

Cloudera Stream Processing 2.0.0

Cloudera Stream Processing Release Notes

Date published: 2019-09-13

Date modified: 2019-09-13

The Cloudera logo is rendered in a bold, orange, sans-serif font. The word "CLOUDERA" is written in all caps. The letter 'E' is stylized with a horizontal bar that extends to the right and then turns downwards, creating a unique graphic element.

<https://docs.cloudera.com/>

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 Stream Processing.....	4
Component support.....	5
Known issues.....	5
Download locations.....	6

What's new in Cloudera Stream Processing

Cloudera Stream Processing (CSP) provides advanced messaging, stream processing and analytics capabilities powered by Apache Kafka as the core stream processing engine. It also provides streams management capabilities.

CSP consists of the following components and features:

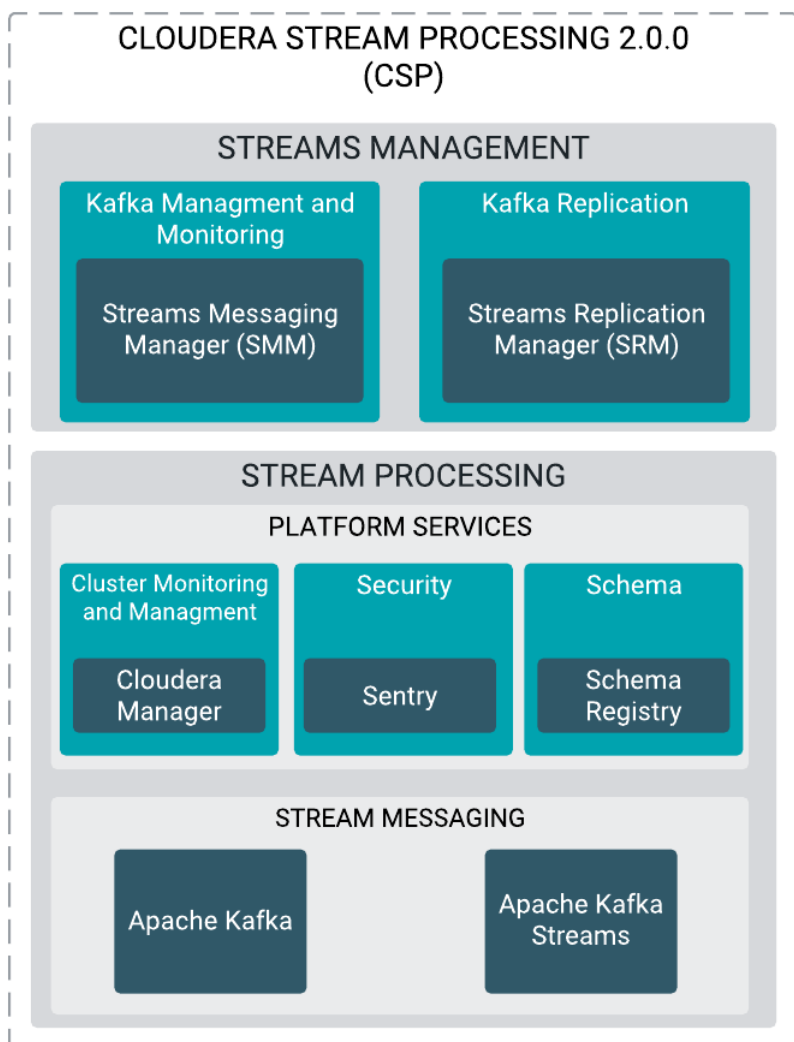
- New Kafka monitoring services powered by Streams Messaging Manager (SMM) delivered via a parcel
- New Kafka replication services powered by Streams Replication Manager (SRM) delivered via parcel
- A Schema Registry parcel that provides a centralized registry to manage schema for Apache Kafka and NiFi
- Support for Apache Kafka 2.2.0 delivered through CDH 6.3.0 or CDK 4.1



Note:

It is strongly recommended that you use CDH 6.3.0 or later with Apache Kafka included. If you are using CDH 5.13.0, you may use CDK 4.1.0 or later.

- Support for Kafka Streams
- Cluster Management and monitoring with Cloudera Manager
- Rich CDH cluster access control with Apache Sentry



Component support

CSP 2.0.0 includes the following components.

- Schema Registry 0.8.0
- Apache Kafka 2.2.0 delivered in the following two ways:
 - If you install CDH 6.3.0 or later, Kafka is included in your CDH installation.
 - If you are using CDH 5.13.0 or later, you must install Kafka using the CDK 4.1 parcel.
- Streams Messaging Manager 2.1.0
- Streams Replication Manager 1.0.0

Known issues

A summary of known issues for this version of Cloudera Stream Processing.

SMM does not trigger alert immediately when SRM goes down

Problem: If the SRM cluster goes down, SMM might take 15 minutes to send the alert through email or display the alert on the UI.

