

Upgrade

Date published: 2019-12-17

Date modified: 2022-02-28

CLOUDERA

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

Upgrading Flink service.....	4
Creating new Flink application versions.....	4
Stopping Flink applications.....	4
Upgrading CSA artifacts for Flink service.....	4
Resuming Flink applications.....	5
 Upgrading SQL Stream Builder service.....	 6
Stopping SQL Stream jobs.....	6
Upgrading CSA artifacts for SQL Stream Builder service.....	7
Resuming SQL jobs.....	8

Upgrading Flink service

To upgrade the Flink service to the latest version in Cloudera Streaming Analytics, you need to create new application versions, stop your Flink applications with savepoint, upgrade the CSA artifacts and resume your Flink applications from savepoint.

Creating new Flink application versions

You need to build new versions of your existing Flink applications by updating the Cloudera Streaming Analytics (CSA) dependencies to the latest version.

You need to verify that your new application versions are compatible with the planned CSA release. For example, if you build your project with Maven, you need to increment your version number in the pom.xml.

Old version: `<flink.version>x.x.x-csax.x.x.x</flink.version>`

New version: `<flink.version>y.y.y-csay.y.y.y</flink.version>`

Stopping Flink applications

You must stop your Flink applications before updating the artifacts. You can use the stop command to create a savepoint of your application. You can also use the cancel command to stop your Flink application without creating a savepoint.

About this task

You can use savepoints to resume the application state after the upgrade.

Procedure

1. Determine the related Flink job IDs.

```
flink list
```

2. Stop your Flink applications.

You have two choices:

- Stop your applications with a savepoint to store the application state.

```
flink stop <Flink job ID>
```

The command returns an HDFS path, which is the automatically created savepoint that stores the application state.

- Cancel your applications without creating a savepoint.

```
flink cancel <Flink job ID>
```

Use this method if you do not need to restore the application state after the upgrade.

Upgrading CSA artifacts for Flink service

You need to upgrade the Cloudera Streaming Analytics (CSA) artifacts to the newer version by manually deleting the old CSD file from the repository, and adding the latest CSD file to the same location.

Procedure

1. Download the new CSD files from the [Download location](#).

In case you are upgrading to the latest CSA version that includes SQL Stream Builder (SSB), you also must download the SSB CSD, and add it to the `/opt/cloudera/csd` folder with the new Flink CSD. For more information, see the [Installing parcel and CSD](#) documentation.

2. Delete the old CSD version from `/opt/cloudera/csd`.
3. Copy the new artifacts to the `/opt/cloudera/csd` folder.



Warning: The upgrade process fails if you only add the new parcel file in Cloudera Manager. You must add the Flink and SSB CSD to the `/opt/cloudera/csd` folder as well.

4. Ensure that you maintain the appropriate file ownership and access the attributes if needed.
5. Restart the Cloudera Manager Server with the following command:
 - RHEL 7 Compatible, SLES 12, Ubuntu: `systemctl restart cloudera-scm-server`
 - RHEL 6 Compatible: `service cloudera-scm-server restart`
6. Restart the Cloudera Management Services.
7. Install and activate the new CSA parcel from Cloudera Manager.
8. Restart the Flink service when prompted in the wizard.

For more information about the general process of Service Management, see the [Cloudera Manager documentation](#).

Resuming Flink applications

After updating the Cloudera Streaming Analytics artifacts, you can resume your application by the Flink run command. In case you have stopped your application by creating a savepoint, you must add the savepoint location to the command.

If you had stopped your Flink applications with a savepoint, you must define the path to the savepoint with an additional `-s <savepointPath>` parameter to the original run command.





```
flink run -d -ynm YarnApplicationName -s <savepointPath> application.jar
```

To start your stateless applications or stateful applications with an empty application state, use the original Flink run command.



Important:

Additional TLS related properties were added to the Gateway role. If you previously enabled TLS on your cluster, you can define Gateway TLS/SSL Client Trust Store File and Gateway TLS/SSL Client Trust Store Password configurations. You can set `{{CM_AUTO_TLS}}` as the property values when using Auto TLS in Cloudera Manager.

