Cloudera Streaming Analytics

Support Matrix

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

CLOUDERA

Legal Notice

Contents

System Requirements

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

Cloudera platform support	CSA 1.2 runs on CDP Private Cloud Base 7.1.1-7.1.5.	
Operating system support	RHEL/CentOS 7.6, 7.7, 7.8	
	Ubuntu 18.04	
JDK support	Open JDK 8	
	Oracle JDK 8	
	Open JDK 11	
Component support	Flink	1.10
Component support for connectors	Kafka[1]	2.4.1
	Schema Registry	0.8.1
	Streams Messaging Manager	2.1.0
	HBase	2.2.3
	HDFS	3.1.1
	Atlas	2.1.0
	Kudu	1.13.0
	Hive	3.1.3

Flink API Support

Cloudera Streaming Analytics (CSA) 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)
• DataStream API	• ProcessFunction • Stream Join • Interval Join • Stateful operators • FsStatebackend with HDFS • RocksDBStateBackend with HDFS

Note: CSA does not support batch processing (DataSet API).

In the CSA 1.2 release, support for a subset of the SQL API is introduced. Cloudera encourages users to interact with the SQL API through the SQL client or embedding it into DataStream applications. However, Cloudera only provides production support for the latter method due to limitations of the SQL client. For a complete list of supported features in the SQL API, see the Release Notes.

---

[1] Connecting to Kafka that is running on remote CDH6 or HDP3 is also supported.

Related Information

SQL Client

DataStream applications