

CSA Operator Reference

Date published:

Date modified:

CLOUDERA

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

Helm chart configuration reference..... 4

Helm chart configuration reference

You can review the configurable properties of the Cloudera Streaming Analytics Operator for Kubernetes Helm chart and their default values.

Table 1: Cloudera Streaming Analytics Operator for Kubernetes Helm chart configuration properties reference

Parameters	Description	Default Value
flink-kubernetes-operator.enabled	Specifies to install the Flink Operator or not.	<i>TRUE</i>
flink-kubernetes-operator.clouderaLicense	Name of the secret if it was created previously.	<i>NULL</i>
flink-kubernetes-operator.clouderaLicense.fileContent	Cloudera license content.	<i>NULL</i>
flink-kubernetes-operator.clouderaLicense.secretRef	The name of the Kubernetes secret that contains the Cloudera license.	<i>NULL</i>
flink-kubernetes-operator.watchNamespaces	List of kubernetes namespaces to watch for FlinkDeployment changes, empty means all namespaces.	
flink-kubernetes-operator.image.repository	The image repository of flink-kubernetes-operator.	<i>GHCR.IO/APACHE/FLINK-KUBERNETES-OPERATOR</i>
flink-kubernetes-operator.image.pullPolicy	The image pull policy of flink-kubernetes-operator.	<i>IFNOTPRESENT</i>
flink-kubernetes-operator.image.tag	The image tag of flink-kubernetes-operator.	<i>LATEST</i>
flink-kubernetes-operator.image.digest	The image tag of flink-kubernetes-operator. If set then it takes precedence and the image tag will be ignored.	
flink-kubernetes-operator.replicas	Operator replica count. Must be 1 unless leader election is configured.	<i>1</i>
flink-kubernetes-operator.strategy.type	Operator pod upgrade strategy. Must be Recreate unless leader election is configured.	<i>RECREATE</i>
flink-kubernetes-operator.rbac.create	Whether to enable RBAC to create for said namespaces.	<i>TRUE</i>
flink-kubernetes-operator.rbac.nodesRule.create	Whether to add RBAC rule to list nodes which is needed for rest-service exposed as NodePort type.	<i>FALSE</i>
flink-kubernetes-operator.operatorPod.annotations	Custom annotations to be added to the operator pod (but not the deployment).	
flink-kubernetes-operator.operatorPod.labels	Custom labels to be added to the operator pod (but not the deployment).	
flink-kubernetes-operator.operatorPod.env	Custom env to be added to the operator pod.	
flink-kubernetes-operator.operatorPod.envFrom	Custom envFrom settings to be added to the operator pod.	
flink-kubernetes-operator.operatorPod.dnsPolicy	DNS policy to be used by the operator pod.	
flink-kubernetes-operator.operatorPod.dnsConfig	DNS configuration to be used by the operator pod.	
flink-kubernetes-operator.operatorPod.nodeSelector	Custom nodeSelector to be added to the operator pod.	
flink-kubernetes-operator.operatorPod.topologySpreadConstraints	Custom topologySpreadConstraints to be added to the operator pod.	

