Cloudera Streams Messaging Operator for Kubernetes 1.7.0

Reference

Date published: 2024-06-11
Date modified: 2026-07-31

https://docs.cloudera.com/

Legal Notice

Contents

Strimzi Cluster Operator Helm chart configuration reference

A reference listing all configurable properties of the Strimzi Cluster Operator Helm chart.

Configure these properties during installation in your helm install command, or update them after installation using helm upgrade. Use --set options to specify properties. For example:

```
helm install strimzi-cluster-operator \
  --namespace [***NAMESPACE***] \
  --set 'image.imagePullSecrets[0].name=[***SECRET NAME***]' \
  --set-file clouderaLicense.fileContent=[***PATH TO LICENSE FILE***] \
  --set watchAnyNamespace=true \
  oci://container.repository.cloudera.com/cloudera-helm/csm-operator/strimzi
-kafka-operator \
  --version 1.7.0-b248
```

Table 1: Strimzi Cluster Operator Helm chart configuration properties reference

Parameter	Description	Default
replicas	Number of replicas of the cluster operator	1
revisionHistoryLimit	Number of replicaSet to keep of the operator deployment	10
watchNamespaces	Comma separated list of additional namespaces for the strimzi-operator to watch	[]
watchAnyNamespace	Watch the whole Kubernetes cluster (all namespaces)	false
defaultImageRegistry	Default image registry for all the images	container.repository.cloudera.com
defaultImageRepository	Default image registry for all the images	cloudera
defaultImageTag	Default image tag for all the images except Kafka Bridge	1.0.1-csmop-1.7.0-b248
deploymentAnnotations	Annotations for the operator deployment	{}
deploymentLabels	Labels for the operator deployment	{}
image.registry	Override default Cluster Operator image registry	nil
image.repository	Override default Cluster Operator image repository	nil
image.name	Cluster Operator image name	kafka-operator
image.tag	Override default Cluster Operator image tag	nil
image.digest	Override Cluster Operator image tag with digest	nil
image.imagePullPolicy	Image pull policy for all pods deployed by Cluster Operator	IfNotPresent
image.imagePullSecrets	List of Docker registry pull secrets	[]
fullReconciliationIntervalMs	Full reconciliation interval in milliseconds	120000
leaderElection.enable	Whether to enable leader election	true
operationTimeoutMs	Operation timeout in milliseconds	300000

Parameter	Description	Default
operatorNamespaceLabels	Labels of the namespace where the operator runs	nil
podSecurityContext	Cluster Operator pod's security context	nil
deploymentStrategy	Adjust the Kubernetes rollout strategy of the Cluster Operator Deployment	nil
priorityClassName	Cluster Operator pod's priority class name	nil
hostUsers	Set hostUsers value on the Cluster Operator container	nil
securityContext	Cluster Operator container's security context	nil
rbac.create	Whether to create RBAC related resources	yes
serviceAccountCreate	Whether to create a service account	yes
serviceAccount	Cluster Operator's service account	strimzi-cluster-operator
podDisruptionBudget.enabled	Whether to enable the podDisruptionBudget feature	false
podDisruptionBudget.minAvailable	Default value for how many pods must be running in a cluster	1
podDisruptionBudget.maxUnavailable	Default value for how many pods can be down	nil
podDisruptionBudget.unhealthyPodEviction Policy	Default value for how to respond to unheathly pods	IfHealthyBudget
extraEnvs	Extra environment variables for the Cluster operator container	[]
kafka.image.registry	Override default Kafka image registry	nil
kafka.image.repository	Override default Kafka image repository	nil
kafka.image.name	Kafka image name	kafka
kafka.image.tagPrefix	Override default Kafka image tag prefix	nil
kafka.image.tag	Override default Kafka image tag and ignore suffix	nil
kafka.image.digest	Override Kafka image tag with digest	nil
kafkaConnect.image.registry	Override default Kafka Connect image registry	nil
kafkaConnect.image.repository	Override default Kafka Connect image repository	nil
kafkaConnect.image.name	Kafka Connect image name	kafka
kafkaConnect.image.tagPrefix	Override default Kafka Connect image tag prefix	nil
kafkaConnect.image.tag	Override default Kafka Connect image tag and ignore suffix	nil
kafkaConnect.image.digest	Override Kafka Connect image tag with digest	nil
cruiseControl.image.registry	Override default Cruise Control image registry	nil
cruiseControl.image.repository	Override default Cruise Control image repository	nil
cruiseControl.image.name	Cruise Control image name	kafka
cruiseControl.image.tagPrefix	Override default Cruise Control image tag prefix	nil
cruiseControl.image.tag	Override default Cruise Control image tag and ignore suffix	nil

