

CSA Operator Installation

Date published: 2024-06-15

Date modified: 2024-02-28



Legal Notice

© Cloudera Inc. 2025. 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

Installation overview.....	4
Installing Cloudera Streaming Analytics - Kubernetes Operator in an environment with internet access.....	5
Installing Cloudera Streaming Analytics - Kubernetes Operator in an air- gapped environment.....	8

Installation overview

Get started with installing Cloudera Streaming Analytics - Kubernetes Operator. Learn about the concept of installing, the installation artifacts, and where these artifacts are hosted.

Cloudera Streaming Analytics - Kubernetes Operator is installed using a Helm chart, which installs the Apache Kubernetes Flink Operator (Flink Operator), Cloudera SQL Stream Builder, and other components.

Installation artifacts and artifact locations

Cloudera Streaming Analytics - Kubernetes Operator ships with various installation artifacts, hosted at two locations: the Cloudera Docker registry and the Cloudera Archive.

Both the Cloudera Docker registry and the Cloudera Archive require your Cloudera credentials (username and password) to access. Credentials are provided to you as part of your license and subscription agreement. You can access both the registry and the archive using the same credentials.

Cloudera Docker registry – container.repository.cloudera.com

The Cloudera Docker registry hosts the Helm chart as well as all Docker images used for the installation.

Table 1: Cloudera Streaming Analytics - Kubernetes Operator artifacts on the Cloudera Docker registry

Artifact	Location	Description
Flink Kubernetes Operator Docker image	container.repository.cloudera.com/cloudera/flink-kubernetes-operator:1.9-csaop1.2.2-b6	Docker image used for deploying the various operator components shipped with the Cloudera Streaming Analytics - Kubernetes Operator.
Flink Docker image	container.repository.cloudera.com/cloudera/flink:1.19.2-csaop1.2.2-b6	Docker image used for deploying Apache Flink and its related components.
SQL Runner Docker image	container.repository.cloudera.com/cloudera/flink-extended:1.19.2-csaop1.2.2-b6	Docker image used for deploying Flink application when SQL query is executed in SSB.
SQL Stream Engine Docker image	container.repository.cloudera.com/cloudera/ssb-sse:1.19.2-csaop1.2.2-b6	Docker image used for deploying Cloudera SQL Stream Builder and its UI.



Note: The images are built for linux/arm64 and linux/amd64 architectures.

Cloudera Archive – archive.cloudera.com/p/csa-operator/1.2.2/

The Cloudera Archive hosts various installation artifacts including the Helm chart, diagnostic tools, and the Maven artifacts.

All artifacts hosted in the Cloudera Archive are supplemental resources, accessing them is not required to complete the installation. The following table collects the Cloudera Streaming Analytics - Kubernetes Operator directories located in the Cloudera Archive, with an overview of what artifacts they contain and how you can use them:

Table 2: CSA Operator artifacts on the Cloudera Archive

Archive Directory	Description
archive.cloudera.com/p/csa-operator/1.2.2/ charts /	The charts directory contains the Helm chart. This is the same chart that is available in the Cloudera Docker registry. Cloudera recommends that whenever possible you install with the chart hosted in the registry. The chart in the archive is provided in case you cannot access the registry or want to download the chart using a browser.
archive.cloudera.com/p/csa-operator/1.2.2/ maven-repository /	The Maven artifacts can be used to develop your own applications or tools for use with Cloudera Streaming Analytics - Kubernetes Operator.
archive.cloudera.com/p/csa-operator/1.2.2/ tools /	The tools directory contains command line tools that you use to collect diagnostic information and to troubleshoot cluster issues.

Installing Cloudera Streaming Analytics - Kubernetes Operator in an environment with internet access

Complete these steps to install Cloudera Streaming Analytics - Kubernetes Operator if your Kubernetes cluster has internet access. Installing the Cloudera Streaming Analytics - Kubernetes Operator enables you to deploy and manage Flink and Cloudera SQL Stream Builder in Kubernetes.

About this task

Cloudera Streaming Analytics - Kubernetes Operator is installed in your Kubernetes cluster with the provided Helm chart through the helm install command. When you install the chart, Helm installs the Custom Resource Descriptors (CRD) included in Cloudera Streaming Analytics - Kubernetes Operator, and deploys the Apache Kubernetes Flink Operator (Flink Operator), Cloudera SQL Stream Builder engine, and a PostgreSQL database for Cloudera SQL Stream Builder.

