

Cloudera Streaming Analytics 1.16.0

## Release Notes

**Date published:** 2019-12-17

**Date modified:** 2025-10-24

**CLOUDERA**

# Legal Notice

© Cloudera Inc. 2026. 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>What's new in Cloudera Streaming Analytics.....</b>	<b>4</b>
<b>Fixed issues.....</b>	<b>4</b>
<b>Deprecation notices in Cloudera Streaming Analytics 1.16.1.....</b>	<b>4</b>
Deprecation Notices for Cloudera SQL Stream Builder.....	5
<b>Known issues and limitations.....</b>	<b>5</b>
<b>Improvements and Behavioral changes.....</b>	<b>7</b>
<b>Unsupported features.....</b>	<b>7</b>
<b>Support Matrix.....</b>	<b>8</b>
Component support.....	8
System Requirements.....	9
Default ports for Flink and Cloudera SQL Stream Builder.....	10
Maven dependencies in Flink.....	11
Flink API Support.....	13

# What's new in Cloudera Streaming Analytics

Cloudera Streaming Analytics 1.16.1 is a maintenance release that focuses on enhancing stability, security, and usability.

## Cloudera platform support

Cloudera Streaming Analytics 1.16.1 is supported on Cloudera on premises 7.3.1.700 (SP3 CHF2). Ensure that you review the [7.3.1.700 \(SP3 CHF2\) Release Notes](#) and the [Support Matrix](#) to understand which operating systems, databases, and JDK versions are supported for Cloudera Streaming Analytics as well.

## Fixed issues

Review the list of Flink and Cloudera SQL Stream Builder issues that are resolved in Cloudera Streaming Analytics 1.16.1.

### **CSA-6032 - Cloudera SQL Stream Builder logs are not displayed in Cloudera Manager**

Log configurations for Streaming SQL Engine (SSE) and Materialized View Engine (MVE) services have been updated to ensure logs are correctly displayed in Cloudera Manager.

### **CSA-6023 - No foreign key constraint on udf table for artifact id**

A foreign key constraint has been added to the udfs table for the artifact ID to ensure data integrity.

### **CSA-5969 - Password placeholder resolution crashes when there are special characters in passwords**

Fixed an issue where special characters in passwords caused service startup failures. Special characters in password placeholders are now properly escaped.

### **CSA-5945 - Cloudera SQL Stream Builder does not check push result when exporting projects to Github**

Cloudera SQL Stream Builder now verifies the result of the push operation when exporting projects. Previously, a rejected push might not have been detected, resulting in a successful sync status despite the failure.

### **CSA-5873 - Cloudera SQL Stream Builder environment variables cannot be used through the V2 API**

Cloudera SQL Stream Builder environment variables are now correctly applied when validating data sources via the V2 API.

### **CSA-5743 - Invalid request error cannot instantiate User Defined Function after restart**

Fixed an issue where Python User Defined Functions (UDFs) failed to instantiate during query analysis. Python UDFs are now correctly configured for temporary sessions used during analysis.

### **CSA-6021 Validation of queries with functions erroneously invoked under SELECT fails with StackOverflow (FLINK-38750)**

Resolved a StackOverflow error that occurred during the validation of SQL queries where functions were invoked under the SELECT clause.

Apache Jira:[FLINK-38750](#)

## Deprecation notices in Cloudera Streaming Analytics 1.16.1

Certain features and functionalities have been removed or deprecated in Cloudera Streaming Analytics 1.16.1. You must review these items to understand whether you must modify your existing configuration. You can also learn about the features that will be removed or deprecated in the future release to plan for the required changes.

## Terminology

Items in this section are designated as follows:

### Deprecated

Technology that Cloudera is removing in a future Cloudera Streaming Analytics release. Marking an item as deprecated gives you time to plan for removal in a future Cloudera Streaming Analytics release.

### Moving

Technology that Cloudera is moving from a future Cloudera Streaming Analytics release and is making available through an alternative Cloudera offering or subscription. Marking an item as moving gives you time to plan for removal in a future Cloudera Streaming Analytics release and plan for the alternative Cloudera offering or subscription for the technology.

### Removed

Technology that Cloudera has removed from Cloudera Streaming Analytics and is no longer available or supported as of this release. Take note of technology marked as removed since it can potentially affect your upgrade plans.