Parameter	Description	Default
cruiseControl.image.digest	Override Cruise Control image tag with digest	nil
topicOperator.image.registry	Override default Topic Operator image registry	nil
topicOperator.image.repository	Override default Topic Operator image repository	nil
topicOperator.image.name	Topic Operator image name	kafka-operator
topicOperator.image.tag	Override default Topic Operator image tag	nil
topicOperator.image.digest	Override Topic Operator image tag with digest	nil
userOperator.image.registry	Override default User Operator image registry	nil
userOperator.image.repository	Override default User Operator image repository	nil
userOperator.image.name	User Operator image name	kafka-operator
userOperator.image.tag	Override default User Operator image tag	nil
userOperator.image.digest	Override User Operator image tag with digest	nil
kafkaInit.image.registry	Override default Init Kafka image registry	nil
kafkaInit.image.repository	Override default Init Kafka image repository	nil
kafkaInit.image.name	Init Kafka image name	operator
kafkaInit.image.tag	Override default Init Kafka image tag	nil
kafkaInit.image.digest	Override Init Kafka image tag with digest	nil
kafkaBridge.image.registry	Override default Kafka Bridge image registry	quay.io
kafkaBridge.image.repository	Override default Kafka Bridge image repository	strimzi
kafkaBridge.image.name	Kafka Bridge image name	kafka-bridge
kafkaBridge.image.tag	Override default Kafka Bridge image tag	1.0.0
kafkaBridge.image.digest	Override Kafka Bridge image tag with digest	nil
kafkaExporter.image.registry	Override default Kafka Exporter image registry	nil
kafkaExporter.image.repository	Override default Kafka Exporter image repository	nil
kafkaExporter.image.name	Kafka Exporter image name	kafka
kafkaExporter.image.tagPrefix	Override default Kafka Exporter image tag prefix	nil
kafkaExporter.image.tag	Override default Kafka Exporter image tag and ignore suffix	nil
kafkaExporter.image.digest	Override Kafka Exporter image tag with digest	nil
kafkaMirrorMaker2.image.registry	Override default Kafka Mirror Maker 2 image registry	nil
kafkaMirrorMaker2.image.repository	Override default Kafka Mirror Maker 2 image repository	nil
kafkaMirrorMaker2.image.name	Kafka Mirror Maker 2 image name	kafka
kafkaMirrorMaker2.image.tagPrefix	Override default Kafka Mirror Maker 2 image tag prefix	nil
kafkaMirrorMaker2.image.tag	Override default Kafka Mirror Maker 2 image tag and ignore suffix	nil

Parameter	Description	Default
kafkaMirrorMaker2.image.digest	Override Kafka Mirror Maker 2 image tag with digest	nil
kanikoExecutor.image.registry	Override default Kaniko Executor image registry	nil
kanikoExecutor.image.repository	Override default Kaniko Executor image repository	nil
kanikoExecutor.image.name	Kaniko Executor image name	kaniko-executor
kanikoExecutor.image.tag	Override default Kaniko Executor image tag	nil
kanikoExecutor.image.digest	Override Kaniko Executor image tag with digest	nil
buildah.image.registry	Override default Buildah image registry	nil
buildah.image.repository	Override default Buildah image repository	nil
buildah.image.name	Buildah image name	buildah
buildah.image.tag	Override default Buildah image tag	nil
buildah.image.digest	Override Buildah image tag with digest	nil
resources.limits.memory	Memory constraint for limits	384Mi
resources.limits.cpu	CPU constraint for limits	1000m
resources.requests.memory	Memory constraint for requests	384Mi
livenessProbe.initialDelaySeconds	Liveness probe initial delay in seconds	10
livenessProbe.periodSeconds	Liveness probe period in seconds	30
readinessProbe.initialDelaySeconds	Readiness probe initial delay in seconds	10
readinessProbe.periodSeconds	Readiness probe period in seconds	30
imageTagOverride	Override all image tag config	nil
createGlobalResources	Allow creation of cluster-scoped resources	true
createAggregateRoles	Create cluster roles that extend aggregated roles to use Strimzi CRDs	false
tolerations	Add tolerations to Operator Pod	[]
topologySpreadConstraints	Add topology spread constraints to Operator Pod	[]
affinity	Add affinities to Operator Pod	{}
annotations	Add annotations to Operator Pod	{}
labels	Add labels to Operator Pod	{}
nodeSelector	Add a node selector to Operator Pod	{}
featureGates	Feature Gates configuration	nil
tmpDirSizeLimit	Set the sizeLimit for the tmp dir volume used by the operator	1Mi
labelsExclusionPattern	Override the exclude pattern for exclude some labels	""
operatorNetworkPolicy.enabled	Generate a networkPolicy for the strimzi-operator (operator not Kafka)	false
generateNetworkPolicy	Controls whether Strimzi generates network policy resources	true
connectBuildTimeoutMs	Overrides the default timeout value for building new Kafka Connect	300000

Parameter	Description	Default
mavenBuilder.image.registry	Override default Maven Builder image registry	nil
mavenBuilder.image.repository	Maven Builder image repository	nil
mavenBuilder.image.name	Override default Maven Builder image name	maven-builder
mavenBuilder.image.tag	Override default Maven Builder image tag	nil
mavenBuilder.image.digest	Override Maven Builder image tag with digest	nil
logConfiguration	Override default log4j.properties content	nil
logLevel	Override default logging level	INFO
dashboards.enabled	Generate configmaps containing the dashboards	false
dashboards.label	How should the dashboards be labeled for the sidecar	grafana_dashboard
dashboards.labelValue	What should the dashboards label value be for the sidecar	"1"
dashboards.extraLabels	Any additional labels you would like on the dashboards	{}
dashboards.namespace	What namespace should the dashboards be loaded into	Follows toplevel Namespace
dashboards.annotations	Any custom annotations (such as folder for the sidecar)	{}
clouderaLicense.fileContent	Cloudera license content	nil
clouderaLicense.secretRef	The name of the Kubernetes secret that contains the Cloudera license	nil

Related Information

Configuring the Strimzi Cluster Operator

Strimzi Drain Cleaner Helm chart configuration reference

A reference listing configurable properties of the Strimzi Drain Cleaner Helm chart.

