

Release Notes

Date published: 2019-12-17

Date modified: 2019-12-17



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

What's new in Cloudera Streaming Analytics.....	4
SQL and Table API supported features.....	4
Component support.....	5
Known issues and limitations.....	5
Unsupported features.....	6

What's new in Cloudera Streaming Analytics

Cloudera Streaming Analytics 1.2 covers new features beside the core streaming functionality of Apache Flink.

Cloudera Streaming Analytics 1.2 introduces the following key features:

- Upgrade to Apache Flink 1.10

For more information about Flink 1.10, see the [Release Notes](#) in the Apache Flink documentation.

- Reworked Dashboard to Flink WebUI
- Reworked Flink CLI
- Apache Atlas integration
- Apache Kudu streaming sink connector
- SQL/Table API and SQL client with the following connectors:
 - Apache Kudu
 - Schema Registry
 - Apache Hive
- Connecting to Kafka 2.4.1 that is running on remote CDH6 or HDP3 is also supported.

SQL and Table API supported features

The following SQL and Table API features are supported in Cloudera Streaming Analytics.

StreamTableEnvironment (Blink Planner)	Catalogs
<ul style="list-style-type: none"> • DataStream conversions: <ul style="list-style-type: none"> • Row • Tuple • Temporary views • SQL update and Query operations • Creating and using catalogs • Listing catalogs, tables, databases and functions 	<ul style="list-style-type: none"> • In-memory - Flink default catalog • Hive • Schema Registry • Kudu
SQL update statements	SQL query statements
<ul style="list-style-type: none"> • CREATE TABLE <ul style="list-style-type: none"> • Computed columns • Watermark definitions • Table connectors <ul style="list-style-type: none"> • Kafka • Kudu • Hive (through catalog) • Data formats (Kafka) <ul style="list-style-type: none"> • JSON • Avro • CSV • INSERT INTO <ul style="list-style-type: none"> • VALUES • SELECT queries 	<ul style="list-style-type: none"> • Basic operations <ul style="list-style-type: none"> • Scan, Project, Filter • Basic aggregations <ul style="list-style-type: none"> • Group By • Group By Window • Distinct (only windowed) • Joins <ul style="list-style-type: none"> • Time-windowed stream-stream join • User defined scalar functions (scalar udf)

Component support

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

CSA version	Component version
CSA 1.2.0	Apache Flink 1.10
CSA 1.1.0	Apache Flink 1.9.1

Related Information

[CSA 1.1.0 Release Notes](#)

Known issues and limitations

The following limitations apply to Cloudera Streaming Analytics 1.2.

- The SQL Client is offered as a preview for user convenience, but in itself is not supported.
- 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 CSA.
- 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/LocalDateTime/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.

Unsupported features

The following features are not supported in Cloudera Streaming Analytics 1.2.

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