## Deprecation Notices for Cloudera SQL Stream Builder

Certain features and functionality are deprecated or removed in Cloudera SQL Stream Builder. You must review these changes along with the information about the features in Cloudera SQL Stream Builder that will be removed or deprecated in a future release.

### Deprecated

#### Support for JavaScript UDFs

Due to the deprecation of the Nashorn engine used in JDK 8 and 11, User-Defined Functions (UDFs) written in JavaScript are deprecated in Cloudera Streaming Analytics. Cloudera recommends that you create new functionality with [Python UDFs](#) and start migrating existing JavaScript UDFs to Python.

Support in Python UDFs for table transformations and the webhook connector is available for Cloudera SQL Stream Builder and Flink.

### Removed

#### Cloudera SQL Stream Builder v1 REST API

The v1 REST API for Cloudera SQL Stream Builder is removed in Cloudera Streaming Analytics 1.16.1 and later releases. Migrate scripts and integrations to the v2 API.

See the [Cloudera SQL Stream Builder REST API reference](#) for v2 API details.

## Known issues and limitations

Learn about the known issues in Flink and Cloudera SQL Stream Builder, the impact or changes to the functionality, and the workaround in Cloudera Streaming Analytics 1.16.1.

### Cloudera SQL Stream Builder

#### CSA-6104 - Validating sync sources does not indicate when push is rejected

When using the API operation `/api/v2/projects/{projectId}/sync/config/validate`, the system performs a dry run push. If the push is rejected, the local branch remains up to date, and the API may return an `UP_TO_DATE` status instead of an error, even if the push failed.

Verify push results through the project export functionality, which now correctly processes push failures.

#### **CSA-5733 - Chart or diagram type dashboard widgets do not work when the label field is the same as the data field**

When creating a diagram type widget in Cloudera SQL Stream Builder, setting the label and data fields to the same value causes the graph to disappear.

None.

#### **CSA-5732 - MV widget not fetching mv data when authenticated via spnego**

When using Cloudera SQL Stream Builder with SPNEGO authentication, creating a Materialized View widget fails with a Data Source Error.

None. Users are advised to authenticate through KNOX.

### **Flink**

In Cloudera Streaming Analytics, the following SQL API features are in preview:

- Match recognize
- Top-N
- Stream-Table join (without rowtime input)

#### **DataStream conversion limitations**

- Converting between Tables and POJO DataStreams is currently not supported in Cloudera Streaming Analytics.
- Object arrays are not supported for Tuple conversion.
- The java.time class conversions for Tuple DataStreams are only supported by using explicit TypeInformation: LegacyInstantTypeInfo, LocalTimeTypeInfo.getInfoFor(LocalDate/LocalDate/LocalTime/LocalTime.class).
- Only java.sql.Timestamp is supported for rowtime conversion, java.time.LocalDateTime is not supported.

#### **Kudu catalog limitations**

- CREATE TABLE
  - Primary keys can only be set by the kudu.primary-key-columns property. Using the PRIMARY KEY constraint is not yet possible.
  - Range partitioning is not supported.
- When getting a table through the catalog, NOT NULL and PRIMARY KEY constraints are ignored. All columns are described as being nullable, and not being primary keys.
- Kudu tables cannot be altered through the catalog other than simply renaming them.

#### **Schema Registry catalog limitations**

- Currently, the Schema Registry catalog / format only supports reading messages with the latest enabled schema for any given Kafka topic at the time when the SQL query was compiled.
- No time-column and watermark support for Registry tables.
- No CREATE TABLE support. Schemas have to be registered directly in the SchemaRegistry to be accessible through the catalog.
- The catalog is read-only. It does not support table deletions or modifications.
- By default, it is assumed that Kafka message values contain the schema id as a prefix, because this is the default behaviour for the SchemaRegistry Kafka producer format. To consume messages with schema written in the header, the following property must be set for the Registry client: store.schema.version.id.in.header: true.

# Improvements and Behavioral changes

Learn about the change in certain functionality of Flink and Cloudera SQL Stream Builder that has resulted in a change in behavior from the previously released version of Cloudera Streaming Analytics.