Configure these properties during installation in your helm install command, or update them after installation using helm upgrade. Use --set options to specify properties. For example:

```
helm upgrade strimzi-drain-cleaner \
  --namespace strimzi-drain-cleaner \
  --set webhook.failurePolicy=Fail \
  oci://container.repository.cloudera.com/cloudera-helm/csm-operator/strimzi
-drain-cleaner \
  --reuse-values
```

Table 2: Strimzi Drain Cleaner Helm chart configuration properties reference

Parameter	Description	Default
replicaCount	Number of replicas of the Drain Cleaner webhook	2
image.registry	Override default Drain Cleaner image registry	container.repository.cloudera.com

Parameter	Description	Default
image.repository	Override default Drain Cleaner image repository	cloudera
image.name	Drain Cleaner image name	drain-cleaner
image.tag	Override default Drain Cleaner image tag	1.6.0
image.imagePullPolicy	Image pull policy for all pods deployed by Drain Cleaner	nil
image.imagePullSecrets	List of Docker registry pull secrets	[]
resources	Configures resources for the Drain Cleaner Pod	[]
tolerations	Add tolerations to Drain Cleaner Pod	[]
affinity	Add affinities to Drain Cleaner Pod	{}
nodeSelector	Add a node selector to Drain Cleaner Pod	{}
deploymentStrategy	Adjust the Kubernetes rollout strategy of the Drain Cleaner Deployment	{}
webhook.faillurePolicy	Override default validating webhook failurePolicy	Ignore
webhook.labels	Additional labels of the ValidatingWebhookConfiguration	{}
webhook.annotations	Additional annotations of the ValidatingWebhookConfiguration	{}
webhook.timeoutSeconds	Override default validating webhook timeoutSeconds	5
securityContext	Set the security context for the Drain Cleaner container	{}
podSecurityContext	Set the pod security context for the Drain Cleaner pod	{}
certManager.create	Enables or disables the creation of cert-manager resources (Issuer and Certificate) used to secure the Drain Cleaner webhook.	true
certManager.duration	Configures the duration of the cert-manager certificate requested by the Helm Chart. Value must be in units accepted by Go time.ParseDuration, e.g. 2160h.	NULL
certManager.renewBefore	Configures the renewal period of the cert-manager certificate requested by the Helm Chart. Value must be in units accepted by Go time.ParseDuration, e.g. 2160h.	NULL

Related Information

Installing Strimzi Drain Cleaner with Helm

Strimzi Drain Cleaner

Cloudera Surveyor for Apache Kafka Helm chart configuration reference

A reference listing all configurable properties of the Cloudera Surveyor Helm chart.

Configure these properties during installation in your helm install command, or update them after installation using helm upgrade. Use the --values (-f) and --set options to specify properties. For example:

```
helm install cloudera-surveyor \
  --namespace [***NAMESPACE***] \
  --values [***VALUES FILE***] \
  --set-file surveyorConfig.surveyor.authentication.keys.active=[***PATH TO
 AUTHENTICATION KEY FILE***] \
  --set 'image.imagePullSecrets=[***CLOUDERA CREDENTIAL SECRET***]' \
  --set-file clouderaLicense.fileContent=[***PATH TO LICENSE FILE***] \
  oci://container.repository.cloudera.com/cloudera-helm/csm-operator/surv
eyor \
  --version 1.7.0-b248
```

Properties are grouped into two tables. The *General Properties* table lists configuration properties for customizing the overall deployment of Cloudera Surveyor, including settings for global Kafka client configurations, authentication, TLS, `Ingress`, `Secret` management, and others.

The *ClusterConfigs Properties* table lists configuration properties that specify which Kafka clusters Cloudera Surveyor connects to and how connections are established. These properties include settings for cluster identification, per-cluster client configuration, alert thresholds, and more.

Table 3: General Properties

Property	Description	Default Value
replicaCount	The ReplicaSet count.	2
fipsMode	Enable FIPS mode	false
instanceName	Specifies the name of the Cloudera Surveyor instance. When set, the name is displayed on the UI and included in the response of the /api /v1/ui-config endpoint.	none
serviceAccountName	Specifies the name of an existing Kubernetes ServiceAccount for the Cloudera Surveyor pod. When left blank, the pod runs under the default ServiceAccount in its namespace.	""
securityProviders[*]	List of security providers to load	[]
securityProviders[*].providerClassName	Class name of the provider	none
securityProviders[*].argument	Argument to pass to provider	none
securityProviders[*].classpath[*]	List of filesystem paths that should be added to the classpath in order for this provider to function.	[]
image.registry	The container image registry where the image is hosted.	container.repository.cloudera.com
image.repository	The repository name of the container image.	cloudera
image.name	The name of the container image.	surveyor
image.tag	The tag of the container image to use.	latest
image.pullPolicy	The policy to use when pulling the container image.	IfNotPresent
image.uiName	The name of the UI container image.	surveyor-app
image.uiTag	The tag of the UI container image.	latest

