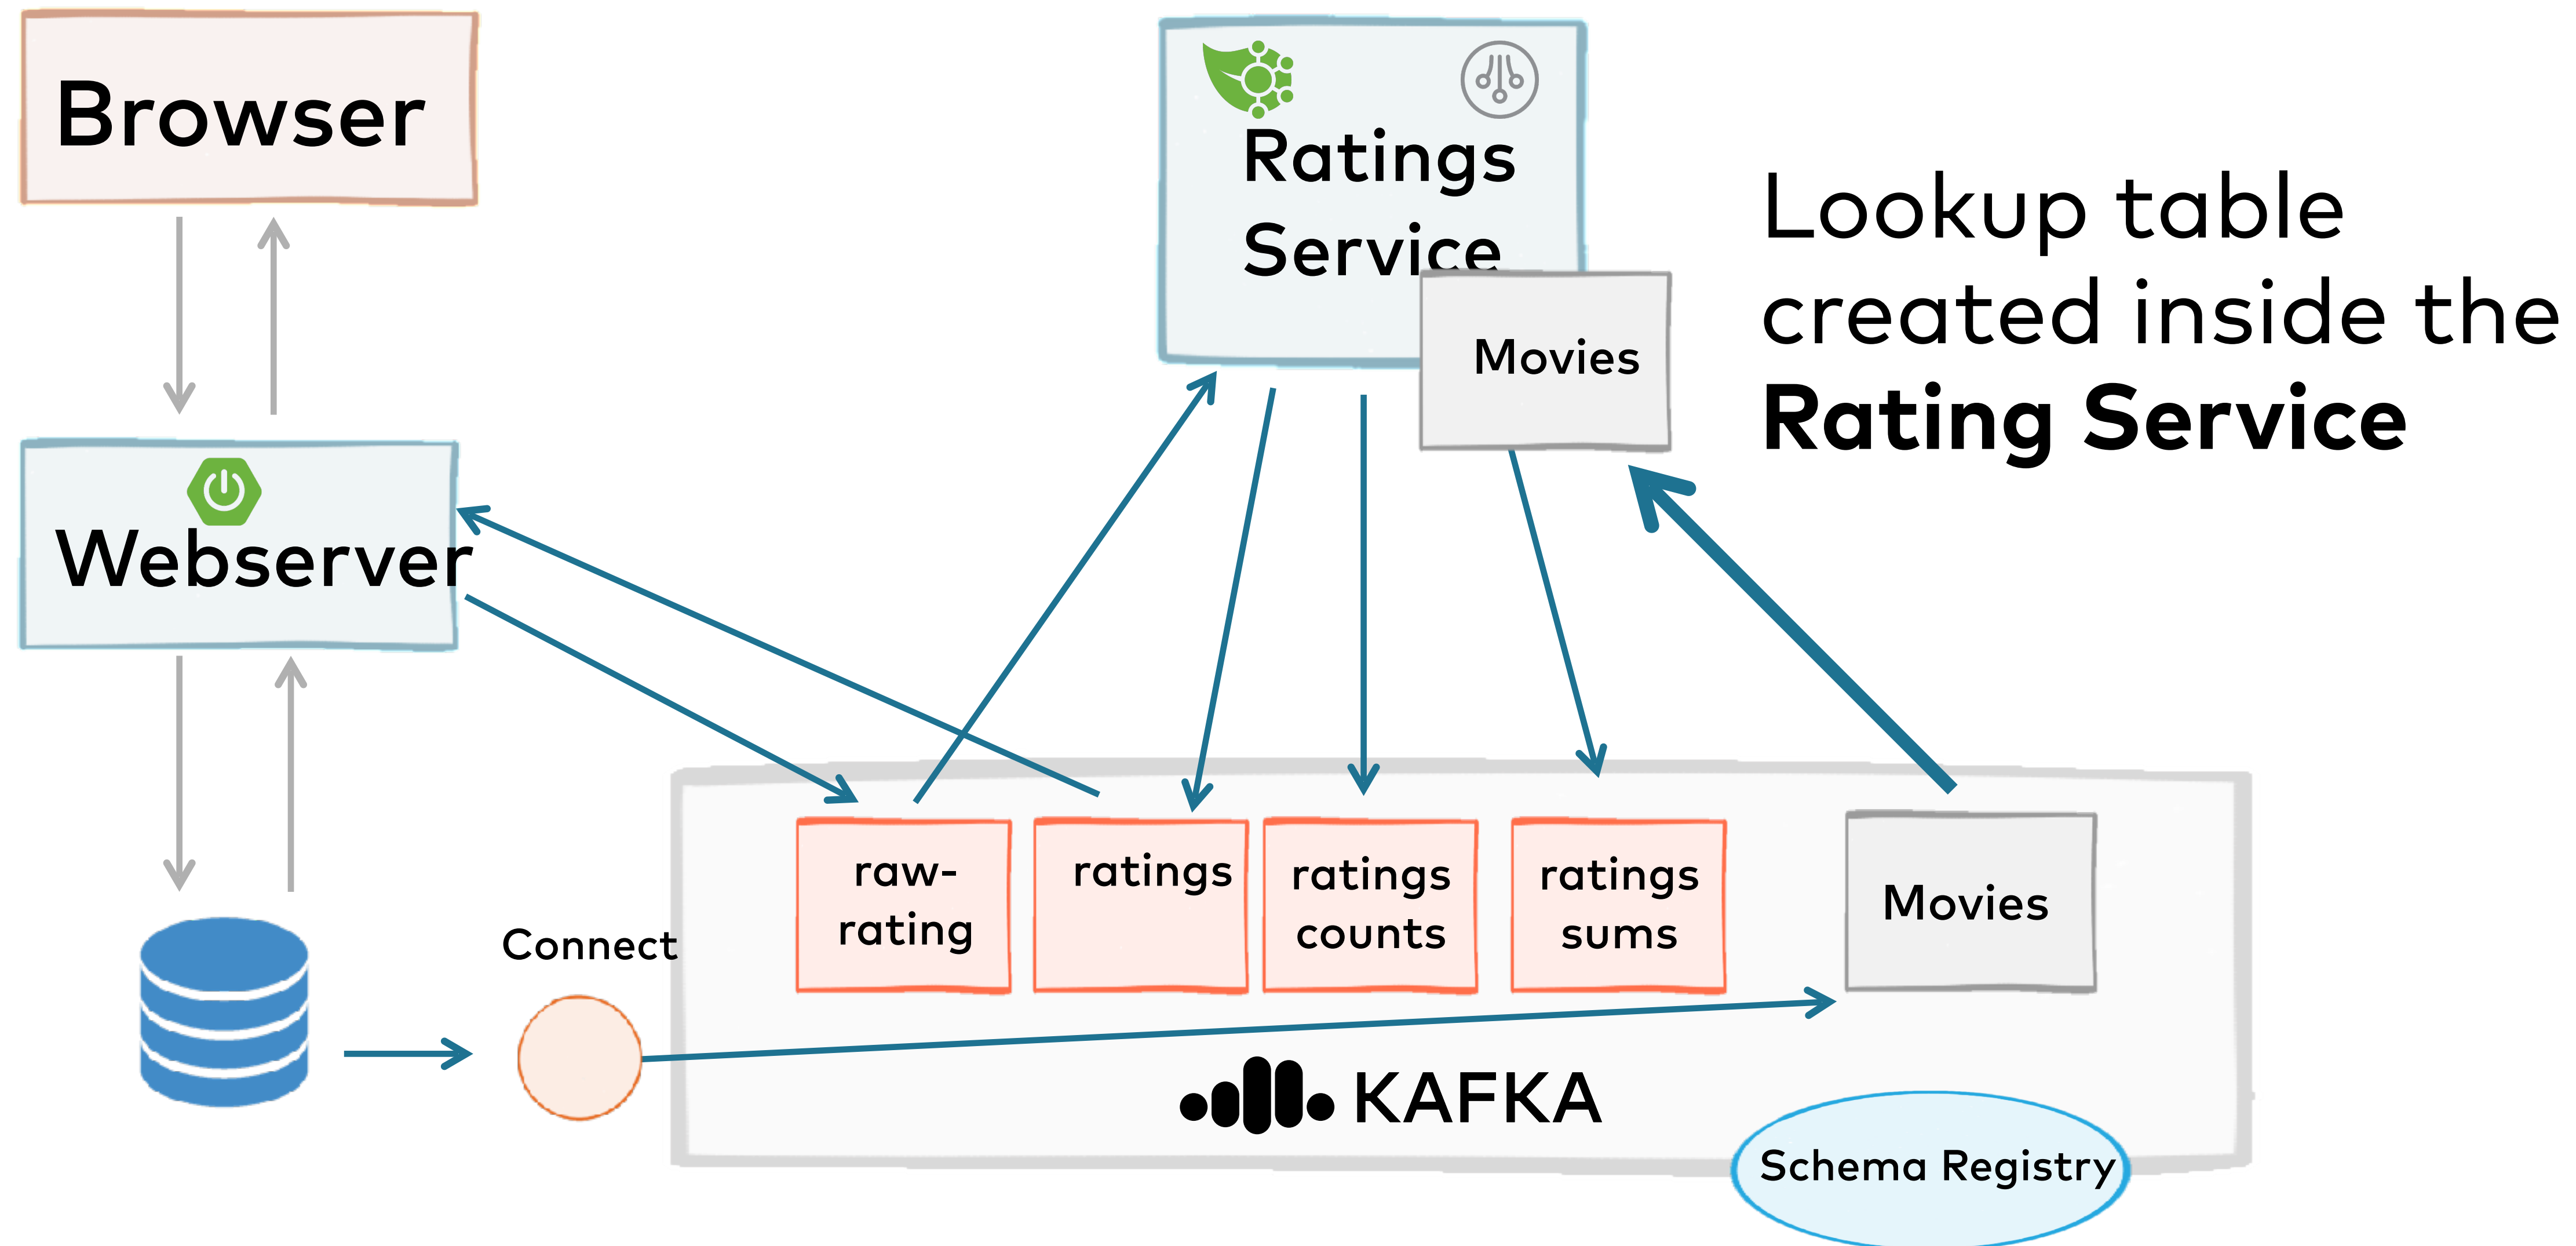

SCHEMA MANAGEMENT



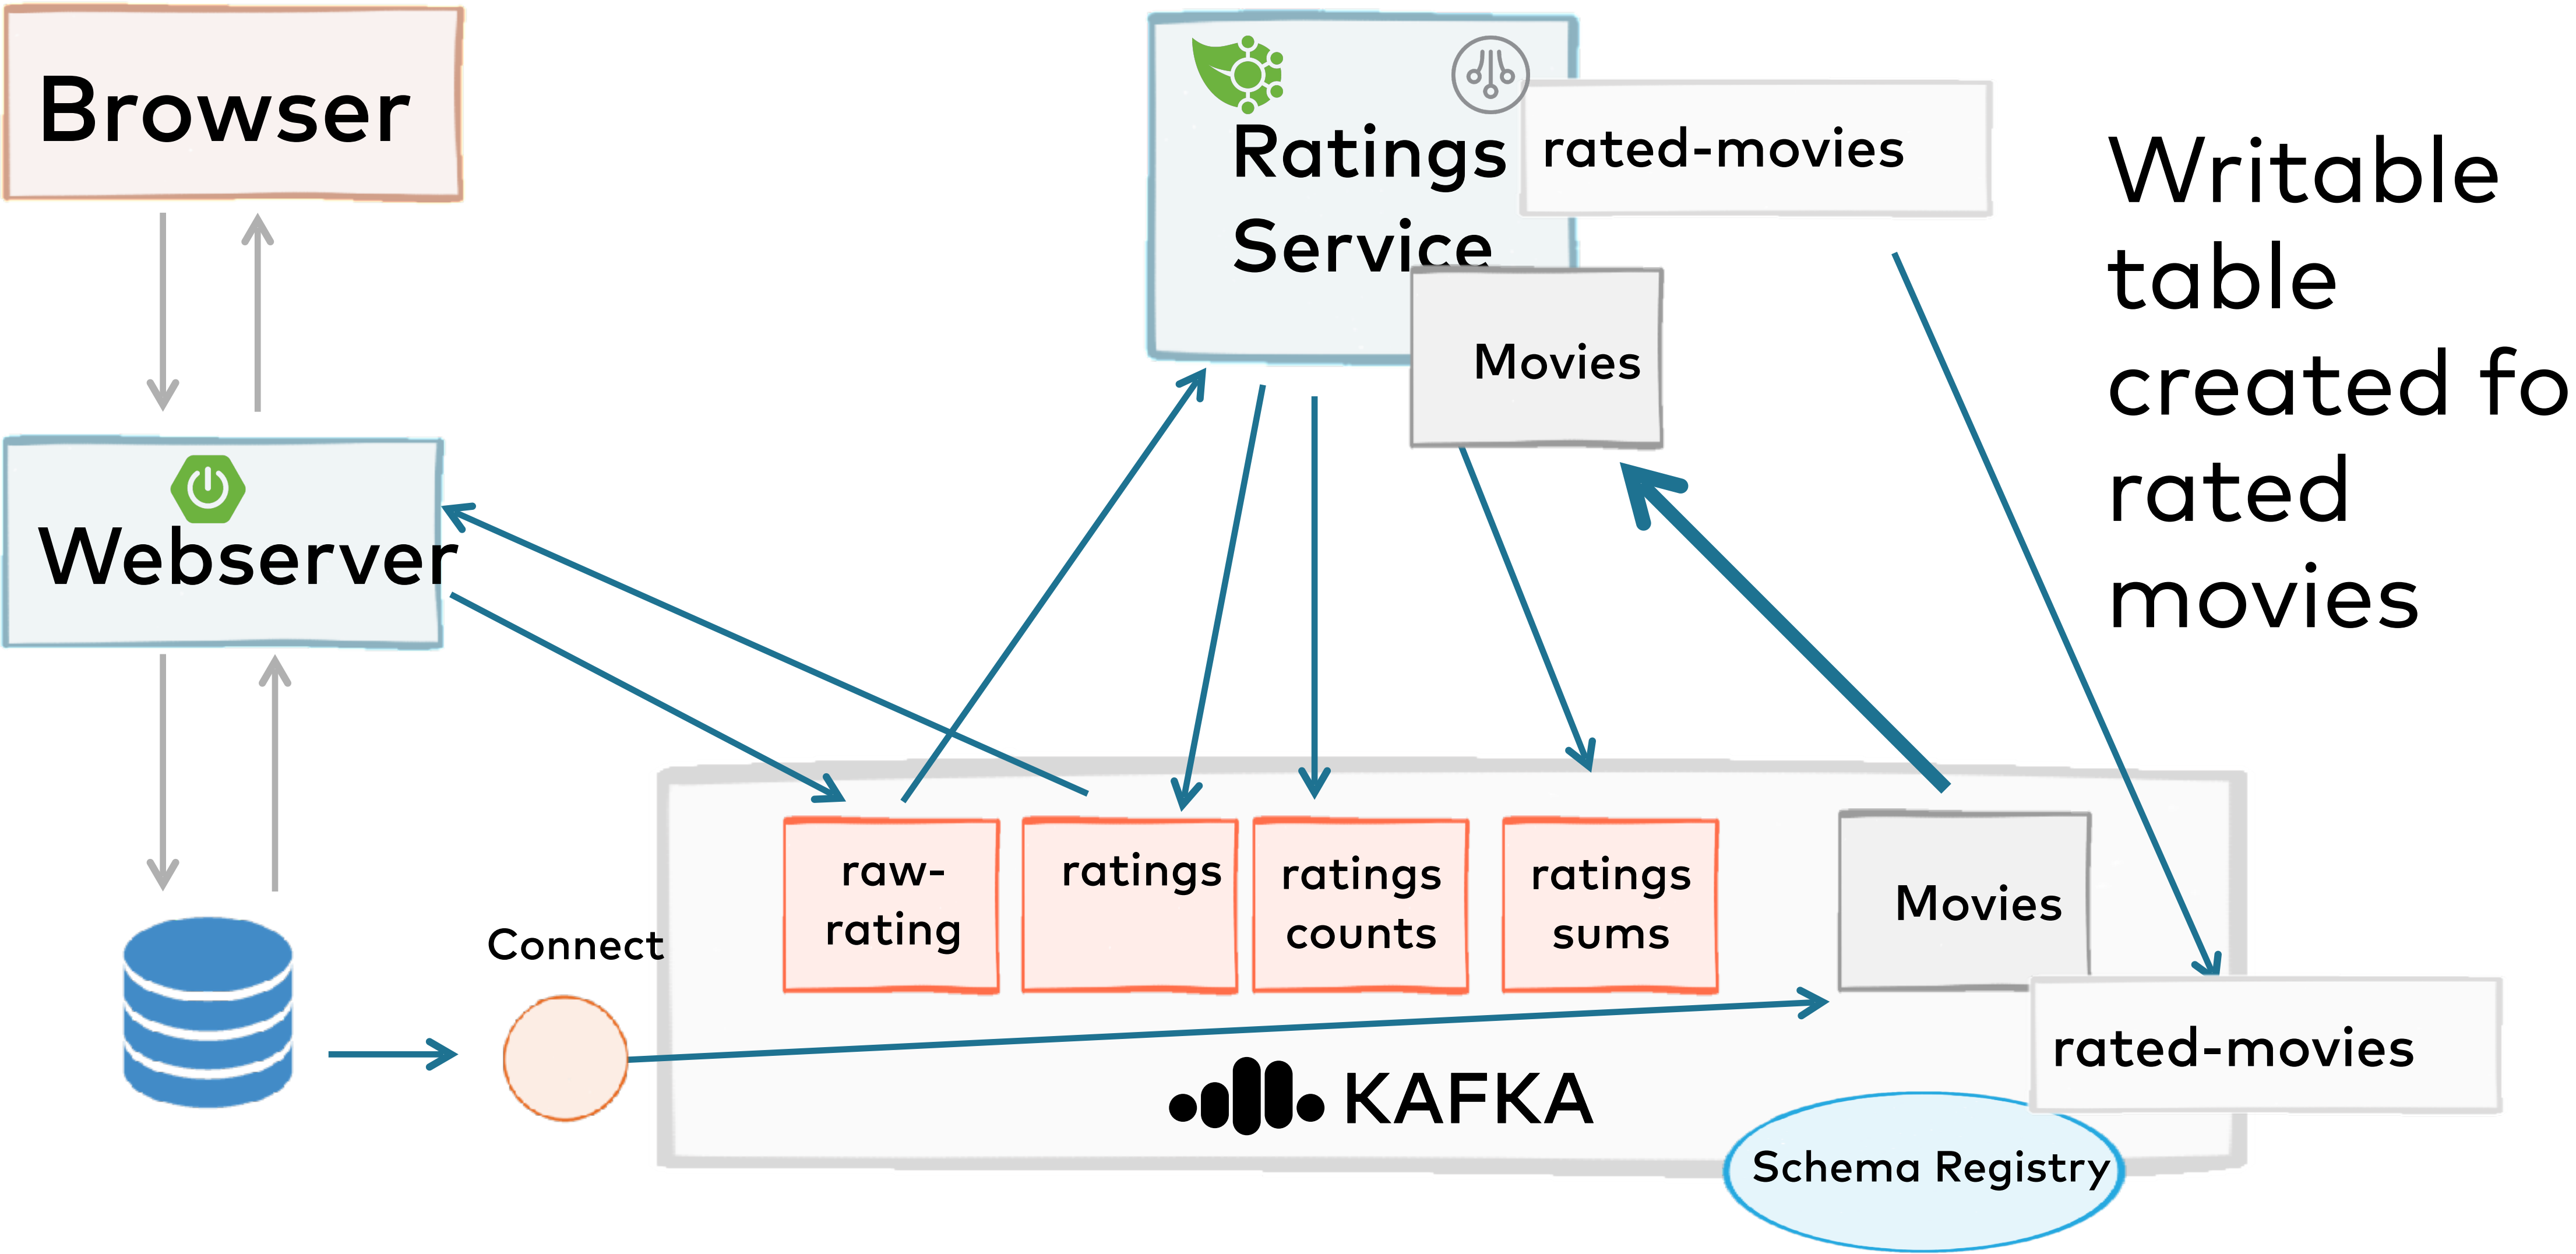

ACCESS LEGACY RELATION DATA



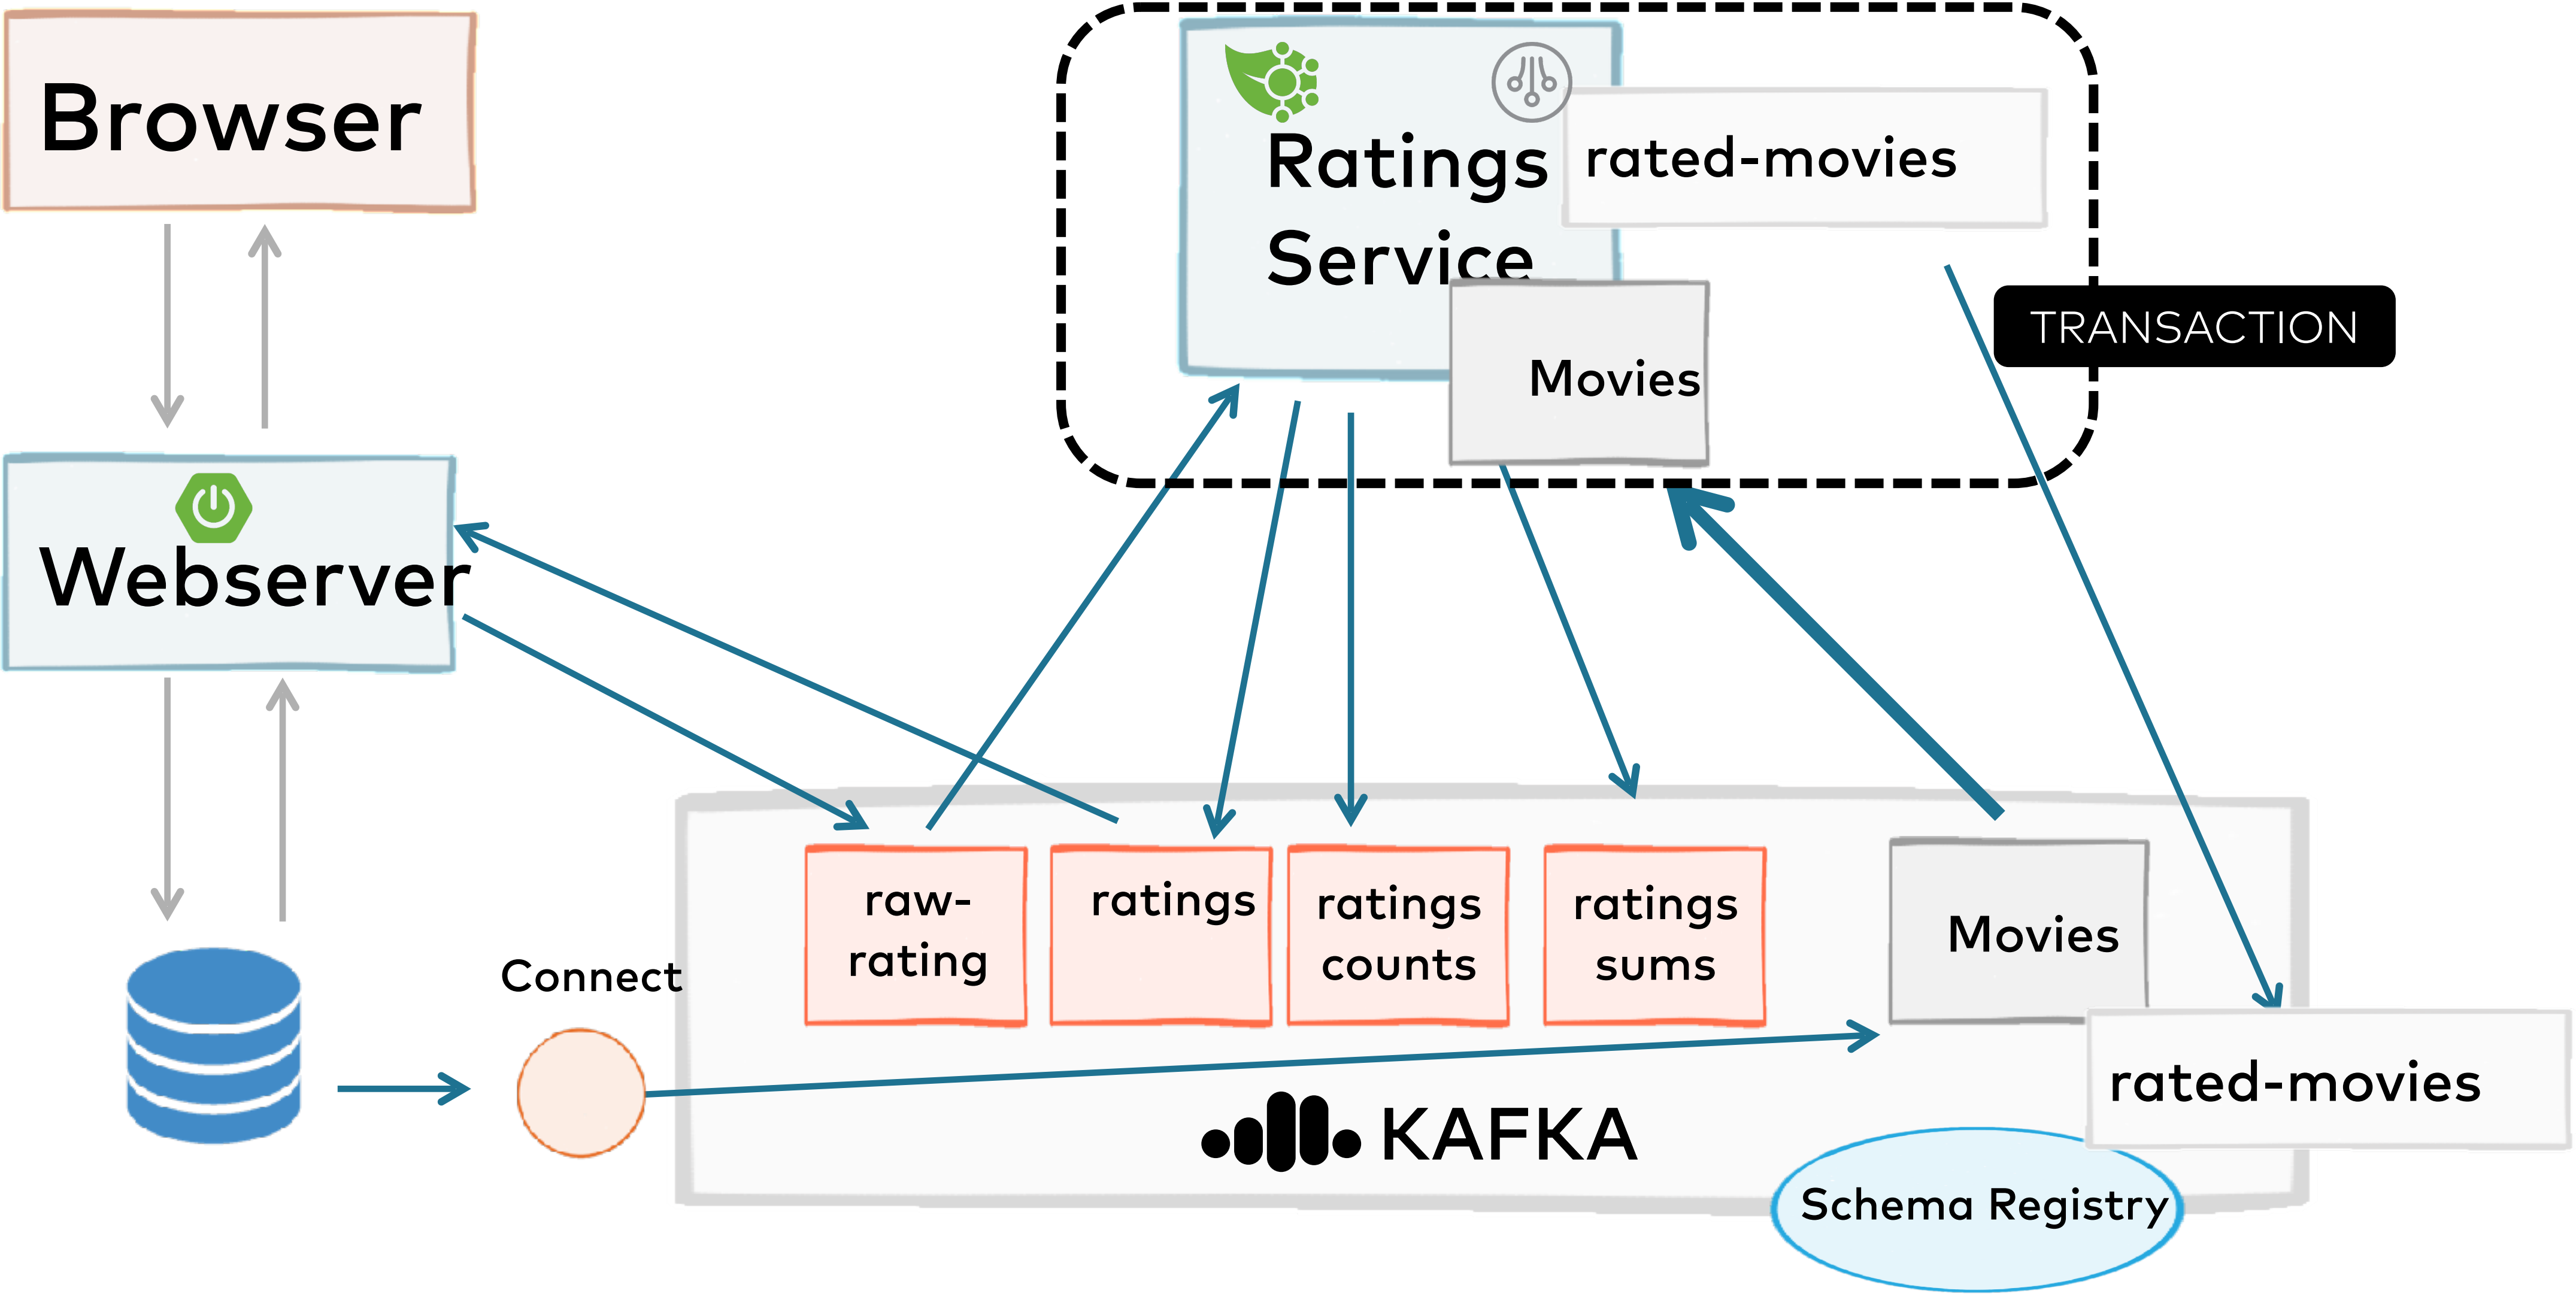

MATERIALIZE TABLES INSIDE THE APP



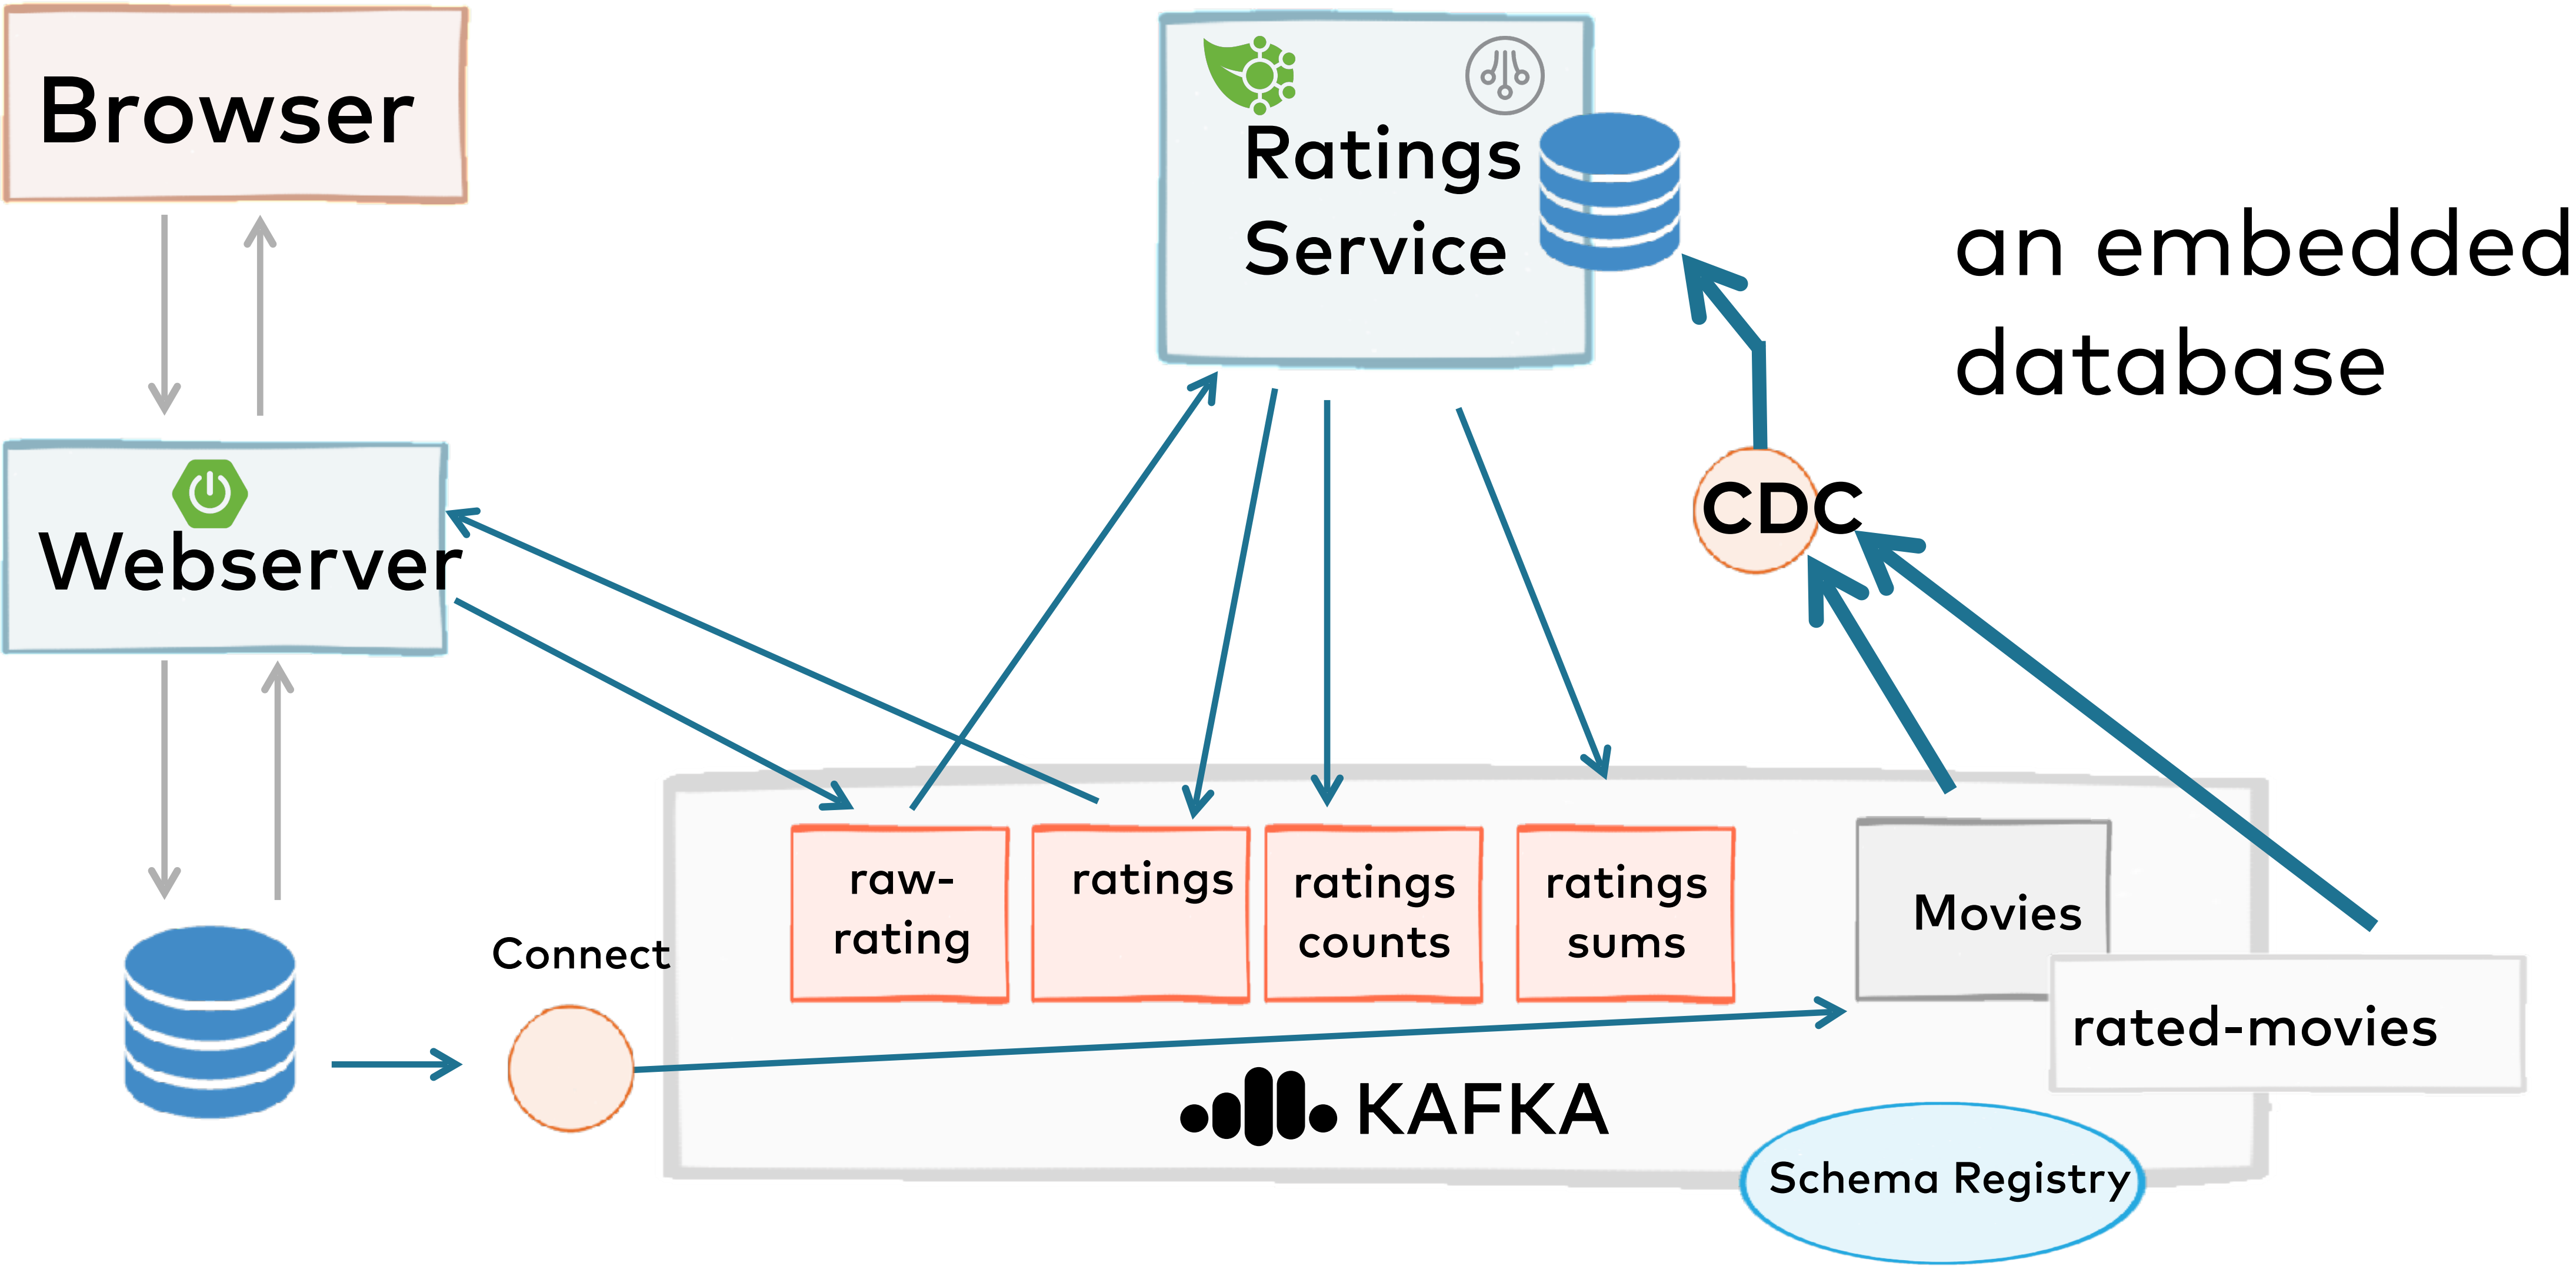

CREATE A NEW TABLE, PERSIST IT TO THE LOG



CREATE A NEW TABLE, PERSIST IT TO THE LOG




HYDRATE A MATERIALIZED VIEW




