Cloudera Streaming Analytics 1.8.0

Monitoring SQL jobs

Date published: 2019-12-17
Date modified: 2022-09-28

https://docs.cloudera.com/

Legal Notice

Contents

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 job. 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. Select SQL Jobs from the main menu.

   By default, all of the jobs are displayed on the **SQL Jobs** page.

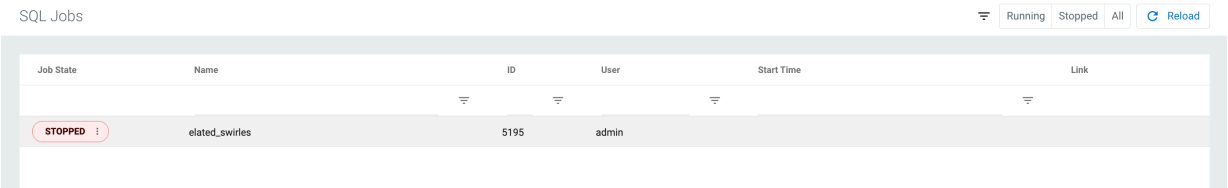

2. Select Running or Stopped.

3. Click on the job you want to monitor.

   The details of the job appear in a new window.

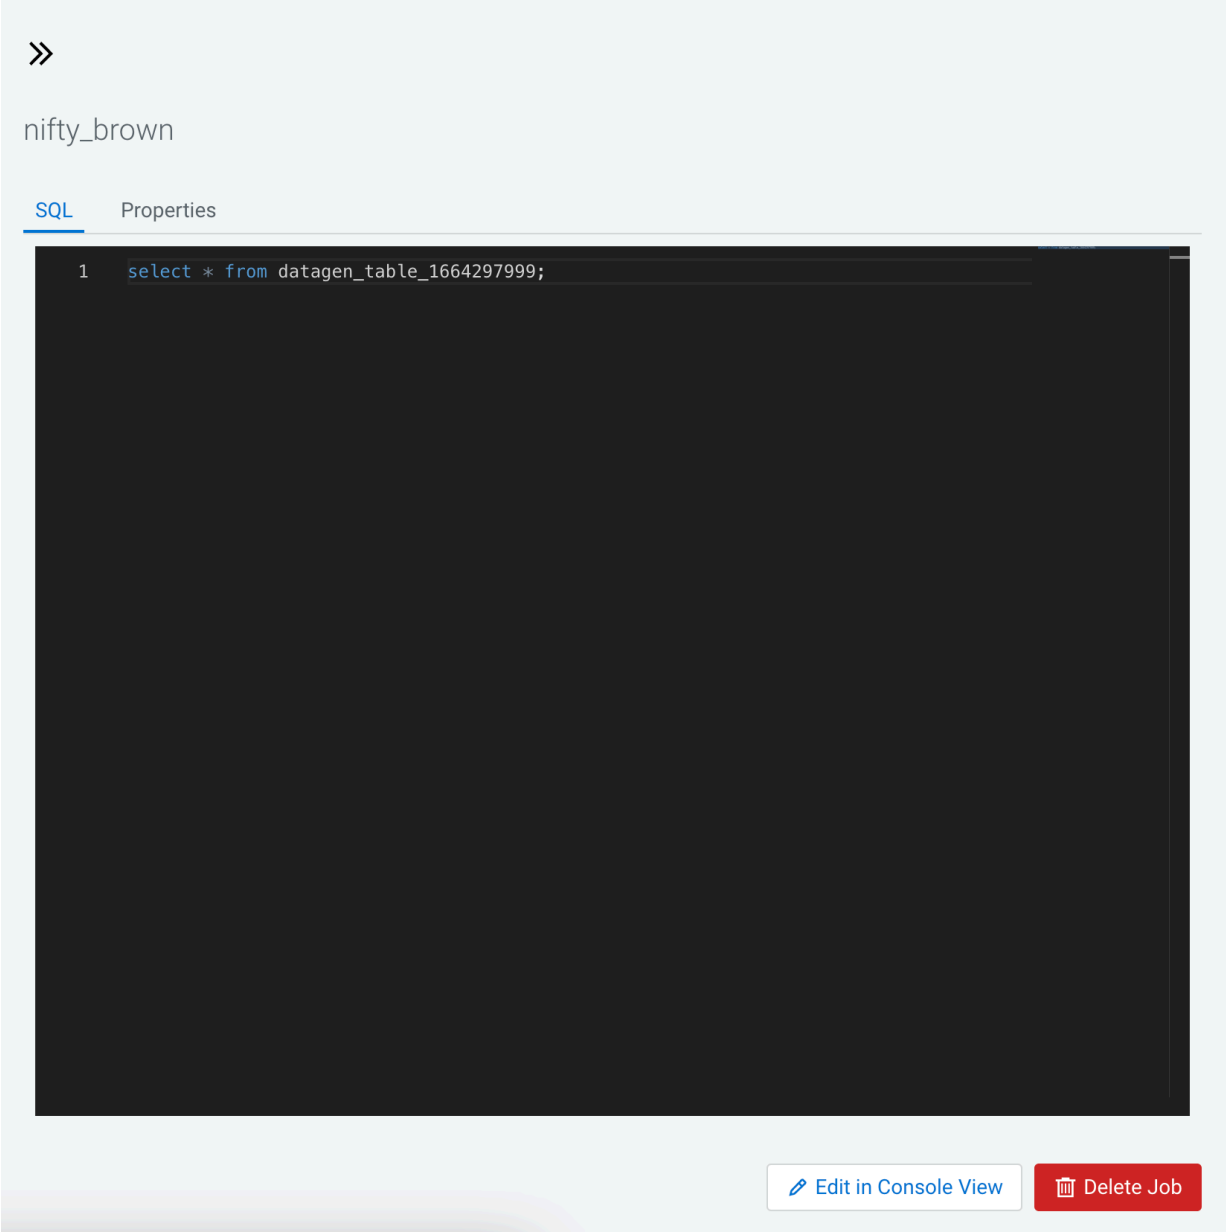

4. Select SQL or Properties tab to display additional details and configurations about the SQL job.

   You can also review the details of the Flink job from the **SQL Jobs** page.

5. Click next to **Job State** of a job on the **SQL Jobs** page.

6. Select Events.

   The **Events** window appears detailing the Flink job information that can be further filtered based on the log level settings.