Property	Description	Default Value
image.imagePullSecrets	The name of the Secret containing credentials that provide access to the registry where Cloudera Surveyor images are hosted. To specify multiple Secrets, provide an array of objects where each object has a name property: [{name: "secret1"}, {name: "secret2"}].	none
service.type	The type of Kubernetes Service.	ClusterIP
service.port	The port exposed by the Service.	8443
surveyorConfig.*	Cloudera Surveyor configuration object.	{}
surveyorConfig.surveyor.commonClientConfig	Kafka client configuration properties applied to all clients. Must contain upstream Kafka client properties as a map. Use clusterConfigs.clusters[*].commonClientConfig to set client configuration on a per-cluster basis.	{}
surveyorConfig.surveyor.globalHiddenTopicNamePattern	Regular expression to specify which topic should be hidden. Use clusterConfigs.clusters[*].hiddenTopicNamePattern to configure it on a per-cluster basis. The default expression filters topic names that follow common internal topic naming conventions.	^(__.+|ATLAS_.+)
surveyorConfig.surveyor.snapshotClientPool.clientCount	Number of Kafka clients used for snapshotting clusters. Use clusterConfigs.clusters[*].snapshotClientPool.clientCount to set the client count on a per-cluster basis. Cloudera recommends using the default value. Increase the client count if snapshotting data from the cluster is slow.	5
surveyorConfig.surveyor.snapshotClientPool.clientConfig	Kafka client configuration properties applied to all snapshot clients. Must contain upstream Kafka client properties as a map. Use clusterConfigs.clusters[*].snapshotClientPool.clientConfig to set snapshot client configuration on a per-cluster basis.	{}
surveyorConfig.surveyor.adminClientPool.clientCount	Number of Kafka clients used for administrative operations. Administrative operations include actions such as managing topics (create, delete, alter) and managing consumer groups (delete, reset-offsets). Use clusterConfigs.clusters[*].adminClientPool.clientCount to set the client count on a per-cluster basis.	2
surveyorConfig.surveyor.adminClientPool.clientConfig	Kafka client configuration properties applied to all admin clients. Must contain upstream Kafka client properties as a map. Use clusterConfigs.clusters[*].adminClientPool.clientConfig to set admin client configuration on a per-cluster basis.	{}
surveyorConfig.surveyor.maxGlobalSnapshotParallelism	Maximum number of usable threads for processing snapshots across all clusters. Note that this does not control the maximum number of threads created rather the number of active threads. This also does not limit the number of threads in use by the kafka clients. Defaults to the number of available processor cores at runtime.	null

Property	Description	Default Value
surveyorConfig.surveyor.globalSnapshotInterval	Interval specifying how frequently snapshots are made for all clusters in ISO 8601 representation. Use clusterConfigs.clusters[*].snapshotInterval to configure snapshot interval on a per-cluster basis.	PT3M (3 minutes)
surveyorConfig.surveyor.globalAclFetchInterval	Interval specifying how frequently ACLs are fetched and refreshed for all clusters in ISO 8601 representation. Use clusterConfigs.clusters[*].authorization.aclFetchInterval to configure the ACL fetching interval on a per-cluster basis.	PT1M (1 minute)
surveyorConfig.surveyor.globalSnapshotTimeout	Interval specifying the timeout of the snapshot operation for all clusters in ISO 8601 representation. Use clusterConfigs.clusters[*].snapshotTimeout to configure the snapshot operation timeout on a per-cluster basis.	PT1M (1 minute)
surveyorConfig.surveyor.globalSnapshotTtl	Interval specifying how long the last successful snapshot is kept when subsequent snapshots are failing for all clusters in ISO 8601 representation. Use clusterConfigs.clusters[*].snapshotTtl to configure the TTL of the last successful snapshot on a per-cluster basis.	PT8M (8 minutes)
surveyorConfig.surveyor.snapshotMaxJitter	Interval specifying the maximum initial jitter of snapshotting for all clusters in ISO 8601 representation.	PT10S (10 seconds)
surveyorConfig.surveyor.authentication.ldap.enabled	Enables or disables LDAP authentication.	true
surveyorConfig.surveyor.authentication.ldap.keys.active	The secure key used to sign authentication tokens. It must be random bytes of at least 32 length.	crptographically secure random 128 bytes (only in non-FIPS mode)
surveyorConfig.surveyor.authentication.ldap.keys.passive	The secure key used to verify authentication tokens during authentication key rolling. When used, it should be set to the previously active key.	none
surveyorConfig.surveyor.authentication.ldap.principalMappingRule	Central principal mapping rule used to transform user principals into short names before evaluating authorization permissions. You can specify multiple rules, which are evaluated in order. The first rule that matches a principal name is applied for mapping, subsequent rules are ignored. The supported format is: 'RULE:pattern/replacement/, RULE:pattern/replacement/, DEFAULT'. This property works the same way as the ssl.principal.mapping.rules property in Kafka.	DEFAULT
surveyorConfig.surveyor.authentication.ldap.userSessionTimeout	Time interval after which users are required to relogin into Cloudera Surveyor in ISO 8601 representation.	P1D (1 day)
surveyorConfig.surveyor.authentication.ldap.inactivityTimeout	User inactivity timeout in ISO 8601 representation. If users are inactive for the duration of this timeout, they are required to relogin into Cloudera Surveyor.	PT1H (1 hour)
surveyorConfig.surveyor.authentication.ldap.tokenRenewalInterval	Authentication token renewal interval in ISO 8601 representation. This must be smaller than inactivityTimeout, recommended not to be larger than half of the inactivityTimeout.	PT10M (10 minutes)
surveyorConfig.surveyor.authentication.trustedProxy.enabled	Enables or disables TrustedProxy authentication.	false