WHAT DID YOU LEARN TODAY?

- **Streaming services** improve on the traditional microservices pattern in a number of ways:
- Decouple ecosystems for **easier change**
- Evolve existing **databases into event streams** with Kafka Connect.
- Build on **asynchronicity first**. Add synchronicity where needed.



WHAT DID YOU LEARN TODAY?



- Remove the need for circuit breakers or other "overload" patterns.
- Quickly combine and operate on data from many different services.
- Embrace an immutable log that can be replayed.
- Keep many databases, in many independent services, in sync.
- Safely manage the evolution of data with the schema registry.



WANT TO LEARN MORE?

Developing Event-Driven Microservices with Spring Boot, Confluent Cloud, Kotlin, and Java

Wed, Sep 2 at 10:00 AM EDT (90 minutes)

A microservices-based approach is now pretty universally recognized as the right way to develop modern applications. Technologies like Spring Boot and Spring Framework are ideal choices to develop and deploy your microservices (regardless of the language of choice, Java or Kotlin), with Apache Kafka® as the microservices communication layer. This approach provides highly scalable and robust communication between services that will allow you to unlock data in your existing monolithic and/or legacy systems via Kafka to a new generation of applications. These are not just cool things even though they are very cool, they are a real solution to a real problem. You will meet the right architectural decisions with many successful projects behind them.

In this workshop, you will:

- Develop a small, functioning, microservices-based web application using Kafka Streams (using Java or Kotlin).
- Use Spring Boot to handle tedious and not-so-exciting tasks (like application configuration or container images generation).
- Use Confluent Cloud (that provides fully managed Apache Kafka as a service) to integrate the services.
- Use Spring for Kafka as the API between the two.

If you want to use Spring and Kafka together, apply the right patterns, avoid the wrong antipatterns, and generally get a good glimpse into a robust toolset for building modern Java web applications, look no further than this hands-on workshop.

This workshop is suited for Spring Boot (Java or Kotlin) developers who want to learn more about Apache Kafka®. We provide a code repository and Confluent Cloud coupon codes.

Sorry, registration is now closed.

If you've already reserved a seat, please
sign on or drop in for the workshop to join.

HTTPS://GAMOV.DEV/SPRING20



Learn Kafka.

Start building with
Apache Kafka at
Confluent Developer.

Watch full version

<https://gamov.dev/developer>



developer.confluent.io

<https://developer.confluent.io>



CONFLUENT