First set of produced messages are not counted in total messages displayed in SMM

Problem: You might notice minor discrepancy in the 'BytesIn', 'BytesOut', and 'MessagesIn' timeline metrics for the newly created topics and in the 'outMessagesCount' timeline metric for the producers when the time period is selected all the way up to topic initialisation or usage time. The discrepancy in the metrics may be off by the first minute's measurement.

The discrepancy happens due to a delay in the initialisation of metrics in the Kafka Brokers and by default the Metric reporter is configured to report the metrics once in a minute to the external systems post topic initialisation or usage.

Workaround: The graphs in the topic and producer profile page, however, reflect the true state of the metrics the way they are recorded.

SRM does not sync re-created source topics until the offsets have caught up with target topic

Problem: Messages written to topics that were deleted and re-created are not replicated until the source topic reaches the same offset as the target topic. For example, if at the time of deletion and re-creation there are a 100 messages on the source and target clusters, new messages will only get replicated once the re-created source topic has 100 messages. This leads to messages being lost.

Workaround: N/A.

SRM may automatically re-create deleted topics

Problem: If `auto.create.topics.enable` is enabled, deleted topics are automatically recreated on source clusters.

Workaround: Prior to deletion, remove the topic from the topic allowlist (whitelist) with the `srm-control` tool. This prevents topics from being re-created.

```
srm-control topics --source [SOURCE_CLUSTER] --target [TARGET_CLUSTER] --remove [TOPIC1][TOPIC2]
```

SRM cannot replicate Sentry authorization rules from one Kafka cluster to another

Problem: When Sentry contains Kafka authorization policies for any ConsumerGroup resource, SRM cannot replicate authorization rules from one Kafka cluster to another in environments where Sentry is enabled. This is due to a Kafka resource conversion error in Sentry. For more information

regarding the underlying Sentry issue, see [SENTRY-2535](#) and the Kafka known issues in [CDH Known Issues](#).

Workaround: Disable authorization policy synchronization in SRM. This can be achieved by setting the `sync.topic.acls.enabled` property to false.

SMM throws error if Service Monitor is not on CM host

Problem: If the Cloudera Manager Service Monitor and Cloudera Manager Server are deployed on different hosts, SMM is unable to fetch metrics correctly. As a result, historic data for consumer offsets and lag are not displayed, only the latest data is available.

You receive the following error message and cause in the `streams-messaging-manager.log`:

```
ERROR
com.hortonworks.smm.kafka.services.clients.ConsumerGroupManagementTask: Error while fetching consumer group information.
Caused by: org.apache.avro.AvroRemoteException: java.net.ConnectException: Connection refused (Connection refused)
```

Workaround: No workaround.

Cannot read property 'indexOf' thrown when ETELatencyMetrics is disabled

Problem: If latency metrics is disabled and you click the latency tab on the Topic Profile page in the SMM UI, then the following error appears:

```
Cannot read property 'indexOf' of null.
```

Workaround: Latency metrics should be configured and enabled before you view it on the SMM UI.

Download locations

The Cloudera Stream Processing (CSP) software artifacts can be downloaded from the Cloudera Archive. To access the CSP artifacts, you must have an active subscription agreement and a license key file along with the required authentication credentials (username and password).

Download instructions

The license key file and authentication credentials are provided in an email sent to customer accounts from Cloudera as part of the subscription fulfillment process.



Important: You must have your download credentials to access the CSP files in the <https://archive.cloudera.com/p/csp/> repository.

Your authentication credentials are not the same as the credentials you use to access the support portal.

In case you did not receive an email or do not have access to your authentication credentials, you can retrieve them in the following ways:

- Contact your Cloudera sales representative.
- File a non-technical case within the [Cloudera support portal](#) for our Support team to assist you.
- Use the CDP Private Cloud Credential Generator (requires a license key):
 1. Log in to the account associated with your license and subscription agreement on [cloudera.com](#).
 2. Go to the [CDP Private Cloud Base Downloads](#) page.
 3. Click Download Now.
 4. Enter the text of the “PGP Signed Message” within your license key in the Generate Credentials text box.
 5. Click Generate Credentials.

The Credential Generator returns your username and password.

After you have obtained your credentials, you can use them to access the software artifacts on the Cloudera Archive.

Download locations

You can directly click the links provided in the following tables and provide your credentials through a browser prompt.