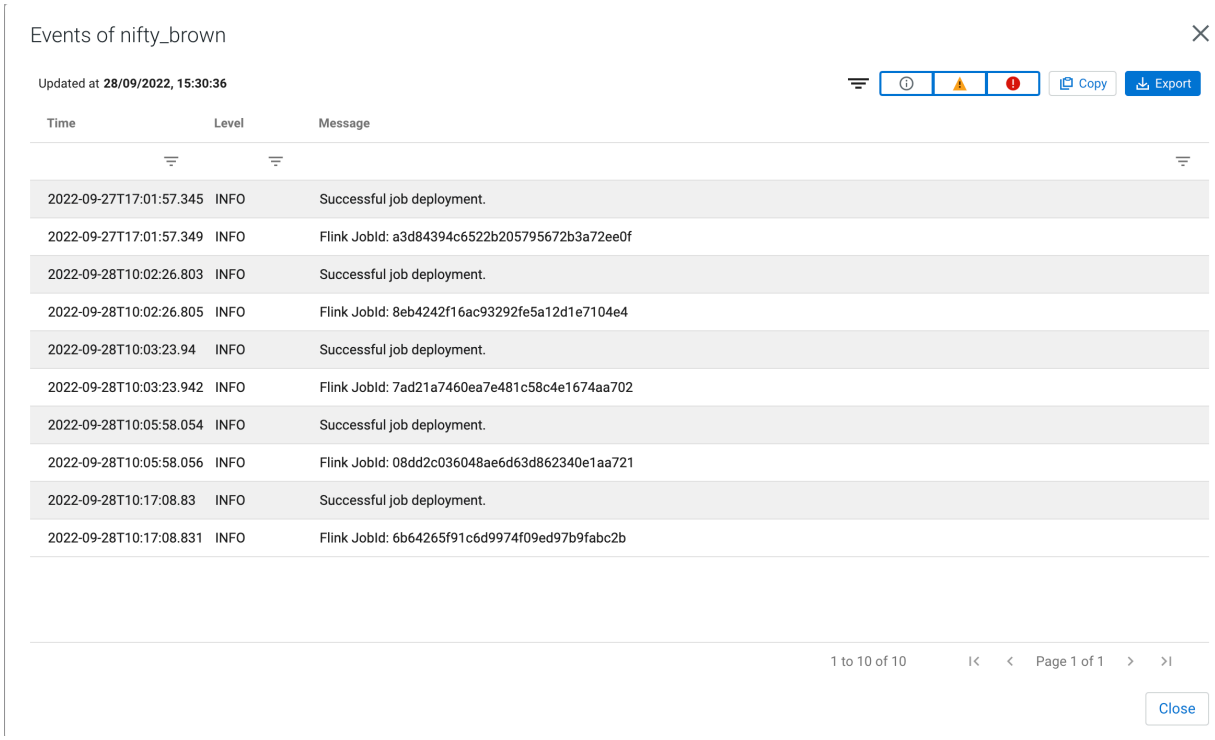

History of SQL queries

You can review and reuse the SQL queries that were previosuly 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.

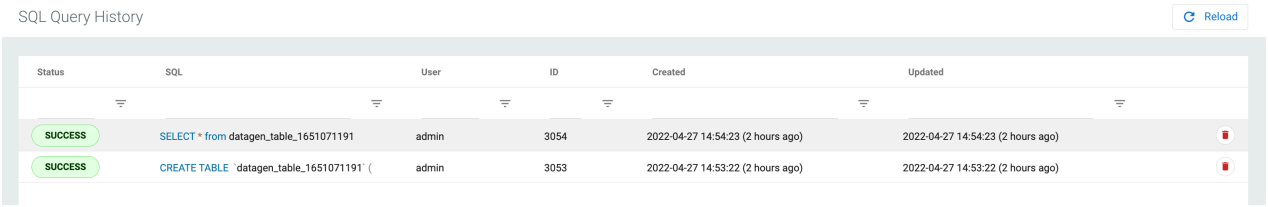

Using the Flink Dashboard

You can also monitor your running SQL jobs using the Flink Dashboard. You can easily reach the Flink Dashboard from the main menu of Streaming SQL Console.

Using SQL job notification

You can set up notifications for SQL jobs using Streaming SQL Console to alert individuals and teams about job status changes. The notifications include email and webhook alerts, which you can also group together for easier organization.

You can add job notifications to SQL jobs to send alerts and information when the status of the job is changed. The information in the notification can include the job name, job ID, cluster ID, last exception, and Flink job ID.

You can create the following notification actions for SQL jobs:

Email

Sending notification email to the provided email address

Webhook

Sending notification to a webhook address using POST and PUT

Group

Grouping email and webhook notifications into one group

After creating and adding the notification to a SQL job, anytime a change occurs for that job, a notification is sent to the specified address.

Enabling job notifications

Before you are able to use the job notifications, you need to enable it in Cloudera Manager and based on the notification type configure the email and webhook parameters.

1. Go to your cluster in Cloudera Manager.
2. Select SQL Stream Builder from the list of services.
3. Click Configuration.
4. Search for *ENABLE JOB NOTIFICATIONS FUNCTIONALITY*.
5. Enable the job notifications by checking the checkbox.

