

## Monitoring SQL jobs

Date published: 2019-12-17

Date modified: 2021-09-29



# 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

|   |          |
|---|----------|
| <b>Monitoring SQL Stream jobs.....</b>                          | <b>4</b> |
| <b>Enabling Flink DEBUG logging for SQL Stream Builder.....</b> | <b>5</b> |

## Monitoring SQL Stream jobs

You can use the Streaming SQL Console to review the status, properties and log of your SQL Stream jobs executed in SQL Stream Builder. Using the Flink Dashboard, you can also monitor the Flink job that is submitted when you execute a SQL query.

### Using the Streaming SQL Console

When using the **SQL Jobs** page in Streaming SQL Console to monitor your SQL jobs, you can review the ID, the Name, the Start time, the State of the submitted jobs, and the User who submitted the SQL jobs. When monitoring running jobs, you are also able to open the Flink Dashboard, and stop the job using the Stop button under Actions. The Flink Dashboard link and Stop button are not available for Stopped Jobs.

1. Click Console on the main menu.
2. Select SQL Jobs tab.

By default, Running Jobs are displayed on the **SQL Jobs** page.

3. Select Running Jobs or Stopped Jobs from the drop-down menu.
4. Click on the job you want to monitor.
  - a. You can further filter down the results, by directly searching for the job name in the Search field.
5. Select Details tab on the bottom panel to display additional details and configurations about the SQL job.

| Property            | Value                           |
|---------------------|---------------------------------|
| Job Name            | nostalgic_varahamihira          |
| Sources             | transactions                    |
| Sinks               | No sinks defined                |
| SQL                 | select * from transactions      |
| Parallelism         | 1 thread(s)                     |
| Restart Strategy    | never                           |
| Restart Retry Delay | 30                              |
| Sample Behavior     | Sample one message every second |

6. Select Log tab on the bottom panel to review the log information of the job submission.



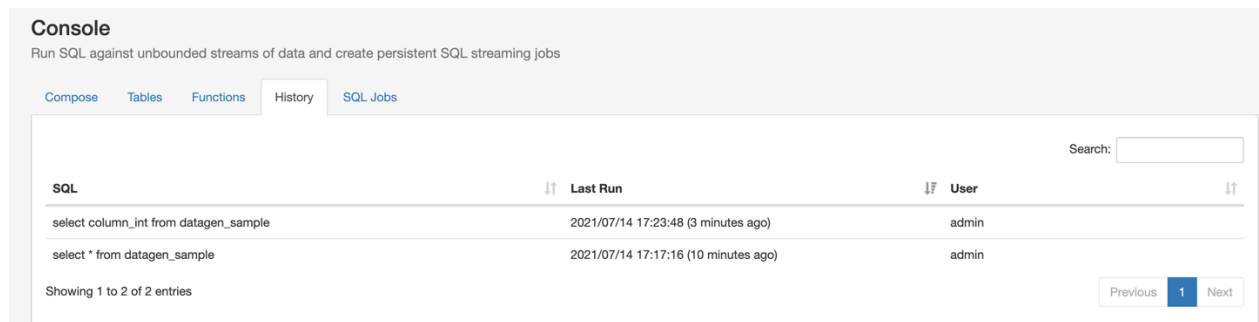
```

[14/07/2021, 18:47:45][INFO] INFO - 2021-07-14 16:42:08.274652 : Job metadata saved
[14/07/2021, 18:47:45][INFO] INFO - 2021-07-14 16:42:08.284747 : SSB version 1626276514638921589 selected for job.
[14/07/2021, 18:47:45][INFO] INFO - 2021-07-14 16:42:10.680009 : Successful job deployment to: http://docs-test-3.docs-test.root.hwx.site:45643
[14/07/2021, 18:47:45][INFO] INFO - 2021-07-14 16:42:10.705042 : Free slots remaining: 0
[14/07/2021, 18:47:45][INFO] INFO - 2021-07-14 16:42:10.719030 : Generated Flink ID: fl1906519f6932472e6bb9b9f31219a5

```

## History of SQL queries

You can review and reuse the SQL queries that were previously executed on the **History** page. When you click on one of the SQL queries, it is automatically imported to the SQL window for execution. You can filter the SQL queries by the time they were last run or by the user who run them. You can also search for a type of SQL query using the Search field.