Table 1: CSD files

File	Location
Schema Registry	https://archive.cloudera.com/p/csp/2.0.0.0/csd/SCHEMAREGISTRY-0.8.0.jar
Streams Messaging Manager	https://archive.cloudera.com/p/csp/2.0.0.0/csd/STREAMS_MESSAGING_MANAGER-2.1.0.jar
Streams Replication Manager	https://archive.cloudera.com/p/csp/2.0.0.0/csd/STREAMS_REPLICATION_MANAGER-1.0.0.jar

Table 2: Manifest file

File	Location
CSD Manifest	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/manifest.json

Table 3: Centos OS parcel files

File	Location
Schema Registry parcel file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/SCHEMAREGISTRY-0.8.0.2.0.0.0-135-el7.parcel
Schema Registry parcel sha file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/SCHEMAREGISTRY-0.8.0.2.0.0.0-135-el7.parcel.sha
Streams Messaging Manager parcel file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/STREAMS_MESSAGING_MANAGER-2.1.0.2.0.0.0-135-el7.parcel
Streams Messaging Manager parcel sha file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/STREAMS_MESSAGING_MANAGER-2.1.0.2.0.0.0-135-el7.parcel.sha
Streams Replication Manager parcel file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/STREAMS_REPLICATION_MANAGER-1.0.0.2.0.0.0-135-el7.parcel
Streams Replication Manager parcel sha file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/STREAMS_REPLICATION_MANAGER-1.0.0.2.0.0.0-135-el7.parcel.sha

Table 4: SLES 12 parcel files

File	Location
Schema Registry parcel file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/SCHEMAREGISTRY-0.8.0.2.0.0.0-135-sles12.parcel
Schema Registry parcel sha file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/SCHEMAREGISTRY-0.8.0.2.0.0.0-135-sles12.parcel.sha
Streams Messaging Manager parcel file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/STREAMS_MESSAGING_MANAGER-2.1.0.2.0.0.0-135-sles12.parcel
Streams Messaging Manager parcel sha file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/STREAMS_MESSAGING_MANAGER-2.1.0.2.0.0.0-135-sles12.parcel.sha
Streams Replication Manager parcel file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/STREAMS_REPLICATION_MANAGER-1.0.0.2.0.0.0-135-sles12.parcel
Streams Replication Manager parcel sha file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/STREAMS_REPLICATION_MANAGER-1.0.0.2.0.0.0-135-sles12.parcel.sha

Table 5: Ubuntu 18 parcel files

File	Location
Schema Registry parcel file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/SCHEMAREGISTRY-0.8.0.2.0.0.0-135-bionic.parcel

File	Location
Schema Registry parcel sha file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/SCHEMAREGISTRY-0.8.0.2.0.0.0-135-bionic.parcel.sha
Streams Messaging Manager parcel file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/STREAMS_MESSAGING_MANAGER-2.1.0.2.0.0.0-135-bionic.parcel
Streams Messaging Manager parcel sha file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/STREAMS_MESSAGING_MANAGER-2.1.0.2.0.0.0-135-bionic.parcel.sha
Streams Replication Manager parcel file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/STREAMS_REPLICATION_MANAGER-1.0.0.2.0.0.0-135-bionic.parcel
Streams Replication Manager parcel sha file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/STREAMS_REPLICATION_MANAGER-1.0.0.2.0.0.0-135-bionic.parcel.sha

Table 6: Ubuntu 16 parcel files

File	Location
Schema Registry parcel file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/SCHEMAREGISTRY-0.8.0.2.0.0.0-135-xenial.parcel
Schema Registry parcel sha file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/SCHEMAREGISTRY-0.8.0.2.0.0.0-135-xenial.parcel.sha
Streams Messaging Manager parcel file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/STREAMS_MESSAGING_MANAGER-2.1.0.2.0.0.0-135-xenial.parcel
Streams Messaging Manager parcel sha file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/STREAMS_MESSAGING_MANAGER-2.1.0.2.0.0.0-135-xenial.parcel.sha
Streams Replication Manager parcel file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/STREAMS_REPLICATION_MANAGER-1.0.0.2.0.0.0-135-xenial.parcel
Streams Replication Manager parcel sha file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/STREAMS_REPLICATION_MANAGER-1.0.0.2.0.0.0-135-xenial.parcel.sha

Table 7: Ubuntu 14 parcel files

File	Location
Schema Registry parcel file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/SCHEMAREGISTRY-0.8.0.2.0.0.0-135-trusty.parcel
Schema Registry parcel sha file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/SCHEMAREGISTRY-0.8.0.2.0.0.0-135-trusty.parcel.sha
Streams Messaging Manager parcel file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/STREAMS_MESSAGING_MANAGER-2.1.0.2.0.0.0-135-trusty.parcel
Streams Messaging Manager parcel sha file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/STREAMS_MESSAGING_MANAGER-2.1.0.2.0.0.0-135-trusty.parcel.sha
Streams Replication Manager parcel file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/STREAMS_REPLICATION_MANAGER-1.0.0.2.0.0.0-135-trusty.parcel
Streams Replication Manager parcel sha file	https://archive.cloudera.com/p/csp/2.0.0.0/parcels/STREAMS_REPLICATION_MANAGER-1.0.0.2.0.0.0-135-trusty.parcel.sha