

Cloudera Streaming Analytics Operator for Kubernetes 1.6.1 Release Notes

Date published:

Date modified:

The Cloudera logo, featuring the word "CLOUDERA" in a bold, orange, sans-serif font. The letter "E" is stylized with a horizontal bar through its center.

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

Cloudera Streaming Analytics Operator for Kubernetes 1.6.1.....	4
What's new in Cloudera Streaming Analytics Operator for Kubernetes 1.6.1.....	4
Fixed issues in Cloudera Streaming Analytics Operator for Kubernetes 1.6.1.....	4
Cloudera Streaming Analytics Operator for Kubernetes 1.6.0.....	4
What's new.....	5
Fixed issues and improvements.....	6
Improvements.....	6
Behavioral changes.....	6
Known issues and limitations.....	7
Unsupported features.....	8
Deprecation notices in Cloudera Streaming Analytics Operator for Kubernetes.....	8
Deprecation Notices for Cloudera SQL Stream Builder.....	9
Component support.....	9
System requirements.....	11

Cloudera Streaming Analytics Operator for Kubernetes 1.6.1

Release notes for Cloudera Streaming Analytics Operator for Kubernetes 1.6.1.

This section documents fixes and improvements for the 1.6.1 patch release. For the initial Cloudera Streaming Analytics Operator for Kubernetes 1.6.0 general availability content, see the release notes topics that follow this section in the table of contents.

What's new in Cloudera Streaming Analytics Operator for Kubernetes 1.6.1

Cloudera Streaming Analytics Operator for Kubernetes 1.6.1 is a patch release that fixes LDAP login, SCRAM authentication, and connector registration issues, and improves Kafka schema detection.

CSADS-491: Kafka schema detection without a consumer group

Previously, Cloudera SQL Stream Builder joined a consumer group when it detected a schema from a Kafka topic. Schema detection now uses topic assignment instead, which simplifies Kafka authorization in environments with strict consumer group policies.

Fixed issues in Cloudera Streaming Analytics Operator for Kubernetes 1.6.1

Review the Cloudera SQL Stream Builder issues resolved in Cloudera Streaming Analytics Operator for Kubernetes 1.6.1.

CSADS-474: LDAP login session cookie not created

Previously, LDAP login did not create a session cookie after successful authentication because Spring Security no longer configured the session automatically. LDAP login now completes with a valid session. This issue is now fixed.

CSADS-490: SCRAM login module not found on Flink jobs

Previously, Kafka jobs that use `org.apache.kafka.common.security.scram.ScramLoginModule` failed because the login module was not translated to the shaded class used in Flink job deployments. SCRAM authentication now works for deployed jobs. This issue is now fixed.

CSADS-492: New connector missing from Connectors list

Previously, after you added a connector, it did not appear under Connectors in Cloudera SQL Stream Builder. New connectors now appear in the list as expected. This issue is now fixed.

Cloudera Streaming Analytics Operator for Kubernetes 1.6.0

Learn about the new features, improvements, known and fixed issues, limitations, and unsupported features in this release of Cloudera Streaming Analytics Operator for Kubernetes.

This section documents the initial 1.6.0 general availability release. For the latest Cloudera Streaming Analytics Operator for Kubernetes 1.6.1 release notes, see the topics above this section in the table of contents.

What's new

This release of Cloudera Streaming Analytics Operator for Kubernetes adds Flink 1.20.5, Gateway API HTTPRoute exposure, external PostgreSQL credentials, install-time sampling Kafka settings, project file upload, PostgreSQL 18.4, and FIPS deployment guidance.

Cloudera Streaming Analytics Operator for Kubernetes 1.6.0

This release of Cloudera Streaming Analytics Operator for Kubernetes is based on Apache Flink Kubernetes Operator 1.13 and Apache Flink 1.20.5.

- [Flink Kubernetes Operator 1.13 release announcement](#)
- [Flink 1.20 release announcement](#)

Apache Flink upgraded to 1.20.5

The embedded Flink runtime and Cloudera SQL Stream Builder images are updated to Apache Flink 1.20.5, including upstream stability and connector fixes.