**Console**  
Run SQL against unbounded streams of data and create persistent SQL streaming jobs

Compose Tables Functions History **SQL Jobs**

Search:

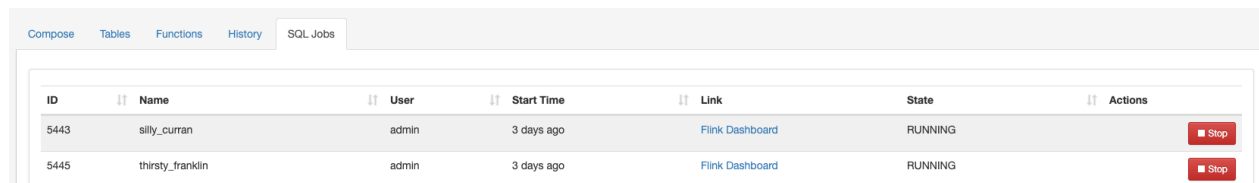
| SQL                                   | Last Run                             | User  |
|---------------------------------------|--------------------------------------|-------|
| select column_int from datagen_sample | 2021/07/14 17:23:48 (3 minutes ago)  | admin |
| select * from datagen_sample          | 2021/07/14 17:17:16 (10 minutes ago) | admin |

Showing 1 to 2 of 2 entries

Previous 1 Next

## Using the Flink Dashboard

You can also monitor your running SQL jobs using the Flink Dashboard. You can easily reach the Flink Dashboard on the SQL jobs tab below the Console main menu. After clicking on the Flink Dashboard, you are redirected to the Flink Dashboard interface.



Compose Tables Functions History **SQL Jobs**

| ID   | Name             | User  | Start Time | Link                            | State   | Actions              |
|------|------------------|-------|------------|---------------------------------|---------|----------------------|
| 5443 | silly_curran     | admin | 3 days ago | <a href="#">Flink Dashboard</a> | RUNNING | <a href="#">Stop</a> |
| 5445 | thirsty_franklin | admin | 3 days ago | <a href="#">Flink Dashboard</a> | RUNNING | <a href="#">Stop</a> |

# Enabling Flink DEBUG logging for SQL Stream Builder

You can review the log text files of the Flink jobs behind SQL Stream Builder (SSB) when an error is detected during the processes. When you set the log level of Flink to DEBUG, you can easily trace the log file for errors.

## About this task

A log file is created about every Flink process that contains messages for the different events happening in the given process. You can use these log files to solve the errors and problems that can occur during SSB and Flink processes. You can access the Flink logs using the Flink Dashboard.

## Procedure

1. Go to your cluster in Cloudera Manager.
2. Select Flink from the list of Services.
3. Select Configuration.

4. Search for Flink Client Advanced Configuration Snippet (Safety Valve) for flink-conf/log4j.properties configuration.
5. Add the following parameters to the Safety Valve:

```
logger.flink.name = org.apache.flink
logger.flink.level = DEBUG (or TRACE)
```

6. Click Save Changes.
7. Restart the Flink service.
8. Restart the SQL Stream Builder service.
9. Navigate to the **YARN Resource Manager** to stop the SQL Stream Builder YARN job.
  - a) Select YARN from the list of Services.
  - b) Select YARN.
  - c) Select Applications.

The running SSB application is displayed.
10. Click the settings represented by three dots.
11. Select the stop option.
12. Navigate to **Streaming SQL Console** and run any job.
  - a) Select Clusters > SQLStreamBuilder .
  - b) Click SQLStreamBuilder Console.

You are redirected to the **Streaming SQL Console**.
  - c) Execute a SQL query in the SQL window.
13. Navigate to **Flink Dashboard** and review the log level for the running job.
  - a) Select SQL jobs on the main tabs of **Streaming SQL Console**.
  - b) Select Running jobs.
  - c) Search for the previously submitted job.
  - d) Click Flink Dashboard.

You are redirected to the Flink Dashboard page.
14. Select Task Managers from the main menu.
15. Select the previously submitted job.
16. Click Logs.