

Cloudera Streaming Analytics Operator 1.0.0

SSB Application Management

Date published: 2024-06-15

Date modified: 2024-06-15

The Cloudera logo is displayed in a bold, orange, sans-serif font. The word "CLOUDERA" is written in all caps, with a stylized 'E' that has a horizontal bar extending to the right.

<https://docs.cloudera.com/>

Legal Notice

© Cloudera Inc. 2024. All rights reserved.

The documentation is and contains Cloudera proprietary information protected by copyright and other intellectual property rights. No license under copyright or any other intellectual property right is granted herein.

Unless otherwise noted, scripts and sample code are licensed under the Apache License, Version 2.0.

Copyright information for Cloudera software may be found within the documentation accompanying each component in a particular release.

Cloudera software includes software from various open source or other third party projects, and may be released under the Apache Software License 2.0 (“ASLv2”), the Affero General Public License version 3 (AGPLv3), or other license terms. Other software included may be released under the terms of alternative open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

Please visit the Cloudera software product page for more information on Cloudera software. For more information on Cloudera support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Cloudera reserves the right to change any products at any time, and without notice. Cloudera assumes no responsibility nor liability arising from the use of products, except as expressly agreed to in writing by Cloudera.

Cloudera, Cloudera Altus, HUE, Impala, Cloudera Impala, and other Cloudera marks are registered or unregistered trademarks in the United States and other countries. All other trademarks are the property of their respective owners.

Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES NOT WARRANT THAT CLOUDERA PRODUCTS NOR SOFTWARE WILL OPERATE UNINTERRUPTED NOR THAT IT WILL BE FREE FROM DEFECTS NOR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION NOR UNAVAILABILITY, NOR THAT IT WILL MEET ALL OF CUSTOMER’S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLOUDERA EXPRESSLY DISCLAIMS ANY AND ALL IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, QUALITY, NON-INFRINGEMENT, TITLE, AND FITNESS FOR A PARTICULAR PURPOSE AND ANY REPRESENTATION, WARRANTY, OR COVENANT BASED ON COURSE OF DEALING OR USAGE IN TRADE.