The following parameters also need to be configured for the SQL Stream Builder (SSB) service in Cloudera Manager based on the type of the notification:

Configuration	Description
Job notifications monitoring interval	Sets the interval of the job monitoring in seconds.
Mail server host for job notifications	The host of the SMTP server for job failure notifications.
Mail server username for job notifications	The username to access the SMTP server for job failure notifications.
Mail server password for job notifications	The password to access the SMTP server for job failure notifications.
SMTP authentication for job notifications	Enable SMTP authentication for job notifications.
StartTLS for job notifications	Use the StartTLS command to establish a secure connection to the SMTP server for job notifications.
Job notifications sender mail address	Sender mail address for job notifications.
Mail server port for job notifications	The port of the SMTP server for job failure notifications.
Job notifications webhook sender parallelism	Number of threads used by the job notification task to call user-specified webhooks when notifying about a failed or missing job.

Note: In case the YARN cluster of a failed job no longer exists, SSB attempts to fetch the information about the failed job from the Flink History Server. If information cannot be fetched by the first attempt, SSB attempts to fetch the information again based on the interval configured for the History server fetch retry time parameter.

Creating job notifications

After enabling the job notification function, you can create email, webhook and group notifications with Streaming SQL Console.

For Email

1. Navigate to Streaming SQL Console.
   a. Go to your cluster in Cloudera Manager.
   b. Select SQL Stream Builder from the list of services.
   c. Click Streaming SQL Console.

   The **Streaming SQL Console** opens in a new window.
2. Click Job Notifications from the main menu on Streaming SQL Console.

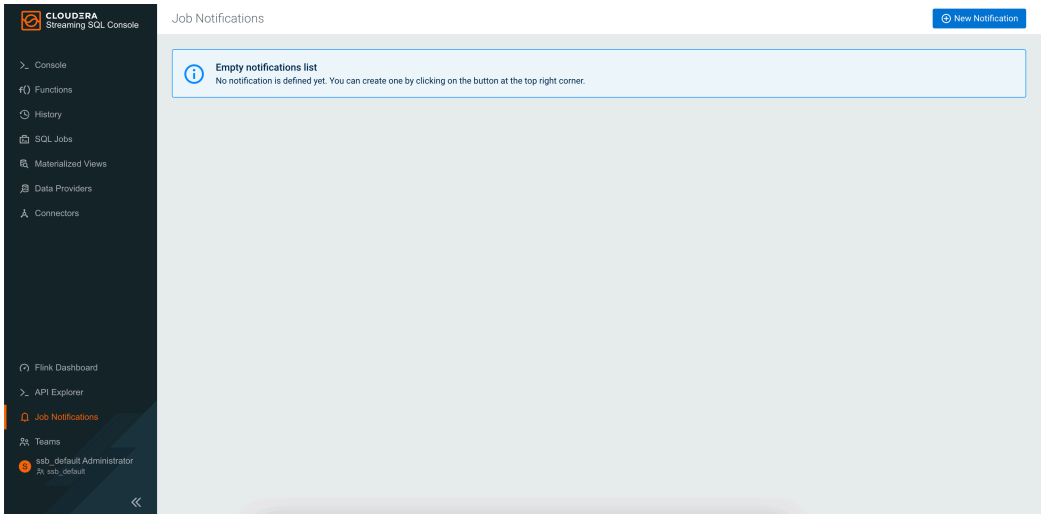

3. Click Create > Email .

   The **Create Email Notification** window appears.