Property	Description	Default Value
surveyorConfig.surveyor.authentication.trustedProxy.userHeader	The name of the HTTP header that contains the user, provided by the proxy application.	none
surveyorConfig.surveyor.authentication.trustedProxy.rolesHeader	The name of the HTTP header that contains the roles, provided by the proxy application.	none
surveyorConfig.surveyor.authentication.trustedProxy.requestIdHeader	The name of the HTTP header that contains the request id, provided by the proxy application.	none
surveyorConfig.quarkus.*	Quarkus specific configurations.	{}
tlsConfigs.enabled	Enables or disables TLS.	true
tlsConfigs.secretRef	The name of the Secret containing TLS configuration properties.	"tls-config"
clouderaLicense.fileContent	The contents of the Cloudera license. Use with --set-file to generate a Secret automatically that contains the Cloudera license.	none
clouderaLicense.secretRef	The name of the Secret containing the Cloudera license file.	none
ingress.enabled	Enables or disables external access through Ingress.	false
ingress.className	The Ingress controller class.	none
ingress.extraAnnotations.*	Extra annotations to apply to the Ingress.	{nginx.ingress.kubernetes.io/backend-protocol: "HTTPS"}
ingress.rules.host	The host of the Ingress rule.	"none"
ingress.rules.path	The Ingress path.	"/"
ingress.tls.enabled	Enables or disables TLS for Ingress.	true
ingress.tls.secretRef	The name of the Secret containing Ingress TLS certificates.	"ingress-tls-cert"
gateway.enabled	Enables or disables external access using the Gateway API. When enabled, an HTTP Route resource is created that routes traffic to Cloudera Surveyor through an existing Gateway resource. Requires Gateway API v1 CRDs, a Gateway Controller, at least one Gateway resource in the cluster, and at least one Gateway reference in gateway.parentRefs. The structure of gateway.parentRefs[*], gateway.hostnames[*], and gateway.rules[*] follows the Kubernetes Gateway API HTTPRoute specification.	false
gateway.parentRefs[*]	Lists existing Gateway resources that handle traffic routing to Cloudera Surveyor. Each entry specifies a Gateway's name, namespace, and optional group/kind. At least one parent must be provided.	[]
gateway.hostnames[*]	Lists DNS hostnames that the HTTPRoute matches for routing to Cloudera Surveyor (for example, surveyor.example.com, surveyor.internal.com). When a request arrives at the Gateway with one of these hostnames, traffic is routed to the Cloudera Surveyor Service. If omitted, the Gateway's default hostname routing applies.	[]

Property	Description	Default Value
gateway.rules[*]	Lists HTTPRoute routing rules for advanced use cases. Each rule can specify path matching (for example, /api/*), request filters (for example, header modification, request redirection), and target backend services. If you do not specify backendRefs in a rule, a single backendRef pointing to the Cloudera Surveyor Service is automatically configured. If you do not specify a matches block in a rule, a single PathPrefix matches block with value / is automatically configured. If you leave gateway.rules empty, a default rule is automatically created that routes all traffic (/) to the Cloudera Surveyor Service.	[]
globalTruststore.secretRef.name	The name of the Secret containing the global Cloudera Surveyor truststore.	none
globalTruststore.secretRef.key	The key in the Secret containing the global Cloudera Surveyor truststore.	none
globalTruststore.type	The type of the global Cloudera Surveyor truststore. Can be JKS, PKCS12 or BCFKS.	PKCS12
globalTruststore.password.secretRef.name	The name of the Secret containing the global truststore password.	none
globalTruststore.password.secretRef.key	The key of the Secret containing the global truststore password.	none
podSecurityContext	Security context settings at the pod level.	{}
containerSecurityContext	Security context settings at the container level.	{}
resources	Resource requests and limits for the Cloudera Surveyor container.	{}
scheduling.nodeSelector	Node selector constraints for pod scheduling.	{}
scheduling.affinity	Affinity rules for more complex pod scheduling. Includes nodeAffinity, podAffinity, and podAntiAffinity.	{}
scheduling.tolerations	Tolerations to allow pod scheduling on nodes with matching taints.	[]
scheduling.topologySpreadConstraints	Constraints for distributing pods across the cluster topology.	[]
env	List of environment variables to be set in the Cloudera Surveyor container. Can be specified as direct values or references to secrets/configmaps.	See below
secretsToMount[*]	List of secrets to mount	[]
secretsToMount[*].secretRef	The name of the secret to mount	none
secretsToMount[*].create	If set to true a new secret will be created	false
secretsToMount[*].items[*]	List of items in the secret to mount. If empty all items will be mounted	[]
secretsToMount[*].items[*].key	Key of the item in the secret	none
secretsToMount[*].items[*].path	Path of the mounted item relative to /opt/secrets/[***SECRET NAME***]/	none
secretsToMount[*].items[*].content	If create is set to true the content of this item in the newly created secret	none

Property	Description	Default Value
extraVolumes[*]	List of additional volumes to attach to the pod. This allows you to mount various types of volumes like secrets, configMaps, persistentVolumeClaims, etc.	[]
extraVolumes[*].name	The name of the volume. Must match the corresponding name in extraVolumeMounts.	none
extraVolumes[*].<volumeType>	The volume configuration. Can be any valid Kubernetes volume type such as secret, conf igMap, persistentVolumeClaim, etc.	none
extraVolumeMounts[*]	List of volume mounts that specify how volumes are mounted into the container.	[]
extraVolumeMounts[*].name	The name of the volume to mount. Must match a volume defined in extraVolumes.	none
extraVolumeMounts[*].mountPath	The path within the container at which the volume should be mounted.	none
extraVolumeMounts[*].readOnly	Whether the volume should be mounted read-only. Defaults to false.	false
podLabels	Additional labels to add to the Cloudera Surveyor pod.	{}

Table 4: ClusterConfigs Properties

Property	Description	Default Value
clusterConfigs.clusters[*]	Array of connected clusters and their client configuration.	[]
clusterConfigs.clusters[*].clusterName	The name of the cluster. This name is displayed on the UI.	``
clusterConfigs.clusters[*].tags	A list of arbitrary tags associated with the cluster. Use tags to logically group and organize clusters. For example, by department, geographic location, or environment. Properly tagging clusters makes it easier to filter and organize them on the UI.	[]
clusterConfigs.clusters[*].bootstrapServers	A comma-separated list of the bootstrap servers for the Kafka cluster that Cloudera Surveyor connects to. Specify multiple servers for highly available connections.	``
clusterConfigs.clusters[*].snapshotInterval	Interval specifying how frequently snapshots are made for this cluster in ISO 8601 representation. Overrides surveyorConfig.s urveyor.globalSnapshotInterval.	null
clusterConfigs.clusters[*].snapshotTimeout	Interval specifying the timeout of the snapshot operation in ISO 8601 representation. Overrides surveyorConfig.surveyor.globalSn apshotTimeout.	null
clusterConfigs.clusters[*].snapshotTtl	Interval specifying how long the last successful snapshot is kept when subsequent snapshots are failing in ISO 8601 representation. Overrides surveyorConfig.s urveyor.globalSnapshotTtl.	null
clusterConfigs.clusters[*].commonClientC onfig	Kafka client configuration properties applied to all clients for this cluster. Must contain upstream Kafka client properties as a map. Properties specified here are merged with, and take precedence over, the client configuration specified in surveyorConfig.surveyor.*.	{}