Contents

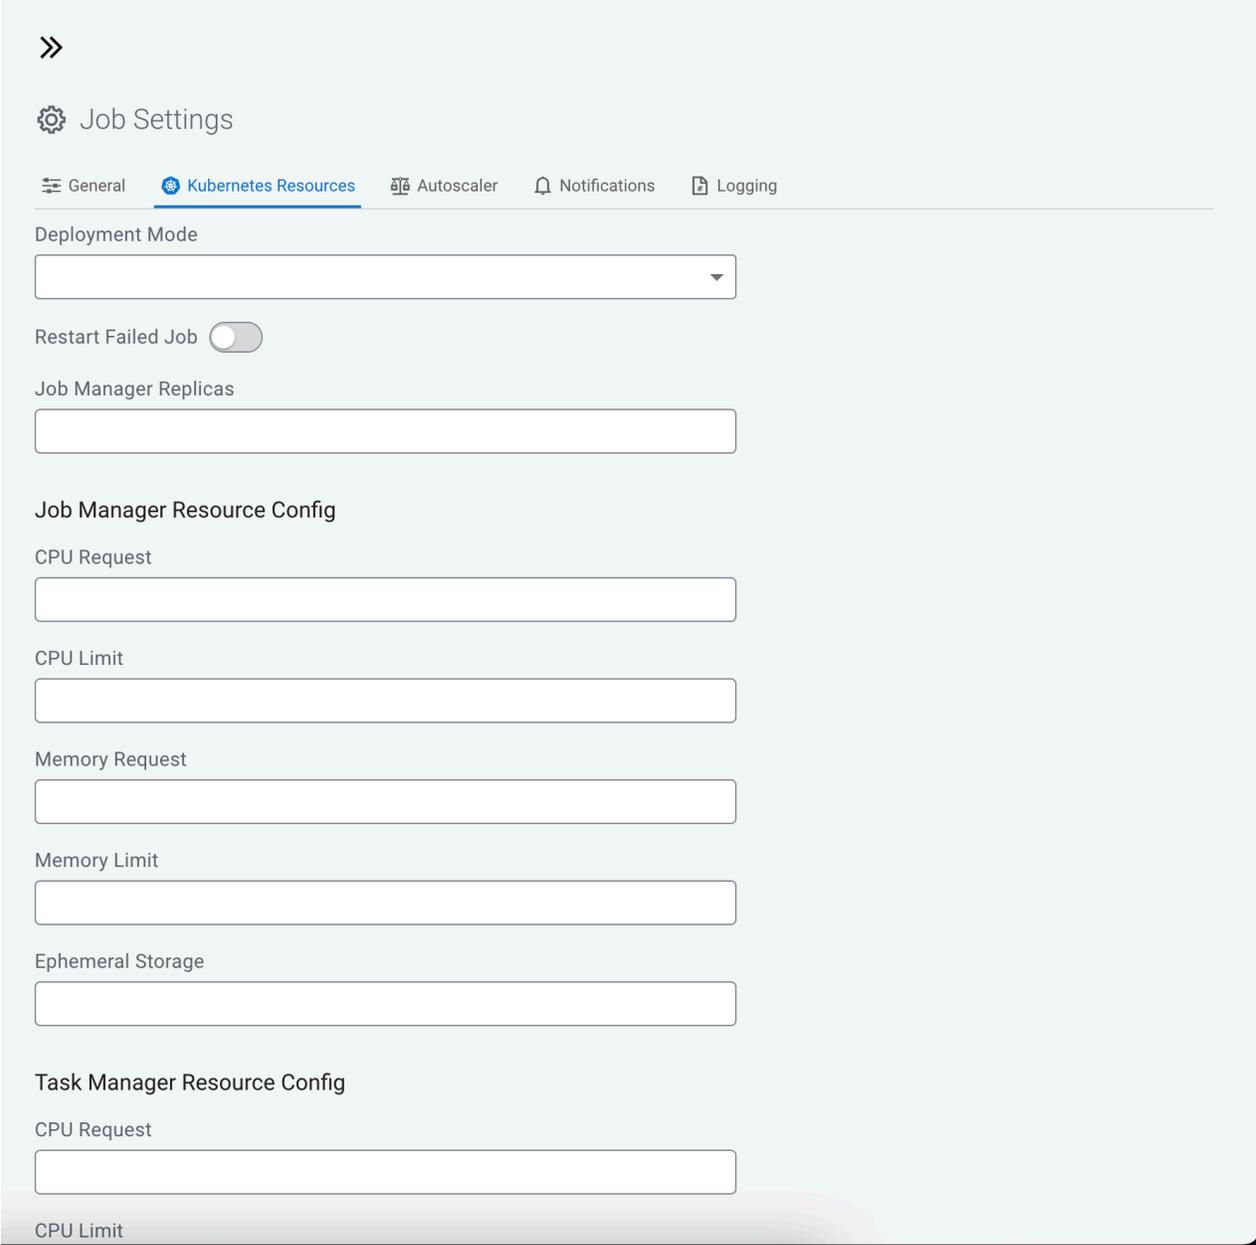
Job settings.....	4
Sample results.....	5

Job settings

Configure jobs before launching on the Stream SQL 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*)



The screenshot shows the 'Job Settings' configuration page. At the top, there is a navigation bar with tabs: 'General', 'Kubernetes Resources' (selected), 'Autoscaler', 'Notifications', and 'Logging'. Below the tabs, the 'Deployment Mode' is set to a dropdown menu. A 'Restart Failed Job' toggle switch is currently turned off. The 'Job Manager Replicas' is a text input field. Under the 'Job Manager Resource Config' section, there are text input fields for 'CPU Request', 'CPU Limit', 'Memory Request', 'Memory Limit', and 'Ephemeral Storage'. Under the 'Task Manager Resource Config' section, there are text input fields for 'CPU Request' and 'CPU Limit'.

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 CSA Operator does not install Kafka, in the SSB 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 SSB and Flink pods, and change `ssbConfiguration` to configure SSB 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
```