Gateway API HTTPRoute for Cloudera SQL Stream Builder

You can expose Cloudera SQL Stream Builder with a Kubernetes Gateway API HTTPRoute in addition to the existing Ingress resource. Configure `ssb.httpRoute` in the Helm chart. For details, see [Security configurations](#).

External and cross-namespace PostgreSQL credentials

The Helm chart supports reading database credentials from an existing Kubernetes Secret, including Secrets managed in another namespace (for example CloudNativePG). Use `ssb.database.auth.secretRef`. For details, see [Database configuration](#).

Sampling Kafka configuration at install time

You can supply sampling Kafka connection settings through a Kubernetes Secret referenced by `ssb.sampling.secretRef` during Helm install. This disables editing those settings in the Cloudera SQL Stream Builder UI. For details, see [SSB service configurations](#).

Default PostgreSQL version updated to 18.4

The default hardened PostgreSQL image deployed with Cloudera SQL Stream Builder is now version 18.4. For upgrade and air-gapped mirroring details, see [Improvements](#).

Project file upload

You can upload certificates, connector JARs, configuration directories, and other artifacts as project files in Cloudera SQL Stream Builder and reference them from Flink SQL DDL and connector properties. Configure `ssb.storageConfiguration` in the Helm chart for remote artifact storage. For procedures, see [Managing project files](#) and [Extensibility](#).

FIPS deployment guidance

Configuration steps for running Cloudera Streaming Analytics Operator for Kubernetes on FIPS-enabled OpenShift clusters are documented in [Configuring FIPS mode](#).

Related Information

[Python UDFs in Cloudera SQL Stream Builder | Cloudera Streaming Analytics Documentation](#)

[Adjusting logging configuration in Advanced Settings | Cloudera Streaming Analytics Documentation](#)

[Managing project files](#)

[Extensibility](#)

Fixed issues and improvements

Resolved defects in this release of Cloudera Streaming Analytics Operator for Kubernetes include Iceberg catalog S3 access, transient FlinkDeployment errors in Cloudera SQL Stream Builder, and Iceberg catalog configuration.

Fixed issues

CSADS-365: Iceberg catalog s3a ClassNotFoundException

Using an Iceberg catalog with S3 storage no longer fails with a ClassNotFoundException for the s3a filesystem when the alternative Flink image with Hadoop connectors is enabled.

CSADS-47: Cloudera SQL Stream Builder intermittently picked up transient FlinkDeployment errors

Previously, intermittent errors from the FlinkDeployment resource could be surfaced by Cloudera SQL Stream Builder, which could lead to application failure. This issue is fixed.

CSA-5960: Missing Iceberg catalog in default catalogs for Cloudera SQL Stream Builder

Cloudera SQL Stream Builder supports Iceberg as a first-class catalog type under Data Sources and Catalogs, including Hive-backed Iceberg catalog configuration. Iceberg tables are no longer presented under Hive catalogs, preventing access errors through the Hive connector.

Improvements

This release improves hardened PostgreSQL and JDK images, Helm chart PostgreSQL mirroring, optional Flink RBAC when the Flink Operator subchart is disabled, and Fernet key Secret handling for Cloudera Streaming Analytics Operator for Kubernetes.

CSADS-292: Chainguard FIPS-enabled JDK base images

Java component images now use the managed cloudera-openjdk:jdk-17-fips base image instead of a manually downloaded JDK. This improves FIPS readiness and removes the pinned Temurin JDK workaround from container builds.

Hardened default PostgreSQL image (18.4)

The bundled PostgreSQL image tag changes from 18.1 to a hardened 18.4 build. New installs use the updated image automatically. Upgrades reuse the existing PersistentVolumeClaim; review image pull settings in air-gapped environments.

CSADS-304: Public PostgreSQL image in release Helm charts

Release Helm charts reference the public hardened PostgreSQL image repository, which simplifies air-gapped mirroring for the bundled database image.

CSADS-169: Optional Flink RBAC when the Flink Operator subchart is disabled

When you install Cloudera SQL Stream Builder without the Flink Kubernetes Operator subchart, you can enable `ssb.flinkRbac.create` so Flink job pods still receive the required Role and RoleBinding in the Cloudera SQL Stream Builder namespace.