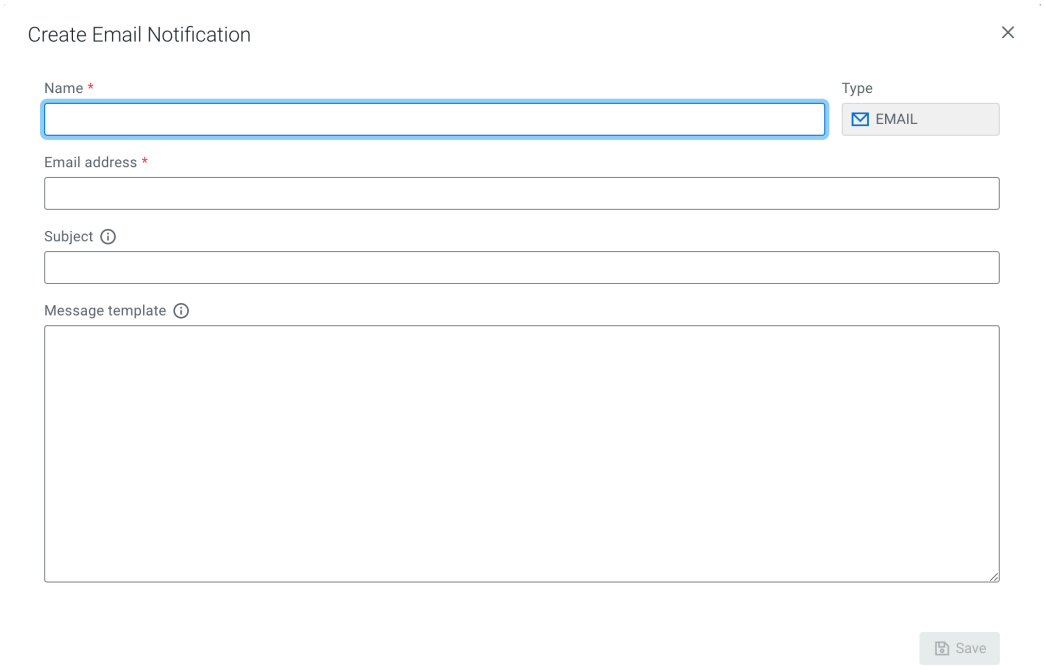

4. Provide a Name for the notification.
5. Add the Email address you want to send the notification.
6. Provide a Subject.

7. Provide a Message.

   You can add placeholder items to the payload that will be converted into the specific job information. You can add the placeholders in ${placeholder} format, and the following placeholders are available:

   - jobName
   - jobStatus
   - jobStatusDescription
   - ssbJobId
   - flinkJobId
   - clusterId
   - lastException

8. Click Save.

For Webhook

1. Navigate to Streaming SQL Console.

   a. Go to your cluster in Cloudera Manager.
   b. Select SQL Stream Builder from the list of services.
   c. Click Streaming SQL Console.

   The **Streaming SQL Console** opens in a new window.

2. Click Job Notifications from the main menu on Streaming SQL Console.

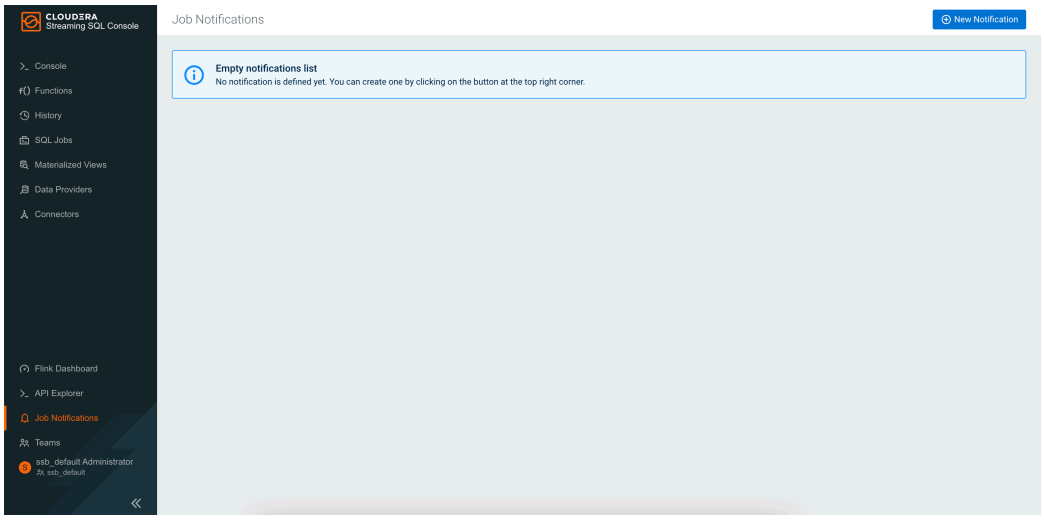

3. Click Create > Webhook .

   The **Create Webhook Notification** window appears.