Property	Description	Default Value
clusterConfigs.clusters[*].hiddenTopicNamePattern	Regular expression to specify which topic should be hidden. Overrides 'surveyorConfig.surveyor.globalHiddenTopicNamePattern'.	null
clusterConfigs.clusters[*].brokerAddressListener	Specifies the name of the Kafka listener to use for resolving broker addresses. When set, Cloudera Surveyor looks up each broker's address from advertised.listeners (falling back to listeners) for the specified listener name, instead of using the addresses from the bootstrap listener. Use this when the listener that bootstrap.servers points to differs from the listener whose addresses you want displayed in the UI.	null
clusterConfigs.clusters[*].snapshotClientPool.clientCount	Number of Kafka clients used for snapshotting this cluster. Overrides surveyorConfig.surveyor.snapshotClientPool.clientCount. Cloudera recommends using the default value. Increase the client count if snapshotting data from the cluster is slow.	null
clusterConfigs.clusters[*].snapshotClientPool.clientConfig	Kafka client configuration properties applied to all snapshot clients for this cluster. Must contain upstream Kafka client properties as a map. Properties specified here are merged with, and take precedence over, the client configuration specified in surveyorConfig.surveyor.* and clusterConfigs.clusters[*].commonClientConfig.	{}
clusterConfigs.clusters[*].adminClientPool.clientCount	Number of Kafka clients to use for administrative operations in this cluster. Administrative operations include actions such as managing topics (create, delete, alter) and managing consumer groups (delete, reset-offsets). Overrides surveyorConfig.surveyor.adminClientPool.clientCount.	null
clusterConfigs.clusters[*].adminClientPool.clientConfig	Kafka client configuration properties applied to all admin clients for this cluster. Must contain upstream Kafka client properties as a map. Properties specified here are merged with, and take precedence over, the client configuration specified in surveyorConfig.surveyor.* and clusterConfigs.clusters[*].commonClientConfig.	{}
clusterConfigs.clusters[*].adminOperationTimeout	Timeout used for administrative operations in ISO 8601 representation.	null
clusterConfigs.clusters[*].readOnlyBrokerConfigs	Enables or disables broker configuration editing in Cloudera Surveyor. When set to true, editing is disabled. Configure this property if the lifecycle of your Kafka clusters is managed by tooling or systems such as the Strimzi Cluster Operator, where changes made through Cloudera Surveyor might be automatically overwritten during reconciliation or deployment.	false
clusterConfigs.clusters[*].allTopicConfigsFetchRetries	Number of retries when fetching all the available topic configurations in the cluster.	null
clusterConfigs.clusters[*].alertConfigs.logDirUsageConcerningThresholdPercent	The percentage of log directory usage that triggers a concerning alert.	null
clusterConfigs.clusters[*].alertConfigs.logDirUsageCriticalThresholdPercent	The percentage of log directory usage that triggers a critical alert.	null

Property	Description	Default Value
clusterConfigs.clusters[*].alertConfigs.diskVolumeImbalanceThresholdPercent	The percentage that a log directory's size can differ from the average size of all log directories before an alert is triggered.	null
clusterConfigs.clusters[*].alertConfigs.brokerVolumeImbalancePercent	The percentage that a broker's total log directory size can differ from the average across all brokers before an alert is triggered.	null
clusterConfigs.clusters[*].alertConfigs.brokerReplicaImbalancePercent	The percentage that the number of replicas hosted on a broker can differ from the average across all brokers before an alert is triggered.	null
clusterConfigs.clusters[*].alertConfigs.brokerLeaderImbalancePercent	The percentage that the number of leader replicas hosted on a broker can differ from the average across all brokers before an alert is triggered.	null
clusterConfigs.clusters[*].alertConfigs.totalLagConcerning	The amount of lag for a consumer group that triggers a concerning alert.	null
clusterConfigs.clusters[*].alertConfigs.totalLagCritical	The amount of lag for a consumer group that triggers a critical alert.	null
clusterConfigs.clusters[*].alertConfigs.consumerGroupOverrides	Specifies per-consumer-group alert threshold overrides. Each override contains a pattern that is evaluated against the full consumer group ID in the order listed. The first matching pattern is applied. Each override can specify optional totalLagConcerning and totalLagCritical thresholds. Omitted thresholds fall back to cluster-level defaults. Set a threshold to -1 to disable that alert for matching groups.	[]
clusterConfigs.clusters[*].alertConfigs.consumerGroupOverrides[*].pattern	Specifies a Java regular expression matched against the full consumer group ID using Pattern.matches(). Invalid patterns are logged as warnings and skipped.	(required)
clusterConfigs.clusters[*].alertConfigs.consumerGroupOverrides[*].totalLagConcerning	Specifies an override for the concerning lag threshold for consumer groups matching this pattern. Set to -1 to disable concerning alerts for matching groups, or omit to use the cluster-level default set in clusterConfigs.clusters[*].alertConfigs.totalLagConcerning.	null
clusterConfigs.clusters[*].alertConfigs.consumerGroupOverrides[*].totalLagCritical	Specifies an override for the critical lag threshold for consumer groups matching this pattern. Set to -1 to disable critical alerts for matching groups, or omit to use the cluster-level default set in clusterConfigs.clusters[*].alertConfigs.totalLagCritical.	null
clusterConfigs.clusters[*].authorization.enabled	Enables authorization.	true
clusterConfigs.clusters[*].authorization.aclFetchInterval	Interval specifying how frequently ACLs are fetched and refreshed for this cluster in ISO 8601 representation. Overrides surveyorConfig.surveyor.globalAclFetchInterval.	null
clusterConfigs.clusters[*].authorization.principalMappingRule	Principal mapping rule. Follows the syntax of Kafka's ssl.principal.mapping.rules property.	DEFAULT
clusterConfigs.clusters[*].authorization.superUsers	List of superusers.	[]
clusterConfigs.clusters[*].authorization.defaultResult	Default result of the authorization. Accepted values are: DENIED or ALLOWED.	DENIED