Parameters	Description	Default Value
flink-kubernetes-operator.operatorPod.resources	Custom resources block to be added to the operator pod on main container.	
flink-kubernetes-operator.operatorPod.webhook.resources	Custom resources block to be added to the operator pod on flink-webhook container.	
flink-kubernetes-operator.operatorPod.tolerations	Custom tolerations to be added to the operator pod.	
flink-kubernetes-operator.operatorServiceAccount.create	Whether to enable operator service account to create for flink-kubernetes-operator.	<i>TRUE</i>
flink-kubernetes-operator.operatorServiceAccount.annotations	The annotations of operator service account.	
flink-kubernetes-operator.operatorServiceAccount.name	The name of operator service account.	<i>FLINK-OPERATOR</i>
flink-kubernetes-operator.jobServiceAccount.create	Whether to enable job service account to create for flink jobmanager/taskmanager pods.	<i>TRUE</i>
flink-kubernetes-operator.jobServiceAccount.annotations	The annotations of job service account.	<i>“HELM.SH/RESOURCE-POLICY”: KEEP</i>
flink-kubernetes-operator.jobServiceAccount.name	The name of job service account.	<i>FLINK</i>
flink-kubernetes-operator.operatorVolumeMounts.create	Whether to enable operator volume mounts to create for flink-kubernetes-operator.	<i>FALSE</i>
flink-kubernetes-operator.operatorVolumeMounts.data	List of mount paths of operator volume mounts.	<i>- NAME: FLINK-ARTIFACTS MOUNTPATH: /OPT/FLINK/ARTIFACTS</i>
flink-kubernetes-operator.operatorVolumes.create	Whether to enable operator volumes to create for flink-kubernetes-operator.	<i>FALSE</i>
flink-kubernetes-operator.operatorVolumes.data	The ConfigMap of operator volumes.	<i>- NAME: FLINK-ARTIFACTS HOSTPATH: PATH: /TMP/FLINK/ARTIFACTS TYPE: DIRECTORYORCREATE</i>
flink-kubernetes-operator.podSecurityContext	Defines privilege and access control settings for a pod or container for pod security context.	<i>RUNASUSER: 9999 RUNASGROUP: 9999</i>
flink-kubernetes-operator.operatorSecurityContext	Defines privilege and access control settings for a pod or container for operator security context.	
flink-kubernetes-operator.webhookSecurityContext	Defines privilege and access control settings for a pod or container for webhook security context.	
flink-kubernetes-operator.webhook.create	Whether to enable validating and mutating webhooks for flink-kubernetes-operator.	<i>TRUE</i>
flink-kubernetes-operator.webhook.mutator.create	Enable or disable mutating webhook, overrides webhook.create	
flink-kubernetes-operator.webhook.validator.create	Enable or disable validating webhook, overrides webhook.create	
flink-kubernetes-operator.webhook.keystore	The ConfigMap of webhook key store.	<i>USEDEFAULTPASSWORD: TRUE</i>
flink-kubernetes-operator.defaultConfiguration.create	Whether to enable default configuration to create for flink-kubernetes-operator.	<i>TRUE</i>
flink-kubernetes-operator.defaultConfiguration.append	Whether to append configuration files with configs.	<i>TRUE</i>

