



About this Document

This document contains sample questions designed to help you prepare for the Confluent Data Streaming Engineer Certification exam. This certification validates your knowledge of building and managing data streaming applications using technologies such as Apache Kafka, Apache Flink, Kafka Connect, Kafka Streams, and Schema Registry.

To learn more about the free certification exam, available learning resources, and the passing score, visit cnfl.io/cert.

The sample questions below reflect the types of questions you can expect on the exam.

1

What is the primary role of a Kafka broker in a Kafka cluster?

A. To manage consumer group offsets

B. To coordinate partition assignments across consumers

C. To store and serve Kafka topics and partitions

D. To serialize and deserialize messages

2

In Kafka, what is a partition?

A. A mechanism to replicate data across brokers

B. A unit of parallelism within a topic

C. A consumer group identifier

D. A message serialization format

3

What is a key feature of Apache Flink's architecture?

A. Batch-only processing

B. Micro-batch processing model

C. True stream processing with low latency

D. Dependency on external schedulers

4

In Flink, what is the purpose of a checkpoint?

A. To monitor system performance

B. To save the state of the application for fault tolerance

C. To partition the data stream

D. To schedule tasks across nodes

5

What is Kafka Connect primarily used for?

A. Managing Kafka consumer groups

B. Monitoring Kafka cluster health

C. Integrating Kafka with external systems

D. Serializing messages in Avro format

6

What is the role of Kafka Connect's worker process?

A. It manages Kafka topic compaction settings

B. It handles message serialization and deserialization

C. It monitors Kafka consumer lag

D. It executes connectors and tasks, and manages configuration and state

7

What is a KTable in Kafka Streams?

A. A stream of immutable records

B. A changelog stream representing the latest value for each key

C. A batch processing table

D. A configuration file for Kafka Streams

8

Which operation would you use in Kafka Streams to group records by key?

A. `groupByKey()`

B. `mapValues()`

C. `filter()`

D. `flatMap()`

9

What is the primary purpose of the Schema Registry in Kafka?

A. To store Kafka topic configurations

B. To manage and enforce data schemas for Kafka messages

C. To monitor Kafka broker health

D. To serialize messages into JSON format

10

Which schema formats are supported by Confluent's Schema Registry?

A. Only Avro

B. Avro, JSON Schema, and Protobuf

C. XML and CSV

D. YAML and TOML

Answer Key	
1. C	6. D
2. B	7. B
3. C	8. A
4. B	9. B
5. C	10. B

 CONFLUENT