Cloudera Streaming Analytics 1.6.0

Migration

Date published: 2019-12-17
Date modified: 2021-12-15



https://docs.cloudera.com/

Legal Notice

© Cloudera Inc. 2025. 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

Migrating Flink service to a different host.....4

Migrating SQL jobs.....4

Migrating Flink service to a different host

The Flink Dashboard role cannot be directly migrated from one host to another. You need to delete the existing role and add a new one to the host where you want to migrate the Flink Dashboard role. As Flink has mandatory dependencies for YARN, ZooKeeper and HDFS, these services should already be on the target host.

Before you begin

- Make sure that you have stopped your Flink jobs with savepoints, and saved your Cloudera Manager configurations.

  For more information, see the Stopping Flink applications and the Backing up the Cloudera Manager configuration documentations.
- Make sure that YARN, Zookeeper and HDFS services are already installed on the target host.

Procedure

1. Go your cluster in Cloudera Manager.
2. Select Flink from the list of services.
3. Click Instances.
4. Check the box for the Flink Dashboard role.
5. Click Action for Selected > Stop to stop the Flink Dashboard role.
6. Click Action for Selected > Restart to delete the Flink Dashboard role.

   When deleting the Flink Dashboard role, an error message will be displayed as the Flink Dashboard role is mandatory for the Flink service:

   ```
   Service FLINK-1 has 0 Flink Dashboards. Flink requires at least 1 Flink Dashboard.
   ```
7. Add a new Flink Dashboard role instance to the target host.
   a) Click Add Role Instances.

      The **Add Role Instances** wizard appears.
   b) Assign the Flink Dashboard role to a host.
   c) Review the changes.
   d) Click Finish.

      The new role is stopped when added.
8. Check the box for the Flink Dashboard role.
9. Click Action for Selected > Restart to restart the new Flink Dashboard role.
10. Migrate the Flink configurations, and submit your Flink jobs to the host where the newly added Flink Dashboard role is running.

What to do next

Start your Flink applications from savepoint and restore the Cloudera Manager configurations.

Migrating SQL jobs

You can migrate your SQL jobs using the SQL Stream Builder (SSB) job migration tool. The migration tool can be used from Command Line Interface (CLI) or using the SSB REST API. The migration tool enables you to export your SQL jobs with configurations, and import them to a different cluster or deployment when needed.

You can migrate you SQL jobs either using the API Explorer or the SSB migration tool from CLI. The migration process consists of exporting and importing the details of the SQL job. No matter which migration option you choose, the underlying process is the same as both solutions are based on calling the SSB REST API endpoints. However, when using the API Explorer, you need to manual copy and paste the job details between the endpoints. The migration tool in CLI automatically takes care of accessing this information.

Before migrating your SQL jobs, you need to stop the job if it is in a running state. To stop the SQL job you can use the Stop button either under the SQL window on the Compose page, or next to the running job on the SQL Jobs tab. Make a note of the SQL job name as you need to provide it for the job migration.

For API Explorer

1. Access the API Explorer.
   a. Click API Explorer on the main menu of Streaming SQL Console.
   b. Go to Cloudera Manager > Cluster > SQL Stream Builder , and click SQLStreamBuilder API Explorer from Web UI.
2. Open the GET/api/v1/ssb/jobs/{jobsName} operation.
3. Paste the copied job name to jobName.
4. Provide your username.
5. Click Execute.

   The details about the job is exported, and returned in a JSON string. To import the job, you need to copy the JSON detail and add it to the POST request.
6. Copy the JSON string.
7. Open the POST/api/v1/ssb/jobs/{jobsName} operation.
8. Provide a job name.
9. Provide your username.
10. Click Execute.

    The job is imported to the Streaming SQL Console, and is listed under the **SQL jobs** tab with a stopped status.

For CLI

1. Connect to your cluster using ssh.
2. Use the following command to export the job:

```
ssb-job-migration -api <cluster_hostname> -o job-export -j <jobname> -f
<target_filename>
```

3. Use the following command to import the job:

```
ssb-job-migration -api <cluster_hostname> -o job-import -j <jobname> -f
<source_filename>
```

The job is imported to the Streaming SQL Console, and is listed under the **SQL jobs** tab with a stopped status.

When using the migration tool in a secure environment, you need to provide your security credentials to submit the commands.

TLS/SSL enabled

You need to provide a private CA signed certificate (global_truststore.jks) using the -ts and -tspw arguments for the truststore path and password.

Authentication

- For keytabs, you can use kinit, and then submit the job migration operation.
- For Knox, you need to provide your Knox credentials with -jt, -pw and -u arguments for the Knox JWT token, password and user.