CSADS-247: Fernet key management

The Helm chart creates a Fernet key Secret only when `ssb.fernetKey` is set to a non-empty value. Clear the placeholder value so Cloudera SQL Stream Builder generates a unique key on first startup. For upgrade planning, see [Behavioral changes](#).

Behavioral changes

Fernet key Secret creation behaves differently in this release of Cloudera Streaming Analytics Operator for Kubernetes than in the previous release.

Fernet key Secret creation

The Helm chart no longer always pre-creates a Fernet key Secret. A Secret is rendered only when `ssb.fernetKey` is set to a non-empty string. The chart still ships a placeholder value; clear it (set to an empty string) before install or upgrade so Cloudera SQL Stream Builder generates a unique key on first startup.

If you previously relied on the chart-managed placeholder key, plan a coordinated upgrade: set an explicit key or clear the value and allow regeneration before rotating running jobs.

Known issues and limitations

Learn about the known issues and limitations in this release of Cloudera Streaming Analytics Operator for Kubernetes.

Stuck session jobs in Cloudera Streaming Analytics Operator for Kubernetes

Session jobs stop running if the session cluster's Job Manager is restarted without High Availability configured. However, because of a Flink bug, such stopped jobs get stuck in RECONCILING/STABLE state and cannot be restarted or deleted.

In such cases, the following is seen when using the `kubectrl get FlinkSessionJobs -n flink` command:

```
kubectrl get FlinkSessionJobs -n flink
NAME                                JOB STATUS  LIFECYCLE STATE
ssb-ssbdefault-testjobname         RECONCILING  STABLE
```

1. Remove the `FlinkDeployment` resource related to Cloudera SQL Stream Builder:

```
kubectrl -n [***NAMESPACE***] delete flinkdeployment [***DEPLOYMENT NAME***]
```

2. Open the Cloudera SQL Stream Builder configurations to edit:

```
kubectrl -n [***NAMESPACE***] edit cm ssb-config
```

3. Add the highlighted line to the `security-context.yaml` entry:

```
allowPrivilegeEscalation: false
capabilities:
  drop:
  - ALL
runAsNonRoot: true
seccompProfile:
  type: RuntimeDefault
runAsUser: 9999
```

It can take a couple of minutes for the changes to take effect after modifying the `ConfigMap`.

4. Execute the SQL job from Cloudera SQL Stream Builder again.

FLINK-33536: S3 filesystem sink and CSV format throws error

When using the Flink Table API CSV streaming sink with the S3 filesystem, the operation fails with `IOException: Stream closed`.

Cloudera SQL Stream Builder

CSADS-85: Data sampling might not start automatically after job reaches RUNNING

In some cases, data sampling does not start automatically when a job transitions to the `RUNNING` state.

Restart the sampler or restart the job manually.

Flink

DataStream conversion limitations

- Converting between Tables and POJO DataStreams is currently not supported in Cloudera Streaming Analytics Operator for Kubernetes.
- 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.

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 behavior 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

Learn what features are unsupported in this release of Cloudera Streaming Analytics Operator for Kubernetes.

Some Apache Flink and Cloudera SQL Stream Builder features exist in Cloudera Streaming Analytics Operator for Kubernetes, but are not supported by Cloudera. The following features are not ready for production deployment. Cloudera encourages you to explore these features in non-production environments and provide feedback on your experiences through the Cloudera Community Forums.

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

Cloudera SQL Stream Builder

- Virtual environments for Python are not supported
- Javascript UDFs are not supported

Deprecation notices in Cloudera Streaming Analytics Operator for Kubernetes

Certain features and functionalities have been removed or deprecated in Cloudera Streaming Analytics Operator for Kubernetes. 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 Operator for Kubernetes release. Marking an item as deprecated gives you time to plan for removal in a future Cloudera Streaming Analytics Operator for Kubernetes release.

Moving

Technology that Cloudera is moving from a future Cloudera Streaming Analytics Operator for Kubernetes 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 Operator for Kubernetes release and plan for the alternative Cloudera offering or subscription for the technology.