Parameters	Description	Default Value
flink-kubernetes-operator.defaultConfiguration.flink-conf.yaml	The default configuration of flink-conf.yaml.	<code>KUBERNETES.OPERATOR.METRICS.REPORTER.SLF4J.FACTORY.ORG.APACHE.FLINK.METRICS.SLF4J.SLF4JREPORTERFACTORY</code> <code>KUBERNETES.OPERATOR.METRICS.REPORTER.SLF4J.INTERVAL: 5 MINUTE</code> <code>KUBERNETES.OPERATOR.RECONCILE.INTERVAL: 15S</code> <code>KUBERNETES.OPERATOR.OBSERVER.PROGRESS-CHECK.INTERVAL: 5 S</code>
flink-kubernetes-operator.defaultConfiguration.log4j-operator.properties	The default configuration of log4j-operator.properties.	
flink-kubernetes-operator.defaultConfiguration.log4j-console.properties	The default configuration of log4j-console.properties.	
flink-kubernetes-operator.metrics.port	The metrics port on the container for default configuration.	
flink-kubernetes-operator.imagePullSecrets	The image pull secrets of flink-kubernetes-operator.	
flink-kubernetes-operator.nameOverride	Overrides the name with the specified name.	
flink-kubernetes-operator.fullnameOverride	Overrides the fullname with the specified full name.	
flink-kubernetes-operator.jvmArgs.webhook	The JVM start up options for webhook.	
flink-kubernetes-operator.jvmArgs.operator	The JVM start up options for operator.	
flink-kubernetes-operator.operatorHealth.port	Operator health endpoint port to be used by the probes.	8085
flink-kubernetes-operator.operatorHealth.livenessProbe	Liveness probe configuration for the operator using the health endpoint. Only time settings should be configured, endpoint is set automatically based on port.	<code>PERIODSECONDS: 10</code> <code>INITIALDELAYSECONDS: 30</code>
flink-kubernetes-operator.operatorHealth.startupProbe	Startup probe configuration for the operator using the health endpoint. Only time settings should be configured, endpoint is set automatically based on port.	<code>FAILURETHRESHOLD: 30</code> <code>PERIODSECONDS: 10</code>
flink-kubernetes-operator.postStart	The postStart hook configuration for the main container.	
flink-kubernetes-operator.tls.create	Whether to mount an optional secret containing a tls truststore for the flink-kubernetes-operator.	<code>FALSE</code>
flink-kubernetes-operator.tls.secretName	The name of the tls secret	<code>FLINK-OPERATOR-CERT</code>
flink-kubernetes-operator.tls.secretKeyRef.name	The name of the secret containing the password for the java keystore/truststore	<code>OPERATOR-CERTIFICATE-PASSWORD</code>
flink-kubernetes-operator.tls.secretKeyRef.key	The key that holds this password	<code>PASSWORD</code>
ssb.enabled	Specifying if Cloudera SQL Stream Builder should be installed or not.	<code>TRUE</code>
ssb.ssbServiceAccount.create	Creating Cloudera SQL Stream Builder service account.	<code>TRUE</code>
ssb.ssbServiceAccount.annotations		
ssb.ssbServiceAccount.name		<code>"SSB"</code>
ssb.rbac.create	Generating RBAC resources.	<code>TRUE</code>
ssb.rbac.ssbRole.create		<code>TRUE</code>
ssb.rbac.ssbRole.name		<code>"SSB"</code>
ssb.rbac.ssbRoleBinding.create		<code>TRUE</code>

Parameters	Description	Default Value
ssb.rbac.ssbRoleBinding.name		'SSB-ROLE'
ssb.fernetKey	This should be exactly 32-byte long randomly generated secret that will be used to encode Flink payloads with.	
ssb.userManagement.registrationEnabled	Predefined list of users who will be automatically registered when SSB-SSE starts.	TRUE
ssb.userManagement.users.name		ADMIN
ssb.userManagement.users.password		ADMIN
ssb.userManagement.users.isAdmin		TRUE
ssb.sqlRunner.image.repository	SQL Runner image information.	DOCKER-PRIVATE.INFRA.CLOUDERA.COM/CLOUDERA/SSB-SQL-RUNNER
ssb.sqlRunner.image.pullPolicy		IFNOTPRESENT
ssb.sqlRunner.image.tag		LATEST
ssb.sqlRunner.image.imagePullSecrets		
ssb.sse.image.repository	Cloudera SQL Stream Builder engine image information.	DOCKER-PRIVATE.INFRA.CLOUDERA.COM/CLOUDERA/SSB-SSE
ssb.sse.image.pullPolicy		IFNOTPRESENT
ssb.sse.image.tag		LATEST
ssb.sse.image.imagePullSecrets		
ssb.pod.annotations	Configurations to customize the SSE pod to your clusters needs.	
ssb.pod.labels		
ssb.pod.resources.requests.cpu		"1"
ssb.pod.resources.requests.memory		"2048MI"
ssb.pod.priorityClassName		NULL
ssb.pod.nodeSelector		
ssb.pod.tolerations		
ssb.pod.topologySpreadConstraints		
ssb.service.labels	Specification of the Service to be created. No Service will be created when left empty.	
ssb.service.annotations		
ssb.service.spec.type		CLUSTERIP
ssb.service.spec.ports.name		"SSE"
ssb.service.spec.ports.port		18121
ssb.service.spec.ports.targetPort		18121
ssb.ingress	Specification of the Ingress to be created. No Ingress will be created when left empty.	
ssb.httpRoute	Specification of a Gateway API HTTPRoute for Cloudera SQL Stream Builder. No HTTPRoute is created when left empty.	
ssb.sampling.secretRef	Name of a Kubernetes Secret that contains sampling Kafka connection parameters. When set, sampling Kafka settings cannot be edited in the Cloudera SQL Stream Builder UI.	