## Improvements

### CSA-5889 - Support custom properties for Kafka data source

You can now add custom properties to Kafka data sources directly within the Cloudera SQL Stream Builder Console. This provides greater flexibility when configuring Kafka connections.

### CSA-5484 - Add Job Manager URL to SqlExecuteJobResponse when on YARN

The Flink Dashboard link in the Cloudera SQL Stream Builder Console now updates dynamically to point to the correct Job Manager URL when running on YARN. This resolves an issue where the link could point to an old, finished YARN job.

## Behavioral changes

### Summary:

CSA-5902 - Cross project Authorization bypass in Cloudera SQL Stream Builder

#### Previous behavior:

Regular users could potentially retrieve project permission information for user IDs other than their own or projects they were not members of.

#### New behavior:

The project-permissions API is now restricted. Regular users can only retrieve information for their own user ID and for projects of which they are a member. Admin users retain access to all user IDs. Additionally, the user property of ProjectPermissionDto returns a BasicUserView which does not expose sensitive information.

### Summary:

CSA-5998 - Update schema registry property name

#### Previous behavior:

The schema registry property name used a legacy format.

#### New behavior:

The schema registry property name has been updated to match the component name in Cloudera.

### Summary:

CSA-6023 - Add foreign key constraint on udfs table

#### Previous behavior:

N/A

#### New behavior:

A foreign key constraint has been added to the udfs table regarding artifact IDs. This database schema change ensures stricter data integrity for UDF artifacts.

# Unsupported features

Some Apache Flink and Cloudera SQL Stream Builder features exist in Cloudera Streaming Analytics 1.16.1, but are not supported by Cloudera. These features are not ready for production deployment, but Cloudera encourages you

to explore them in non-production environments and provide feedback on your experiences through the Cloudera Community Forums.

### Cloudera SQL Stream Builder

- Virtual environments for Python are not supported.

### Flink

- Apache Flink batch (DataSet) API
- GPU Resource Plugin
- SQL Client
- The following features are not supported in SQL and Table API:
  - HBase Table Connector
  - Old Planner
  - Non-windowed (unbounded) joins, distinct

## Support Matrix

Before installing Cloudera Streaming Analytics, review the supported components, databases, connectors and the default ports in use for Flink and Cloudera SQL Stream Builder.

### Component support

Learn more about which Apache Flink component version is supported in the Cloudera Streaming Analytics releases.

Cloudera Streaming Analytics version	Component version
Cloudera Streaming Analytics 1.16.1	Apache Flink 1.20.1
Cloudera Streaming Analytics 1.15.1	Apache Flink 1.20.0
Cloudera Streaming Analytics 1.15.0	Apache Flink 1.20.0
Cloudera Streaming Analytics 1.14.1	Apache Flink 1.19.1
Cloudera Streaming Analytics 1.14.0	Apache Flink 1.19.1
Cloudera Streaming Analytics 1.13.0	Apache Flink 1.19.1
Cloudera Streaming Analytics 1.12.0	Apache Flink 1.18.0
Cloudera Streaming Analytics 1.11.0	Apache Flink 1.16.2
Cloudera Streaming Analytics 1.10	Apache Flink 1.16.1
Cloudera Streaming Analytics 1.9.0	Apache Flink 1.15.1
Cloudera Streaming Analytics 1.8.0	
Cloudera Streaming Analytics 1.7.0	
Cloudera Streaming Analytics 1.6.2	
Cloudera Streaming Analytics 1.6.1	Apache Flink 1.14
Cloudera Streaming Analytics 1.6.0	
Cloudera Streaming Analytics 1.5.3	
Cloudera Streaming Analytics 1.5.1	Apache Flink 1.13
Cloudera Streaming Analytics 1.5.0	

Cloudera Streaming Analytics version	Component version
Cloudera Streaming Analytics 1.4.1	
Cloudera Streaming Analytics 1.4.0	Apache Flink 1.12
Cloudera Streaming Analytics 1.3.0	
Cloudera Streaming Analytics 1.2.0	Apache Flink 1.10
Cloudera Streaming Analytics 1.1.0	Apache Flink 1.9.1

