
Kafka Event Streaming Applications

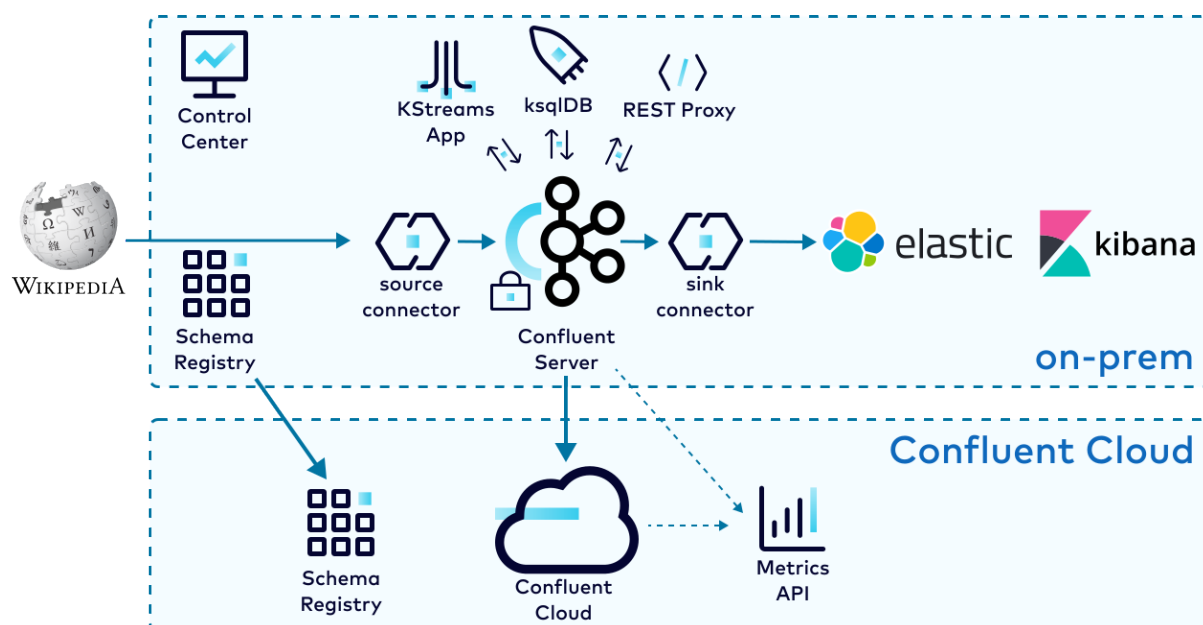
This example and accompanying tutorial show users how to deploy an Apache Kafka® event streaming application using ksqlDB and Kafka Streams for stream processing. All the components in the Confluent Platform have security enabled end-to-end. Run the example with the tutorial.

Table of Contents

- Overview
- Documentation

Overview

The use case is a Kafka event streaming application for real-time edits to real Wikipedia pages. Wikimedia’s EventStreams publishes a continuous stream of real-time edits happening to real wiki pages. Using Kafka Connect, a Kafka source connector `kafka-connect-sse` streams raw messages for the server sent events (SSE), and a custom Kafka Connect transform `kafka-connect-json-schema` transforms these messages and then the messages are written to a Kafka cluster. This example uses ksqlDB and a Kafka Streams application for data processing. Then a Kafka sink connector `kafka-connect-elasticsearch` streams the data out of Kafka and is materialized into Elasticsearch for analysis by Kibana. All data is using Confluent Schema Registry and Avro. Confluent Control Center is managing and monitoring the deployment.



Documentation

You can find the documentation for running this example and its accompanying tutorial at <https://docs.confluent.io/platform/current/tutorials/cp-demo/docs/index.html>.

Additional Examples

For additional examples that showcase streaming applications within an event streaming platform, please refer to the examples GitHub repository.