Installing Cloudera Streaming Analytics - Kubernetes Operator does not create or deploy a Flink cluster. The Flink cluster is created after the installation by deploying the Flink Deployment resource in the Kubernetes cluster with kubectl or oc, or when you execute a SQL job in Streaming SQL Console.



Note: Cloudera recommends that you install Cloudera Streaming Analytics - Kubernetes Operator once per Kubernetes cluster.

By default the Flink Operator (deployed with installation) watches and manages all the Flink clusters that are deployed in the single same namespace as the Flink Operator. However, you can also configure it to watch and manage multiple namespaces. This allows you to manage multiple Flink clusters deployed in different namespaces, using a single Cloudera Streaming Analytics - Kubernetes Operator installation.



Important: Cloudera SQL Stream Builder can only be deployed in one namespace. In case you require multiple instances, you must install Cloudera SQL Stream Builder in every namespace. For more information about deploying Flink and Cloudera SQL Stream Builder in multiple namespaces, see the [Namespace management](#) documentation.

Before you begin

- Ensure that your Kubernetes environment meets requirements listed in [System requirements](#).

**Note:**

If you're installing Cloudera Streaming Analytics - Kubernetes Operator on a Longhorn storage provider (Rancher/RKE2) the following line has to be added to the Helm values.yaml file:

```
ssb:
  database:
    pod:
      securityContext:
        fsGroup: 999
```

- Your Kubernetes cluster requires internet connectivity to complete these steps, as it must be able to reach the Cloudera Docker registry.
- Ensure that you have access to a valid Cloudera license.
- Ensure that you have access to your Cloudera credentials (username and password, that are provided together with the Cloudera license) required to access the Cloudera Docker registry (and, if needed, the Cloudera Archive), where installation artifacts are hosted.
- Review the [Helm chart reference](#) before installation.

The Helm chart accepts various configuration properties that you can set during installation. You can customize your installation using these properties.

- If you want to use the Webhook of Flink Operator, ensure that you have cert-manager installed on your Kubernetes cluster, which you can install using the following command:

```
kubectl create -f https://github.com/jetstack/cert-manager/releases/download/v1.8.2/cert-manager.yaml
kubectl wait -n cert-manager --for=condition=Available deployment --all
```

- The webhook functionality is enabled by default. You can disable it using the following command, and skip the cert-manager installation:

```
--set flink-kubernetes-operator.webhook.create=false
```

Procedure

1. Create a namespace in your Kubernetes cluster where you will install and use the Cloudera Streaming Analytics - Kubernetes Operator.

```
kubectl create namespace [***NAMESPACE***]
```

This is the namespace where you install Flink and Cloudera SQL Stream Builder. Use this namespace you create in all installation steps that follow.

2. Create a Kubernetes secret to contain your Cloudera credentials.

```
kubectl create secret docker-registry [***SECRET NAME***] \
  --docker-server container.repository.cloudera.com \
  --docker-username [***USERNAME***] \
  --docker-password [***PASSWORD***] \
  --namespace [***NAMESPACE***]
```

Ensure that the placeholders are replaced with your specific information:

- a) Provide a desired name for [***SECRET NAME***].
- b) Replace [***USERNAME***] and [***PASSWORD***] with your Cloudera credentials.
- c) Provide the same name for [***NAMESPACE***] that you created in the previous step.

3. Log in to Cloudera Docker registry with helm.

```
helm registry login container.repository.cloudera.com
```

Enter your Cloudera credentials when prompted.

4. Install Cloudera Streaming Analytics - Kubernetes Operator with helm install.

```
helm install csa-operator --namespace [***NAMESPACE***] \
  --set 'flink-kubernetes-operator.imagePullSecrets[0].name=[***SECRET NAME***]' \
  --set 'ssb.sse.image.imagePullSecrets[0].name=[***SECRET NAME***]' \
  --set 'ssb.sqlRunner.image.imagePullSecrets[0].name=[***SECRET NAME***]' \
  --set-file flink-kubernetes-operator.clouderaLicense.fileContent=[***PATH TO LICENSE FILE***] \
  oci://container.repository.cloudera.com/cloudera-helm/csa-operator/csa-operator --version 1.2.2-b6
```



Important: When you install Cloudera Streaming Analytics - Kubernetes Operator, Cloudera SQL Stream Builder will also be installed by default. In case you want to skip installing it, add `--set ssb.enable=false` to the helm install command.

