Cloudera Streaming Analytics Operator 1.5.0

SSB Application Management

Date published: 2024-06-15
Date modified: 2026-02-18

https://docs.cloudera.com/

Legal Notice

Contents

Job settings

Configure jobs before launching on the Cloudera SQL Stream Builder.

Before launching a new SQL job, you can set the following configurations for your job on Streaming SQL Console on the Kubernetes Resources tab:

- JobManager CPU, memory and ephemeral storage
- JobManager replicas (for high availability)
- TaskManager CPU, memory and ephemeral storage
- Flink deployment mode (*native* or *standalone*)

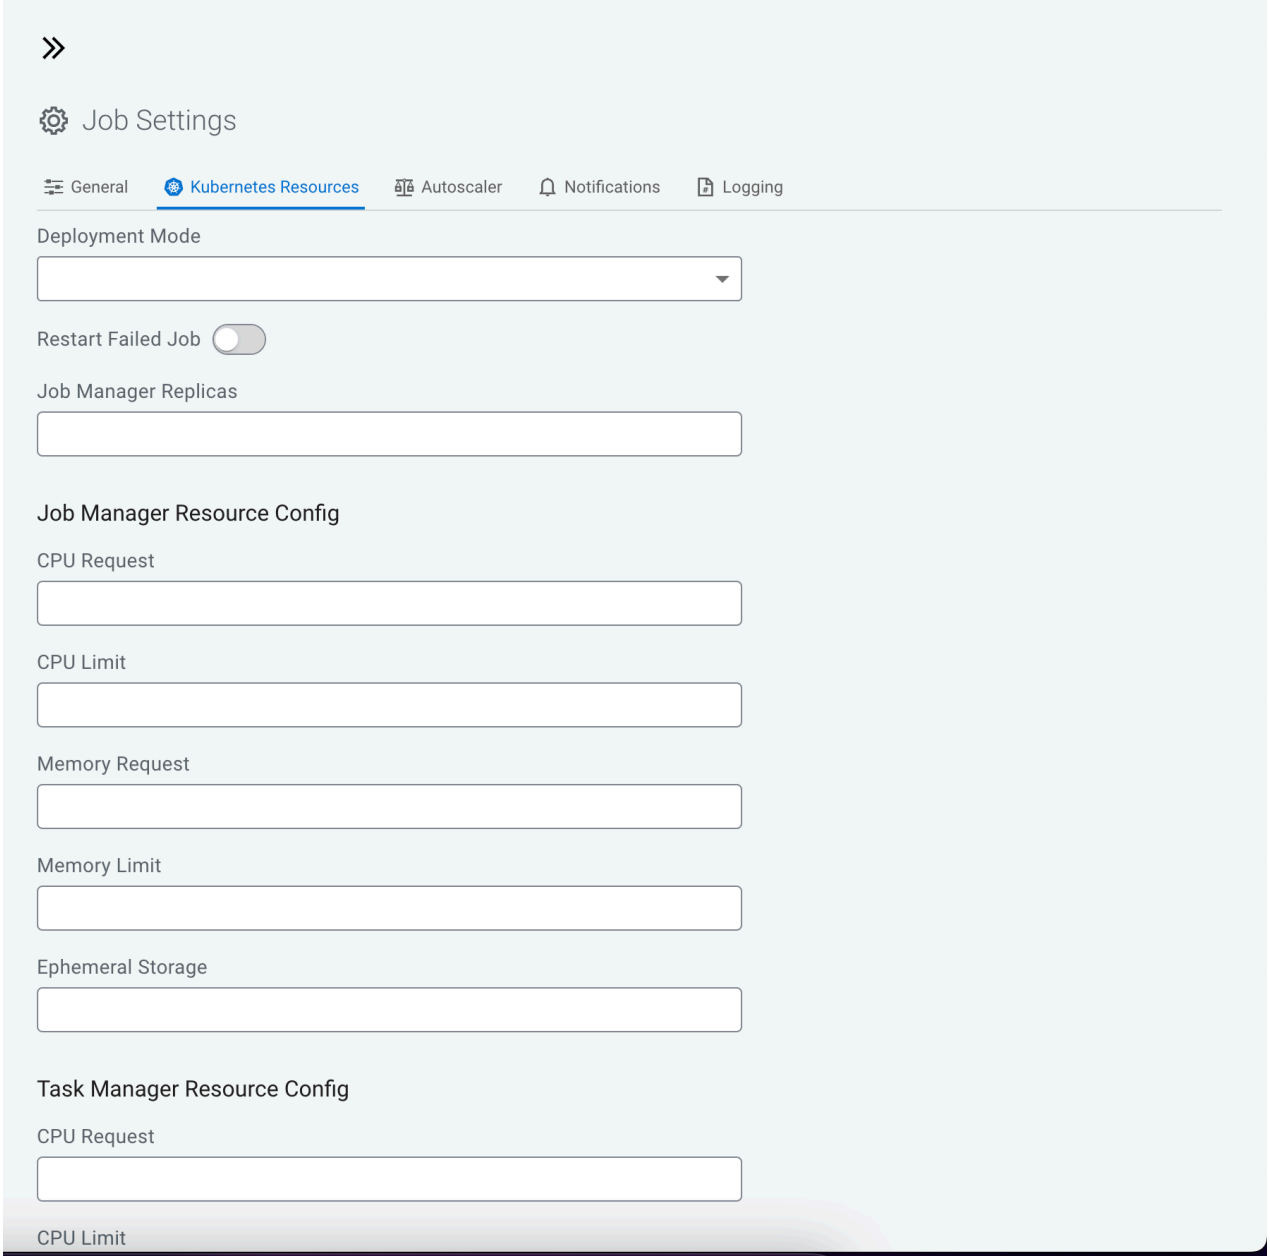

You also have the option to configure the job related settings on the General tab, and adjust the autoscaler configurations on the Autoscaler tab.

Important: For session-mode jobs, these configurations must be set before starting the session cluster. If you modify the default configurations, you need to shut down the existing session cluster and start the job for the changes to take effect.

Sample results

How to view sample results from your SQL queries.

Because Cloudera Streaming Analytics Operator for Kubernetes does not install Kafka, in the Cloudera SQL Stream Builder UI you are not able to see any rows from the Flink jobs. To see sampled results from your SQL queries, you need to have a Kafka cluster installed and accessible by both Cloudera SQL Stream Builder and Flink pods, and change ssbConfiguration to configure Cloudera SQL Stream Builder to use Kafka for data sampling:

```
ssbConfiguration:
  application.properties: |+
 kafka.enabled=true
 spring.kafka.bootstrap-servers=example-kafka:9092
 spring.kafka.jaas.enabled=false
 spring.kafka.properties.security.protocol=PLAINTEXT
```