Property	Description	Default Value
clusterConfigs.clusters[*].clientConfigs	Specifies an optional client configuration bundle for a cluster. Configure this object and its child properties to add cluster-related files to the bundle. The files are made available for download on the Client Configurations tab on the Cluster Details page in the UI. The chart auto-generates the Pod volume and mount for Secrets referenced in clientConfigs.files[*].secretRef. You do not need manual extraVolumes or extraVolumeMounts entries. When omitted, the Client Configurations tab displays a message that no client configuration files are available for download.	null
clusterConfigs.clusters[*].clientConfigs.files	Files included in the bundle. Each entry declares a file name, MIME type, optional description, and a secretRef to an existing Kubernetes Secret holding the file content.	[]
clusterConfigs.clusters[*].clientConfigs.files[*].name	Specifies the file name shown in the file list and used inside the zip archive. This value is also the basename of the file on the Cloudera Surveyor Pod under /etc/surveyor/client-configs/[***CLUSTER NAME***]/, matching clusterConfigs.clusters[*].clusterName.	(required)
clusterConfigs.clusters[*].clientConfigs.files[*].contentType	Specifies the MIME type of a file in the client configuration bundle. The UI previews files with a text/* content type or application/x-pem-file. Other types are download-only.	(required)
clusterConfigs.clusters[*].clientConfigs.files[*].description	Specifies an optional human-readable description of a file in the client configuration bundle, displayed in the UI alongside the file name.	null
clusterConfigs.clusters[*].clientConfigs.files[*].secretRef.name	Name of the Kubernetes Secret in the Cloudera Surveyor release's namespace that contains the file content. The chart projects this Secret onto the Cloudera Surveyor pod automatically. Cert-manager–issued or operator-managed Secrets work in place — rotations propagate without a Helm upgrade.	(required)
clusterConfigs.clusters[*].clientConfigs.files[*].secretRef.key	Specifies the key in the Secret set by clusterConfigs.clusters[*].clientConfigs.files[*].secretRef.name that contains the file content.	(required)

Note: All Kafka clients used by Cloudera Surveyor are configured to use the `DirectoryConfigProvider`. This enables you to reference the contents of a file as a property value instead of hard-coding it. The references are replaced with the file's content. Use the following syntax to reference files:

```
${dir:[***FULL DIR PATH***]:[***FILENAME***]}
```

Files must be mounted as `Secret` before you can reference values. For more information, see *Registering Kafka clusters*.

Related Information

Registering Kafka clusters
Cloudera Surveyor configuration overview

Schema Registry Helm chart configuration reference

A reference listing all configurable properties of the Schema Registry Helm chart.

Configure these properties during installation in your helm install command, or update them after installation using helm upgrade. Use the --values (-f) and --set options to specify properties. For example:

```
helm install schema-registry \
  --namespace [***NAMESPACE***] \
  --values [***VALUES FILE***] \
  --set 'image.imagePullSecrets=[***REGISTRY CREDENTIALS SECRET***]' \
  oci://container.repository.cloudera.com/cloudera-helm/csm-operator/schema-
registry \
  --version 1.7.0-b248
```

Table 5: Schema Registry Helm chart configuration properties reference

Property	Description	Default Value
replicaCount	The ReplicaSet count.	1
image.registry	The container image registry where the image is hosted.	container.repository.cloudera.com
image.repository	The repository name of the container image.	cloudera
image.name	The name of the container image.	schema-registry
image.tag	The tag of the container image to use.	latest
image.pullPolicy	The policy to use when pulling the container image.	IfNotPresent
image.imagePullSecrets	An array of Secrets containing credentials that provide access to the registry where Schema Registry images are hosted.	[]
env	A list of environment variables to set in the Schema Registry Container. Can be specified as direct values or references to Secrets or Configmaps.	[]
database.type	The type of database to use. Valid values are postgresql and in-memory. When using postgresql, you must configure database .jdbcUrl, database.user, and database.pas sword properties. The in-memory option starts Schema Registry with an ephemeral in-memory database that requires no additional configuration. However, in-memory mode is only suitable for testing and evaluation as all schemas will be lost when Pods restart.	postgresql
database.jdbcUrl	The JDBC URL that points to your PostgreSQL database.	jdbc:postgresql://localhost:5432/schema_registry
database.user	The PostgreSQL username for Schema Registry database connections.	none
database.password.secretKeyRef.name	The name of the Secret containing the PostgreSQL database password.	none
database.password.secretKeyRef.key	The key in the Secret specified by database .password.secretKeyRef.name that contains the PostgreSQL database password.	none