Parameters	Description	Default Value
ssb.flinkRbac.create	Create Flink ServiceAccount and RoleBinding in the Cloudera SQL Stream Builder namespace when the Flink Operator subchart is disabled.	<i>FALSE</i>
ssb.storageConfiguration	This is the configuration we use to store user artifacts, such as Java UDFs and can be used as durable storage for the Flink jobs as well, for example S3 sink, checkpoint/savepoint storage.	
ssb.podVolumes.create	You can use podVolumes and podVolumeMounts to customize SSB and Flink pods. You can specify custom ConfigMaps, Secrets and any other volume types that will be mounted to both Cloudera SQL Stream Builder and all created Flink pods.	<i>FALSE</i>
ssb.podVolumes.data		
ssb.podVolumeMounts.create		<i>FALSE</i>
ssb.podVolumeMounts.data		
ssb.ssbConfiguration.application.properties	Cloudera SQL Stream Builder configuration overrides	
ssb.ssbConfiguration.log4j-operator.properties	Cloudera SQL Stream Builder Logging Overrides	
ssb.flinkConfiguration.flink-conf.yaml	Default Flink Configuration, this will affect all Flink jobs that are started by Cloudera SQL Stream Builder.	
ssb.database.create	Installs postgres deployment, service and PersistentVolumeClaim for Cloudera SQL Stream Builder.	<i>TRUE</i>
ssb.database.image.repository		<i>POSTGRES</i>
ssb.database.image.tag		<i>18.4</i>
ssb.database.image.digest		
ssb.database.pod.securityContext		
ssb.database.pvcSpec.accessModes	Specification of the created PersistentVolumeClaim for Postgres data.	<i>READWRITEONCE</i>
ssb.database.pvcSpec.resources.requests.storage		<i>100MI</i>
ssb.database.auth.type	Configurations to use if Postgres deployment is not created.	<i>POSTGRESQL</i>
ssb.database.auth.jdbcUrl		<i>"JDBC:POSTGRESQL://SSB-POSTGRESQL:5432/SSB_ADMIN"</i>
ssb.database.auth.username		
ssb.database.auth.password		
ssb.database.auth.secretRef.name	Read database credentials from an existing Secret. When namespace is set, credentials are copied from a cross-namespace Secret.	
ssb.database.auth.secretRef.namespace		
ssb.database.auth.secretRef.usernameKey		<i>USERNAME</i>
ssb.database.auth.secretRef.passwordKey		<i>PASSWORD</i>
ssb.database.service.labels	Labels added to the PostgreSQL Service when database.create is true.	

Cloudera SQL Stream Builder connectors installed by default

The following connectors are available in the default Cloudera SQL Stream Builder image. Hadoop, Hive, Iceberg, and Kudu require the alternative Flink image; see *Installation overview*.

Table 2: Default connectors

Connector	Notes
Kafka	Source and sink; Schema Registry catalog support.
JDBC	Relational databases.
CDC	MySQL, Oracle, PostgreSQL, Db2, SQL Server.
Amazon S3 / Azure Blob / Google Cloud Storage	Object storage sinks and filesystem access.
Iceberg	Alternative image; first-class catalog in Cloudera SQL Stream Builder. See <i>Iceberg tables and connectors</i> .
Hive / HDFS / Kudu	Alternative image only.