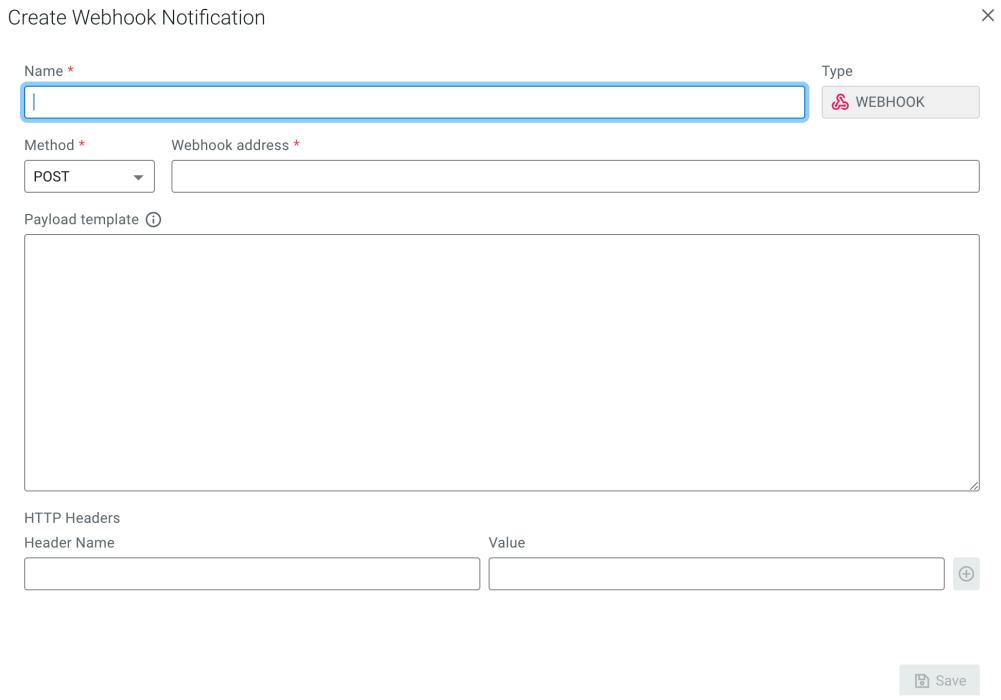

4. Provide a Name for the notification.
5. Select between POST or PUT method.
6. Provide a Webhook address.
7. Provide a Payload to the webhook.

   Warning: The payload must be a valid JSON object.

   You can add placeholder items to the payload that will be converted into the specific job information. You can add the placeholders in ${placeholder} format, and the following placeholders are available:

   - jobName
   - jobStatus
   - jobStatusDescription
   - ssbJobId
   - flinkJobId
   - clusterId
   - lastException
8. Add **HTTP Headers** if needed by specifying the Header Name and Value.
9. Click Save.

For Groups

1. Navigate to Streaming SQL Console.
   a. Go to your cluster in Cloudera Manager.
   b. Select SQL Stream Builder from the list of services.
   c. Click Streaming SQL Console.

   The **Streaming SQL Console** opens in a new window.