### Related Information

[Cloudera Streaming Analytics 1.16.1 Release Notes](#)  
[Cloudera Streaming Analytics 1.16.0 Release Notes](#)  
[Cloudera Streaming Analytics 1.15.1 Release Notes](#)  
[Cloudera Streaming Analytics 1.15.0 Release Notes](#)  
[Cloudera Streaming Analytics 1.14.1 Release Notes](#)  
[Cloudera Streaming Analytics 1.14.0 Release Notes](#)  
[Cloudera Streaming Analytics 1.13.2 Release Notes](#)  
[Cloudera Streaming Analytics 1.13.1 Release Notes](#)  
[Cloudera Streaming Analytics 1.13.0 Release Notes](#)  
[Cloudera Streaming Analytics 1.12.0 Release Notes](#)  
[Cloudera Streaming Analytics 1.11.0 Release Notes](#)  
[Cloudera Streaming Analytics 1.10.0 Release Notes](#)  
[Cloudera Streaming Analytics 1.9.0 Release Notes](#)  
[Cloudera Streaming Analytics 1.8.0 Release Notes](#)  
[Cloudera Streaming Analytics 1.7.0 Release Notes](#)  
[Cloudera Streaming Analytics 1.6.2 Release Notes](#)  
[Cloudera Streaming Analytics 1.6.1 Release Notes](#)  
[Cloudera Streaming Analytics 1.6.0 Release Notes](#)  
[Cloudera Streaming Analytics 1.5.3 Release Notes](#)  
[Cloudera Streaming Analytics 1.5.1 Release Notes](#)  
[Cloudera Streaming Analytics 1.5.0 Release Notes](#)  
[Cloudera Streaming Analytics 1.4.1 Release Notes](#)  
[Cloudera Streaming Analytics 1.4.0 Release Notes](#)  
[Cloudera Streaming Analytics 1.3.0 Release Notes](#)  
[Cloudera Streaming Analytics 1.2.0 Release Notes](#)  
[Cloudera Streaming Analytics 1.1.0 Release Notes](#)

## System Requirements

Before installing Cloudera Streaming Analytics, you should verify that you meet the system requirements. Other than Cloudera Base on premises, you should also check the latest supported version of the needed components.

For detailed information about the supported versions of Cloudera Base on premises, operating systems and databases, see the [Cloudera Support Matrix](#).

Apache Flink support	1.20.1
Cloudera Runtime component support in Cloudera Base on premises 7.3.1	
Atlas	3.0.0
HBase	2.4.17
HDFS	3.1.1

Apache Flink support	1.20.1
<b>Cloudera Runtime component support in Cloudera Base on premises 7.3.1</b>	
Hive	3.1.3
Kafka <sup>1</sup>	3.4.1
Kudu	1.17.0
Schema Registry	0.10.0
Streams Messaging Manager	2.3.0
Apache Iceberg	1.5.2
<b>Connector support</b>	
JDBC PostgreSQL	9.6-16
JDBC MySQL	5.7, 8
JDBC Hive	3.1.3
JDBC Oracle	19, 19c, 21c, 23ai
JDBC Db2	11.5
JDBC SQL Server	2007-2022
CDC PostgreSQL	9.6-16
CDC MySQL	5.7, 8
CDC Oracle	19, 19c, 21c, 23ai
CDC Db2	11.5
CDC SQL Server	2007-2022
Apache Iceberg	1.3.0

## Default ports for Flink and Cloudera SQL Stream Builder

You need to use the default ports of Flink and Cloudera SQL Stream Builder when you need to reach or connect to their services. The default port are set in Cloudera Manager, but can be changed if required.

The following table lists the default ports and the corresponding property file names for Flink and Cloudera SQL Stream Builder. The ports are set by default in Cloudera Manager. You can change the ports as required using the configuration properties.

Component	Service	Port	Configuration property
Flink	Flink Dashboard	18211	historyserver.web.port
Cloudera SQL Stream Builder	Streaming SQL Engine	18121	server.port
	Materialized View Engine	18131	server.port
Cloudera SQL Stream Builder with Load Balancer	Streaming SQL Engine	8080	ssb.sse.loadbalancer.server.port
	Secured Streaming SQL Engine	8445	ssb.sse.loadbalancer.server.secure.port
	Materialized View Engine	8081	ssb.mve.loadbalancer.server.port
	Secured Materialized View Engine	8444	ssb.mve.loadbalancer.server.secure.port

<sup>1</sup> Connecting to Kafka that is running on remote CDH6 or HDP3 is also supported.