Gateway TLS/SSL Client Trust Store File security.ssl.rest.truststore	Gateway Default Group  <input type="text" value="{{CM_AUTO_TLS}}"/>	
Gateway TLS/SSL Client Trust Store Password security.ssl.rest.truststore-password	Gateway Default Group  <input type="password" value="*****"/>	

You can verify your settings by running the `flink list -yD security.ssl.rest.enabled=true` command.

Related Information

[Running a Flink job](#)

Upgrading SQL Stream Builder service

To upgrade the SQL Stream Builder service to the latest version in Cloudera Streaming Analytics, you need to stop your SQL jobs, update the CSA artifacts and resume your SQL jobs on the Streaming SQL Console.

Related Information

[Downloading CSA](#)

[Installing CSD and parcels](#)

[Adding SSB as a service](#)

Stopping SQL Stream jobs

Before upgrading SQL Stream Builder, you need to stop your running jobs on the SQL Jobs page of Streaming SQL Console, and stop all Flink and SQL Stream Builder related YARN sessions. Otherwise Flink is not able to submit applications to YARN after upgrading the artifacts.

Procedure

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

2. Select SQL Jobs tab.

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

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

Compose Tables Functions History **SQL Jobs**

Search by name Running Jobs

ID	Name	User	Start Time	Link	State	Actions
5209	loving_wright	admin	2021-09-28 17:54:51 (30 minutes ago)	Flink Dashboard	RUNNING	Stop

Previous 1 Next

Details [Log](#)

No Job Selected

Click on a Job and its details will display here

3. Click on the job you want to stop.

You can further filter down the results, by directly searching for the job name in the Search field.
4. Click Stop under Actions.

After you stop the SQL job on Streaming SQL Console, go back to Cloudera Manager to stop the YARN session.
5. 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 applications are displayed.
6. Select the application you need to stop.

7. Click Settings.
8. Click Kill application.

Upgrading CSA artifacts for SQL Stream Builder service

You need to upgrade the Cloudera Streaming Analytics (CSA) artifacts to the newer version by manually deleting the old CSD file from the repository, deactivating the old parcel and installing the latest CSD file and parcel.

Procedure

1. Stop the SQL Stream Builder service on your cluster.
 - a) Go to your cluster in Cloudera Manager.
 - b) Select SQL Stream Builder from the list of services.
 - c) Click Action > Stop.
2. Download the new CSD files from the [Download location](#).
3. Delete the old CSD version from /opt/cloudera/csd.
 - a) Access the host where the CSD is stored.

```
ssh root@<hostname>
cd /opt/cloudera/csd
rm FLINK-1.14.0-csa1.6.1.0-cdh7.1.7.0-551-20897251.jar
rm SQL_STREAM_BUILDER-1.14.0-csa1.6.1.0-cdh7.1.7.0-551-20897251.jar
```

4. Copy the new artifacts to the /opt/cloudera/csd folder.



Warning: The upgrade process fails if you only add the new parcel file in Cloudera Manager. You must add the Flink and SSB CSD to the /opt/cloudera/csd folder as well.

```
scp <download_location>/FLINK-1.14.0-csa1.6.2.0-cdh7.1.7.0-551-23013538.jar root@<hostname>/opt/cloudera/csd:
scp <location>/SQL_STREAM_BUILDER-1.14.0-csa1.6.2.0-cdh7.1.7.0-551-23013538.jar root@<hostname>/opt/cloudera/csd:
```

5. Ensure that you maintain the appropriate file ownership and access the attributes if needed.

```
chown cloudera-scm:cloudera-scm /opt/cloudera/csd/FLINK-1.14.0-csa1.6.2.0-cdh7.1.7.0-551-23013538.jar
chown cloudera-scm:cloudera-scm /opt/cloudera/csd/SQL_STREAM_BUILDER-1.14.0-csa1.6.2.0-cdh7.1.7.0-551-23013538.jar
```