Ensure that the placeholders are replaced with your specific information:

- Provide the same name for `[***NAMESPACE***]` that you created in Step 1.
 - Provide the same name for `[***SECRET NAME***]` that you created in the previous step. `imagePullSecrets` specifies what secret is used to pull images from the Cloudera registry. Setting this property is mandatory, otherwise Helm will be unable to pull the necessary images from the Cloudera Docker registry.
 - Replace `[***PATH TO LICENSE FILE***]` with the full (absolute) path to your Cloudera license file. `clouderaLicense.fileContent` is used to register your license. When this property is set, a secret is generated that contains the license you specify. Setting this property is mandatory. The Cloudera Streaming Analytics - Kubernetes Operator will not function without a valid license.
 - You can use `--set` to set various other properties of the Helm chart. This enables you to customize your installation. For example, by default the Flink Operator has access to watch all namespaces. However, you can configure a list of specific namespaces to watch using `watchNamespaces`. For example, in case you created multiple namespaces, you can configure the Flink Operator to only watch specific ones with `--set flink-kubernetes-operator.watchNamespaces=[***NAMESPACE1***], [***NAMESPACE2***]`. For more information about deploying Flink and SSB in multiple namespaces, see the [Namespace management](#) documentation.
5. Check that the Flink Operator, and the Cloudera SQL Stream Builder engine with its database are running.

```
kubectl get pods -n [***NAMESPACE***]
```

NAME	READY	STATUS	RESTARTS	AGE
flink-kubernetes-operator	1/2	Running	0	7s
ssb-postgresql	1/1	Running	0	7s
ssb-sse	1/1	Running	0	7s

What to do next

After successfully installing the Cloudera Streaming Analytics - Kubernetes Operator, you can start using Flink and Cloudera SQL Stream Builder on Kubernetes. The [Getting Started with Flink](#) and [Getting Started with Cloudera SQL Stream Builder](#) guides can help you with the basic operations.

Installing Cloudera Streaming Analytics - Kubernetes Operator in an air-gapped environment

Complete these steps to install Cloudera Streaming Analytics - Kubernetes Operator if your Kubernetes cluster does not have internet access, or if you want to install from a self-hosted registry. Installing the Cloudera Streaming Analytics - Kubernetes Operator enables you to deploy and manage Flink and Cloudera SQL Stream Builder in Kubernetes.

About this task

Cloudera Streaming Analytics - Kubernetes Operator is installed in your Kubernetes cluster with the provided Helm chart through the helm install command. When you install the chart, Helm installs the Custom Resource Descriptors (CRD) included in Cloudera Streaming Analytics - Kubernetes Operator, and deploys the Apache Kubernetes Flink Operator (Flink Operator), Cloudera SQL Stream Builder engine, and a PostgreSQL database for Cloudera SQL Stream Builder.

Installing Cloudera Streaming Analytics - Kubernetes Operator does not create or deploy a Flink cluster. The Flink cluster is created after the installation by deploying the Flink Deployment resource in the Kubernetes cluster with kubectl or oc, or when you execute a SQL job in Streaming SQL Console.



Note: Cloudera recommends that you install Cloudera Streaming Analytics - Kubernetes Operator once per Kubernetes cluster.

By default the Flink Operator (deployed with installation) watches and manages all the Flink clusters that are deployed in the single same namespace as the Flink Operator. However, you can also configure it to watch and manage multiple namespaces. This allows you to manage multiple Flink clusters deployed in different namespaces, using a single Cloudera Streaming Analytics - Kubernetes Operator installation.



Important: Cloudera SQL Stream Builder can only be deployed in one namespace. In case you require multiple instances, you must install Cloudera SQL Stream Builder in every namespace. For more information about deploying Flink and Cloudera SQL Stream Builder in multiple namespaces, see the [Namespace management](#) documentation.

Before you begin

- Ensure that your Kubernetes environment meets requirements listed in [System requirements](#).