Property	Description	Default Value
database.tls.secretRef	The name of a Secret containing TLS configuration for PostgreSQL connections (certificates, truststores, and so on). All keys from the Secret are mounted to /etc/schema-registry/postgres/tls. Reference mounted files in your JDBC URL (database.jdbcUrl) to configure SSL connections if SSL is required for PostgreSQL.	none
service.type	The type of Kubernetes Service used for exposing the Schema Registry application.	ClusterIP
service.port	The port exposed by the Service for accessing the Schema Registry application.	9090
tls.enabled	Enables or disables TLS for the Schema Registry application.	true
tls.keystore.secretKeyRef.name	The name of the Secret containing the TLS keystore.	none
tls.keystore.secretKeyRef.key	The key in the Secret specified by tls.keystore .secretKeyRef.name that contains the TLS keystore.	none
tls.keystore.password.secretKeyRef.name	The name of the Secret containing the TLS keystore password.	none
tls.keystore.password.secretKeyRef.key	The key in the Secret specified by tls.keystore .password.secretKeyRef.name that contains the TLS keystore password.	none
tls.keystore.type	The type of the TLS keystore. Can be PKCS12 or JKS.	PKCS12
additionalApplicationConnectorConfig	Additional Dropwizard configuration for the applicationConnector. For available properties, see https://www.dropwizard.io/en/stable/manual/configuration.html#http.	none
ingress.enabled	Enables or disables external access through Ingress.	false
ingress.className	The Ingress controller class.	none
ingress.extraAnnotations.*	Extra annotations to apply to the Ingress.	{nginx.ingress.kubernetes.io/backend-protocol: "HTTPS"}
ingress.rules.host	The host of the Ingress rule.	"none"
ingress.rules.path	The Ingress path.	"/"
ingress.tls.enabled	Enables or disables TLS for Ingress.	true
ingress.tls.secretRef	The name of the Secret containing Ingress TLS certificates.	none
podSecurityContext	Security context settings at the Pod level.	{}
containerSecurityContext	Security context settings at the Container level.	{}
resources	Resource requests and limits for the Schema Registry Container.	{}
scheduling.nodeSelector	Node selector constraints for Pod scheduling.	{}
scheduling.affinity	Affinity rules for more complex Pod scheduling. Includes nodeAffinity, podAffinity, and podAntiAffinity.	{}
scheduling.tolerations	Tolerations to allow Pod scheduling on nodes with matching taints.	[]

Property	Description	Default Value
scheduling.topologySpreadConstraints	Constraints for distributing Pods across the cluster topology.	[]
authentication.oauth.enabled	Enables OAuth authentication for the Schema Registry server.	true
authentication.oauth.jwt.principalClaimName	The name of the claim in the JWT token that contains the principal (username) used for authorization.	sub
authentication.oauth.jwt.expectedAudience	The expected audience value. If the JWT token contains an aud claim, it must match this value, otherwise the token is considered invalid.	none
authentication.oauth.jwt.expectedIssuer	The expected issuer value. If the JWT token contains an iss claim, it must match this value, otherwise the token is considered invalid. When configured, Schema Registry only accepts tokens issued by this specific issuer.	none
authentication.oauth.jwks.url	The URL to the JWKS endpoint.	none
authentication.oauth.jwks.configMapKeyRef.name	The name of a ConfigMap containing the JWKS output of your OAuth server. When configured, Schema Registry loads the JWKS from the mounted file instead of calling the JWKS endpoint. This is useful in air-gapped environments where JWKS endpoints are not accessible.	none
authentication.oauth.jwks.configMapKeyRef.key	The key in the ConfigMap specified by authentication.oauth.jwks.configMapKeyRef.name that contains the JWKS output.	none
authentication.oauth.jwks.tls.truststore.secretKeyRef.name	The name of the Secret that contains the truststore for accessing the JWKS endpoint. Configure this property if the backend of your JWKS has self-signed certificates.	none
authentication.oauth.jwks.tls.truststore.secretKeyRef.key	The key in the Secret specified by authentication.oauth.jwks.tls.truststore.secretKeyRef.name that contains the truststore for accessing the JWKS endpoint.	none
authentication.oauth.jwks.tls.truststore.type	The type of the TLS truststore for the JWKS endpoint. Can be PKCS12 or JKS.	PKCS12
authentication.oauth.jwks.tls.truststore.password.secretKeyRef.name	The name of the Secret that contains the truststore password for accessing the JWKS endpoint.	none
authentication.oauth.jwks.tls.truststore.password.secretKeyRef.key	The key in the Secret specified by authentication.oauth.jwks.tls.truststore.password.secretKeyRef.name that contains the truststore password for accessing the JWKS endpoint.	none
authentication.oauth.additionalConfig	Additional OAuth configuration values. Configuration specified here is added directly to the YAML configuration of Schema Registry.	{}
logging.level	The level of the root logger in the Schema Registry application.	INFO
logging.config	Advanced logging configuration for packages, appenders, and other logging components.	{}

Property	Description	Default Value
schema.avro.compatibility	The default Avro schema compatibility mode applied when no compatibility setting is provided during schema creation. Valid values: BACKWARD, FORWARD, BOTH, NONE.	BACKWARD
schema.avro.validationLevel	The default Avro schema validation level applied when no validation level is provided during schema creation. Valid values: ALL, LATEST.	ALL
authorization.simple.enabled	Enables or disables authorization.	true
authorization.simple.adminUsers	A list of admin usernames. Admin users can perform any operation in Schema Registry.	[]
authorization.simple.readOnlyUsers	A list of read-only usernames. Read-only users can only perform read operations in Schema Registry.	[]