6. Restart the Cloudera Manager Server with the following command:
 - RHEL 7 Compatible, SLES 12, Ubuntu: `systemctl restart cloudera-scm-server`
 - RHEL 6 Compatible: `service cloudera-scm-server restart`
7. Restart the Cloudera Management Services.
8. Select Parcels on the **Home > Hosts** tab in the Cloudera Manager main navigation bar.
9. Click on Parcel Repositories & Network Settings tab.
10. Add the new Remote Parcel Repository URL for CSA.

```
https://user:password@archive.cloudera.com/p/csa/1.6.2.0/parcels/
```



Note: Make sure that the Remote Parcel Repository URL uses HTTPS link. To install a different version of the parcel, you can change the URL as needed.

11. Click Save & Verify Configuration to commit the change.

12. Click Close.

You are redirected to the Parcels page.

13. Deactivate the old version of the parcel on the **Parcels** page.**14.** Download, activate and distribute the new version of the parcel.**15.** Select the "Activate Only (no restart)" option when prompted in the wizard.

For more information about the general process of Service Management, see the [Cloudera Manager documentation](#).

16. Stop the SQL Stream Builder service.

a) Click **Actions > Stop** next to the SQL Stream Builder service name.

17. Update the Admin Database for the SSB service.

a) Click **Actions > Update Admin Database** next to the SQL Stream Builder service name.

18. Start the SQL Stream Builder service.

a) Click **Actions > Start** next to the SQL Stream Builder service name.

19. Refresh your keytab in Streaming SQL Console.

a) Select SQL Stream Builder from the list of services.

b) Click SQLStreamBuilder Console.

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

c) Select your username on the left main menu of Streaming SQL Console.

d) Click Manage keytab.

e) Provide your password again.

f) Click Refresh Keytab.

Resuming SQL jobs

After upgrading the SQL Stream Builder service in Cloudera Manager, resume your SQL jobs by restarting them using the SQL jobs page of Streaming SQL Console.

Procedure

1. Navigate to the Streaming SQL Console.

a) Go to your cluster in Cloudera Manager.

b) Select SQL Stream Builder from the list of services.

c) Click SQLStreamBuilder Console.

2. Select SQL Jobs tab.

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

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

Compose Tables Functions History **SQL Jobs**

Search by name Running Jobs

ID	Name	User	Start Time	Link	State	Actions
5209	loving_wright	admin	2021-09-28 17:54:51 (30 minutes ago)	Flink Dashboard	RUNNING	Stop

Previous 1 Next

Details [Log](#)

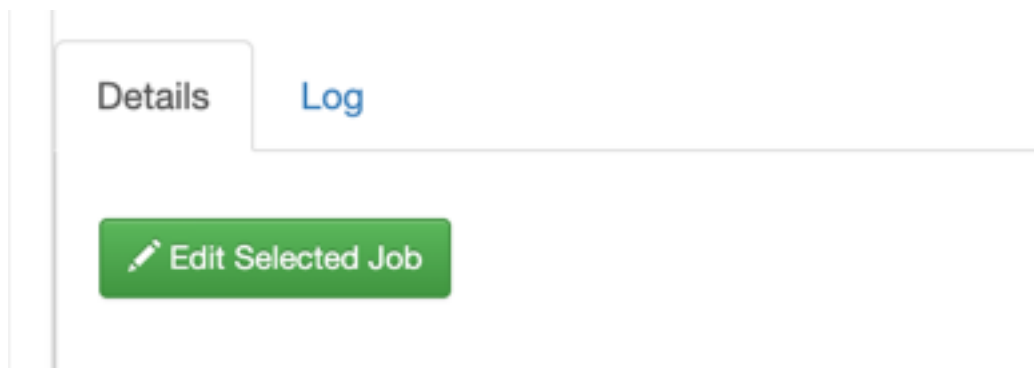
No Job Selected

Click on a Job and its details will display here

3. Click on the job you want to restart.

You can further filter down the job list by searching for the job name, or locate a specific job ID by prefixing your search with id.

4. Click Edit Selected Job under the **Details** tab.



5. Click on Execute.