Note:

If you're installing Cloudera Streaming Analytics - Kubernetes Operator on a Longhorn storage provider (Rancher/RKE2) the following line has to be added to the Helm values.yaml file:

```
ssb:
  database:
    pod:
      securityContext:
        fsGroup: 999
```

- A self-hosted Docker registry is required. Your registry must be accessible by your Kubernetes cluster.
- While the Kubernetes cluster does not need internet access in an air-gapped environment, the preparation steps to create the local (offline) repository, from which you can install Cloudera Streaming Analytics - Kubernetes Operator, require that you can download and move the artifacts hosted on the Cloudera Docker registry and Cloudera Archive.
- Access to docker or equivalent utility that you can use to pull and push images is required. The Cloudera-recommended way is using docker. Replace commands where necessary, if you use a different utility.
- Ensure that you have access to your Cloudera credentials (username and password). Credentials are required to access the Cloudera Docker registry (and, if needed, the Cloudera Archive) where installation artifacts are hosted.

- Ensure that you have access to a valid Cloudera license.
- Review the [Helm chart reference](#) before installation.

The Helm chart accepts various configuration properties that you can set during installation. Using these properties you can customize your installation.

- If you want to use the Webhook of Flink Operator, ensure that you have cert-manager installed on your Kubernetes cluster, which you can install using the following command:

```
kubect1 create -f https://github.com/jetstack/cert-manager/releases/download/v1.8.2/cert-manager.yaml
kubect1 wait -n cert-manager --for=condition=Available deployment --all
```

- The webhook functionality is enabled by default. You can disable it using the following command, and skip the cert-manager installation:

```
--set flink-kubernetes-operator.webhook.create=false
```

Procedure

1. Copy the following installation artifacts to your self-hosted registry.

Table 3: Cloudera Streaming Analytics - Kubernetes Operator artifacts on the Cloudera Docker registry

Artifact	Location	Description
Flink Kubernetes Operator Docker image	container.repository.cloudera.com/cloudera/flink-kubernetes-operator:1.9-csaop1.2.2-b6	Docker image used for deploying the various operator components shipped with the Cloudera Streaming Analytics - Kubernetes Operator.
Flink Docker image	container.repository.cloudera.com/cloudera/flink:1.19.2-csaop1.2.2-b6	Docker image used for deploying Apache Flink and its related components.
SQL Runner Docker image	container.repository.cloudera.com/cloudera/sb-sql-runner:1.19.2-csaop1.2.2-b6	Docker image used for deploying Flink application when SQL query is executed in SSB.
SQL Stream Engine Docker image	container.repository.cloudera.com/cloudera/sb-sse:1.19.2-csaop1.2.2-b6	Docker image used for deploying Cloudera SQL Stream Builder and its UI.



Note: The images are built for linux/arm64 and linux/amd64 architectures.

This step involves pulling the artifacts from the Cloudera Docker registry, retagging them, and then pushing them to your self-hosted registry. The exact steps you need to carry it out depend on your environment and how you set up your registry. The following substeps demonstrate a basic workflow using docker and helm.

- a) Log in to the Cloudera Docker registry with both docker and helm.

Provide your Cloudera credentials when prompted.

```
docker login container.repository.cloudera.com
```

```
helm registry login container.repository.cloudera.com
```

- b) Pull the Docker images from the Cloudera Docker registry.

```
docker pull \
  --platform [***PLATFORM/ARCHITECTURE***] \
```

```
container.repository.cloudera.com/cloudera/[***IMAGE
NAME***]:[***VERSION***]
```

- c) Pull the Cloudera Streaming Analytics - Kubernetes Operator Helm chart.

```
helm pull \
oci://container.repository.cloudera.com/cloudera-helm/csa-operator/csa
-operator \
--version 1.2.2-b6
```

- d) Retag the Docker images you pulled so that they contain the address of your registry.

```
docker tag \
[***ORIGINAL IMAGE TAG***] \
[***REGISTRY HOSTNAME***]:[***PORT***]/cloudera/[***IMAGE NAME***]:
[***VERSION***]
```

- e) Push the images and chart to your self-hosted registry.

```
docker push \
[***REGISTRY HOSTNAME***]:[***PORT***]/cloudera/[***IMAGE
NAME***]:[***VERSION***]
```

```
helm push \
csa-operator-1.2.2-b6.tgz \
oci://[***REGISTRY HOSTNAME***]:[***PORT***]/cloudera-helm/csa-operator/
```

2. Create a namespace in your Kubernetes cluster where you will install and use the Cloudera Streaming Analytics - Kubernetes Operator.

```
kubectl create namespace [***NAMESPACE***]
```

This is the namespace where you install Flink and Cloudera SQL Stream Builder. Use this namespace you create in all installation steps that follow.