Removed

Technology that Cloudera has removed from Cloudera Streaming Analytics Operator for Kubernetes 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.

Removed

v1 REST API

The v1 REST API for Cloudera SQL Stream Builder is removed in Cloudera Streaming Analytics Operator for Kubernetes 1.4.0 and higher.

Cloudera recommends migrating to the v2 API.

For more information on the v2 API, see the [Cloudera SQL Stream Builder REST API reference](#).

Support for JavaScript UDFs

User-Defined Functions (UDFs) written in JavaScript are removed in Cloudera Streaming Analytics Operator for Kubernetes 1.4.0.

Cloudera recommends using Python UDFs and migrating JavaScript UDFs to Python. Support in Python UDFs for table transformations and the webhook connector is available for Cloudera SQL Stream Builder and Flink.

Component support

You can review the Cloudera Streaming Analytics Operator for Kubernetes components and their versions shipped in this release of the Cloudera Streaming Analytics Operator for Kubernetes.

Table 1: Cloudera Streaming Analytics Operator for Kubernetes component versions

Component	Version
Flink	1.20.5-csaop1.6.1-b6
Flink Operator	1.13-csaop1.6.1-b6
Cloudera SQL Stream Builder	1.20.5-csaop1.6.1-b6
PostgreSQL	18.4
OpenJDK	17

Supported Flink versions

Cloudera Streaming Analytics Operator for Kubernetes supports the following Flink versions:

Table 2: Supported Flink versions

Latest (default)	Other
1.20.5-csaop1.6.1-b6	None

The default version is the Cloudera-recommended current and latest supported Flink version. This version is used by default to deploy clusters if an explicit version is not provided in your `FlinkDeployment` resource.

Notice that the Flink versions are specific to Cloudera. Their version numbers consist of two parts: the first three digits specify the Apache Flink version, and the following specify the major and minor version of Cloudera Streaming Analytics Operator for Kubernetes. When deploying a cluster, you must use the Cloudera versions for Flink listed here. Specifying upstream versions is not supported.

Component support

Flink file system support

By default, Cloudera Streaming Analytics Operator for Kubernetes only supports local and locally mounted NFS/SAN file systems for Flink. Pluggable file systems (for example S3, HDFS, etc.) can be used by adding plugins to the Apache Flink Operator. For more information and a list of supported pluggable file systems, see [Using pluggable file systems](#).

Cloudera SQL Stream Builder database support

You can use the following databases with Cloudera SQL Stream Builder:

- MySQL/MariaDB
- Oracle
- PostgreSQL

Cloudera SQL Stream Builder connector support

With Cloudera SQL Stream Builder, you will get the following connectors installed by default:

- Kafka
- JDBC
- CDC (MySQL, Oracle, Postgres, Db2, SqlServer)
- Amazon S3
- Azure Blob Storage
- Google Cloud Storage
- HDFS (when using the alternative image, see [Installation overview](#))
- Hive (when using the alternative image, see [Installation overview](#))
- Iceberg (when using the alternative image, see [Installation overview](#))
- Kudu (when using the alternative image, see [Installation overview](#))

You will also get the following formats installed with Cloudera SQL Stream Builder:

- JSON
- Avro
- ORC
- Parquet

Related Information

[Using pluggable file systems](#)

System requirements

Cloudera Streaming Analytics Operator for Kubernetes requires a Kubernetes cluster environment that meets the provided system requirements and prerequisites. You must ensure to meet these requirements to be able to install and use the Cloudera Streaming Analytics Operator for Kubernetes or its components.

- A Kubernetes 1.25 or later cluster:
 - OpenShift 4.12 or later
 - Any Cloud Native Computing Foundation (CNCF) certified Kubernetes distribution. For more information, see cncf.io.
- Administrative rights on the Kubernetes cluster.
- Access to kubectl or oc. These command line tools must be configured to connect to your running cluster.
- Access to helm.
- Log collection is enabled for the Kubernetes cluster.

Cloudera requires that the logs of Cloudera Streaming Analytics Operator for Kubernetes components are stored long term for diagnostic and supportability purposes. Review [Log collection](#).