For the default port list of the Cloudera Runtime components, see the *Ports Used by Cloudera Runtime Components* document.

## Maven dependencies in Flink

Review the list of Maven dependencies to ensure the correct connector versions in your Flink applications.

### Avro

```
<dependency>
  <groupId>org.apache.flink</groupId>
  <artifactId>flink-avro</artifactId>
  <version>1.20.1-csa1.16.1.0</version>
</dependency>
```

### Confluent Registry

```
<dependency>
  <groupId>org.apache.flink</groupId>
  <artifactId>flink-connector-confluent-registry</artifactId>
  <version>1.0-csa1.15.0.0</version>
</dependency>
```

### CSV

```
<dependency>
  <groupId>org.apache.flink</groupId>
  <artifactId>flink-csv</artifactId>
  <version>1.20.1-csa1.16.1.0</version>
</dependency>
```

### Hive

```
<dependency>
  <groupId>org.apache.flink</groupId>
  <artifactId>flink-connector-hive_2.12</artifactId>
  <version>1.20.1-csa1.16.1.0</version>
</dependency>
```

### HBase

```
<dependency>
  <groupId>org.apache.flink</groupId>
  <artifactId>flink-connector-hbase-1.4</artifactId>
  <version>3.0-csa1.15.0.0</version>
</dependency>
```

```
<dependency>
  <groupId>org.apache.flink</groupId>
  <artifactId>flink-connector-hbase-2.4</artifactId>
  <version>4.0-csa1.16.1.0</version>
</dependency>
```

### Iceberg

```
<dependency>
  <groupId>org.apache.iceberg</groupId>
  <artifactId>iceberg-flink-runtime-1.16</artifactId>
  <version>1.5.2.7.3.1.500-182</version>
</dependency>
```

```
</dependency>
```

## JDBC

```
<dependency>
  <groupId>org.apache.flink</groupId>
  <artifactId>flink-connector-jdbc</artifactId>
  <version>3.2-csa1.15.0.0</version>
</dependency>
```

## JSON

```
<dependency>
  <groupId>org.apache.flink</groupId>
  <artifactId>flink-json</artifactId>
  <version>1.20.1-csa1.16.1.0</version>
</dependency>
```

## Kafka

```
<dependency>
  <groupId>org.apache.flink</groupId>
  <artifactId>flink-connector-kafka</artifactId>
  <version>3.3-csa1.15.0.0</version>
</dependency>
```

## Kudu

```
<dependency>
  <groupId>org.apache.bahir</groupId>
  <artifactId>flink-connector-kudu_2.12</artifactId>
  <version>1.1.0-csa1.15.0.0</version>
</dependency>
```

## Schema Registry

```
<dependency>
  <groupId>org.apache.flink</groupId>
  <artifactId>flink-connector-cloudera-registry</artifactId>
  <version>1.0-csa1.15.0.0</version>
</dependency>
```

## Table API

```
<dependency>
  <groupId>org.apache.flink</groupId>
  <artifactId>flink-table-api-java-bridge</artifactId>
  <version>1.20.1-csa1.16.1.0</version>
</dependency>
<dependency>
  <groupId>org.apache.flink</groupId>
  <artifactId>flink-table-planner-loader-bundle</artifactId>
  <version>1.20.1-csa1.16.1.0</version>
</dependency>
```

For more information about how to use Maven in Flink, see the [Apache documentation](#).

## Flink API Support

Cloudera Streaming Analytics offers support for three fundamental layers of the Apache Flink API. You can use DataStream API, the ProcessFunction API and a selected subset of the SQL API to develop your Flink streaming applications.

From the DataStream and ProcessFunction APIs, the following are supported based on the support annotations provided by the Apache Flink community.

Stable (@Public)	Evolving (@PublicEvolving)
<ul style="list-style-type: none"><li>DataStream API</li></ul>	<ul style="list-style-type: none"><li>ProcessFunction</li><li>Stream Join</li><li>Interval Join</li><li>Stateful operators</li><li>FsStatebackend with HDFS</li><li> RocksDBStateBackend with HDFS</li></ul>



**Note:** Cloudera Streaming Analytics does not support batch processing (DataSet API).