2. Click Job Notifications from the main menu on Streaming SQL Console.

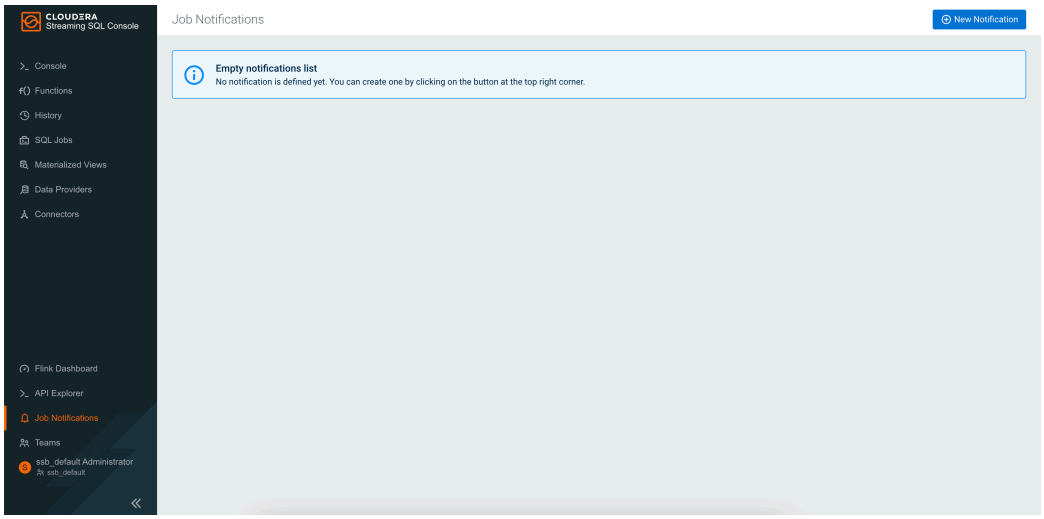

3. Click Create > Group .

   The **Create Group Notification** window appears.

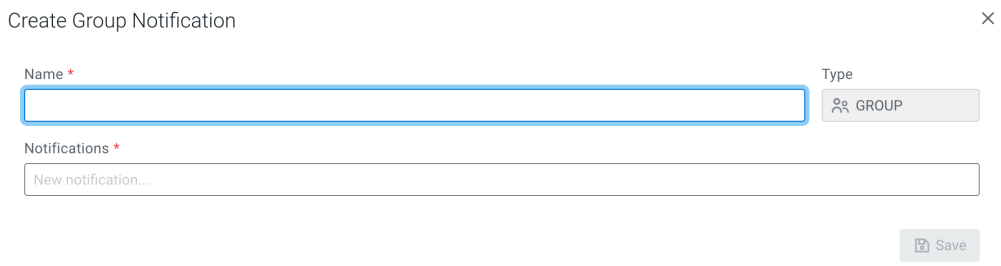

4. Select from the already created Email and Webhook notifications that you would like to add to this group.
5. Click Save.

Enabling notifications for a SQL job

After creating a job notification, you can enable it for a selected job from the **Job Settings** window.

1. Select Job Settings on the Console page.
2. Click Job Notifications.
3. Search for the name that you have provided when creating the notification.
4. Click on the notification.

   The selected notification is added to the job and the notification is sent with the information provided in the message template.

Collecting diagnostic data

You can collect diagnostic data from SQL Stream Builder (SSB) and Flink using Cloudera Manager, and send it to Cloudera Support for easier troubleshooting.

When using SSB and Flink, Cloudera Manager automatically collects the usage and diagnostic data from the services. The collected data is analyzed to further improve the software, and to help you in support cases when you encounter

an error. You can configure the frequency of data collection, and you can also manually trigger sending the diagnostic data to Cloudera.

For more information about the diagnostic data collection and configuration, see the Cloudera Manager documentation.

The following telemetry data is collected when using SSB and Flink:

- Number of users
- Number of jobs
- Number of registered data providers and connectors
- Number of Materialized Views
- Number of API keys
- Number of tables
- Number of catalogs
- Total tasks slots in use
- Log files from Streaming Engine and Materialized View Engine
- Log files from YARN to collect Flink Job and Task Manager processes