3. Create a Kubernetes secret to credentials for your self-hosted registry.

```
kubectl create secret docker-registry [***SECRET NAME***] \
--docker-server [***REGISTRY HOSTNAME***]:[***PORT***] \
--docker-username [***USERNAME***] \
--docker-password [***PASSWORD***] \
--namespace [***NAMESPACE***]
```

Ensure that the placeholders are replaced with your specific information:

- Provide a desired name for [***SECRET NAME***].
 - Replace [***REGISTRY HOSTNAME***]:[***PORT***] with your self-hosted registry hostname and port.
 - Replace [***USERNAME***] and [***PASSWORD***] with your Cloudera credentials.
 - Provide the same name for [***NAMESPACE***] that you created in the previous step.
4. Log in to your self-hosted registry with helm.

```
helm registry login [***REGISTRY HOSTNAME***]:[***PORT***]
```

Enter your credentials when prompted.

5. Install Cloudera Streaming Analytics - Kubernetes Operator with helm install.

```
helm install csa-operator \
--namespace [***NAMESPACE***] \
--set 'flink-kubernetes-operator.image.repository=[***REGISTRY
HOSTNAME***]:[***PORT***]/cloudera/[***IMAGE NAME***]'\
```

```
--set 'ssb.sqlRunner.image.repository=[***REGISTRY
HOSTNAME***]:[***PORT***]/cloudera/[***IMAGE NAME***]'\
--set 'ssb.sse.image.repository=[***REGISTRY
HOSTNAME***]:[***PORT***]/cloudera/[***IMAGE NAME***]'\
--set 'flink-kubernetes-operator.imagePullSecrets[0].name=[***SECRET
NAME***]'\
--set 'ssb.sse.image.imagePullSecrets[0].name=[***SECRET NAME***]'\
--set 'ssb.sqlRunner.image.imagePullSecrets[0].name=[***SECRET
NAME***]'\
--set-file flink-kubernetes-operator.clouderaLicense.fileCon
tent=[***PATH TO LICENSE FILE***] \
oci://[***REGISTRY HOSTNAME***]:[***PORT***]/cloudera-helm/csa-operator/
csa-operator --version 1.2.2-b6
```



Important: When you install Cloudera Streaming Analytics - Kubernetes Operator, Cloudera SQL Stream Builder will also be installed by default. In case you want to skip installing it, add `--set ssb.enable=false` to the helm install command.

Ensure that the placeholders are replaced with your specific information:

- Provide the same name for `[***NAMESPACE***]` that you created in Step 1.
- Replace `[***REGISTRY HOSTNAME***]:[***PORT***]` with your self-hosted registry hostname and port.
- Provide the same name for `[***SECRET NAME***]` that you created in the previous step. `imagePullSecrets` specifies what secret is used to pull images from the Cloudera registry. Setting this property is mandatory, otherwise, Helm cannot pull the necessary images from the Cloudera Docker registry.
- Replace `[***PATH TO LICENSE FILE***]` with the full (absolute) path to your Cloudera license file. `clouderaLicense.fileContent` is used to register your license. When this property is set, a secret is generated that contains the license you specify. Setting this property is mandatory. The Cloudera Streaming Analytics - Kubernetes Operator will not function without a valid license.
- You can use `--set` to set various other properties of the Helm chart. This enables you to customize your installation. (For more information on the available properties, see [Helm chart reference](#).) For example, by default the Flink Operator has access to watch all namespaces. However, you can configure a list of specific namespaces to watch using `watchNamespaces`. For example, in case you created multiple namespaces, you can configure the Flink Operator to only watch specific ones with `--set flink-kubernetes-operator.watchNamespaces=[***NAMESPACE1***], [***NAMESPACE2***]`. For more information about deploying Flink and Cloudera SQL Stream Builder in multiple namespaces, see the [Namespace management](#) documentation.

- Check that the Flink Operator, and the Cloudera SQL Stream Builder engine with its database are running.

```
kubectl get pods -n [***NAMESPACE***]
```

NAME	READY	STATUS	RESTARTS	AGE
flink-kubernetes-operator	1/2	Running	0	7s
ssb-postgresql	1/1	Running	0	7s
ssb-sse	1/1	Running	0	7s

What to do next

After successfully installing the Cloudera Streaming Analytics - Kubernetes Operator, you can start using Flink and Cloudera SQL Stream Builder on Kubernetes. The [Getting Started with Flink](#) and [Getting Started with Cloudera SQL Stream Builder](#) guides can help you with the basic operations.