Cloudera Streaming Analytics Operator 1.0.0

CSA Operator Release Notes

Date published: 2024-06-15
Date modified: 2024-06-15

CLOUDERA

Legal Notice

Contents

Release Notes

Learn about the new features, improvements, known and fixed issues, limitations, and unsupported features in this release of CSA Operator.

What's new

Learn about the new features and notable changes in this release of CSA Operator.

CSA Operator 1.0.0

This is the initial release of CSA Operator.

To learn more about CSA Operator and its typical deployment architecture, see Overview and for the installation instructions, see Installation.

This release of the CSA Operator is based on the Apache Flink Kubernetes Operator (Flink Operator) 1.8.0 and Apache Flink (Flink) 1.18.1. For more information about the supported versions, see the following upstream resources:

- Flink Operator 1.8.0
- Flink 1.18.1

This version of CSA Operator also includes a Technical Preview of the SQL Stream Builder (SSB). SSB is a comprehensive interactive user interface for creating stateful stream processing jobs using SQL. For more information about SSB and its features, see the What is SQL Stream Builder page.

Known issues

Learn about the known issues in this release of CSA Operator.

SSB Flink job deployment can fail on non-OpenShift clusters

On non-OpenShift clusters, SSB Flink job deployment can fail with the following error:

```
 CreateContainerConfigError: container has runAsNonRoot and imag
e will run as root. A default non-root user is not set so the co
ntainer starts as root but this is forbidden by the security pol
icy.
```

1. Remove the `FlinkDeployment` resource related to SSB:

   ```
   kubectl -n [***NAMESPACE***] delete flinkdeployment [***DEPL
   OYMENT NAME***]
   ```

2. Open the SSB configurations to edit:

   ```
   kubectl -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 SSB again.

Unsupported features

Learn what features are unsupported in this release of CSA Operator.

The following upgrade scenarios are not supported in this version of CSA Operator:

• Upgrading from the Technical Preview version to the Generally Available version
• Upgrading from the Apache Flink Kubernetes Operator to CSA Operator

This version of the CSA Operator also includes the Technical Preview of SQL Stream Builder (SSB). The following SSB features are not included in this version of CSA Operator:

• Materialized Views
• Notifications

Component versions

You can review the CSA Operator components and their versions shipped in this release of theCSA Operator.

Table 1: CSA Operatorcomponent versions

Component	Version
Flink	1.18.1-csaop1.0.0-b317
Flink Operator	1.8-csaop1.0.0-b317
SQL Stream Builder	1.18.1-csaop1.0.0-b317

Supported Flink versions

CSA Operator supports the following Flink versions:

Table 2: Supported Flink versions

Latest (default)	Other
1.18.1-csaop1.0.0-b317	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 CSA Operator. 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, CSA Operator 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*.

SSB database support

You can use the following databases with SQL Stream Builder (SSB):

- MySQL/MariaDB
- Oracle
- PostgreSQL

SSB connector support

With SSB, 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

You will also get the following formats installed with SSB:

- JSON
- Avro
- ORC
- Parquet

Related Information
Using pluggable file systems

System requirements

CSA Operator 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 CSA Operator or its components.

- A Kubernetes 1.23 or later cluster:

  - OpenShift 4.10 or later.
- 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 CSA Operator components are stored long term for diagnostic and supportability purposes. Review Log collection.