

Cloudera Runtime 7.2.17

Monitoring Kafka Connect using Streams Messaging Manager

Date published: 2019-08-22

Date modified: 2023-06-27

CLouDERA

<https://docs.cloudera.com/>

Legal Notice

© Cloudera Inc. 2024. 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

Kafka Connect in SMM.....	4
The Kafka Connect UI.....	4
Deploying and managing Kafka Connect connectors in SMM.....	8
Deploying a Kafka Connect connector in SMM.....	8
Pausing, resuming, restarting, and deleting a Kafka Connect connector in SMM.....	12
Reconfiguring Kafka Connect connectors in SMM.....	13
Connector configuration features in SMM.....	14

Kafka Connect in SMM

Get started with Kafka Connect in SMM.

Kafka Connect is a tool for streaming data between Apache Kafka and other systems in a reliable and scalable fashion. Kafka Connect makes it simple to quickly define connectors that move large collections of data into and out of Kafka. Source connectors can ingest entire databases or collect metrics from all your application servers into Kafka topics, making the data available for stream processing with low latency. Sink connectors can deliver data from Kafka topics into secondary storage and query systems or into batch systems for offline analysis.

Kafka Connect in CDP is shipped with many different Cloudera developed as well as publicly available sink and source connectors. Each of which cover a specific use case for streaming data. In addition to the connectors available by default, installing custom developed or third-party connectors is also possible. All connectors can be deployed, managed, and monitored using the Streams Messaging Manager UI (recommended), Streams Messaging Manager REST API, or Kafka Connect REST API.


The following sections walk you through the Kafka Connect related pages of the SMM UI and give an introduction to the basics of connector management and monitoring. For a comprehensive list of connectors shipped in CDP, as well as detailed instructions on how to install connectors, see *Connectors*.

Related Information

[Connectors](#)

The Kafka Connect UI

Learn about the Kafka Connect section in the SMM UI, which you can use to deploy, manage, and monitor Kafka Connect connectors.

The  **Connect** section in SMM is your main hub in CDP where you deploy, manage, and monitor Kafka Connect connectors. This section of the UI is only available if you deployed Kafka Connect role instances under your Kafka service and SMM is configured to integrate with Kafka Connect. For detailed instructions on how to set up Kafka Connect and integrate it with SMM, see *Kafka Connect Setup*.

Cluster: KAFKA-1

Overview


Producers 7 | Brokers 3 | Topics 29 | Consumer Groups 4

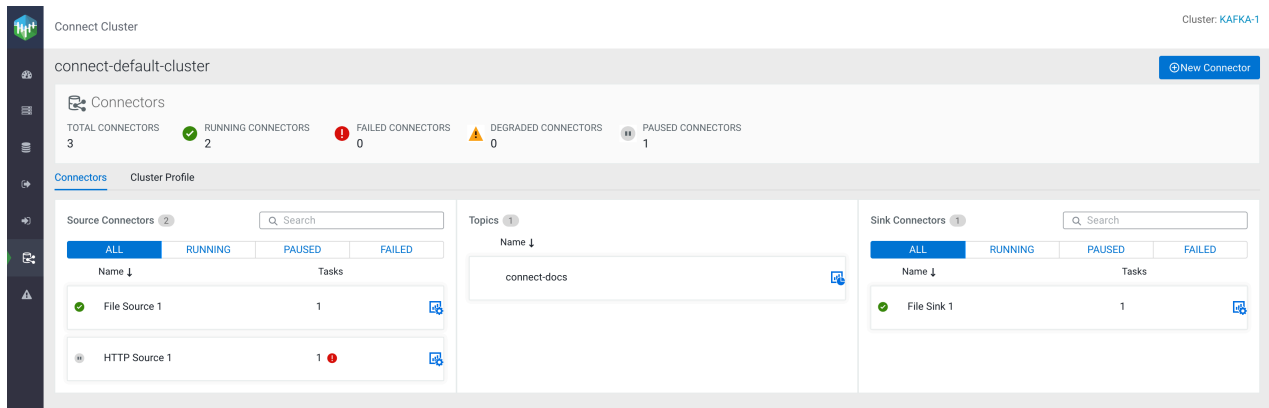
TOPICS 29 | BROKERS 3


NAME	DATA IN	DATA OUT	MESSAGES IN	CONSUMER GROUPS	CURRENT LOG SIZE
takeoffs	0B	0B	0	0	0B
takeoff	16 MB	36 MB	0.2m	0	21 MB
connect-status	0B	0B	0	0	0B
connect-secrets	0B	0B	0	0	463 B
connect-offsets	0B	0B	0	0	0B
connect-configs	0B	0B	0	0	5 KB
airport-weather	24 MB	57 MB	0.2m	0	36 MB
__ssb_sample_e9c68e60-5592-43be-bd7a-0288f1b17bf7	253 KB	0B	1.5k	0	332 KB
__ssb_sample_c696db70-69d9-4b94-afa0-6754f9707c53	10 KB	10 KB	55	1	14 KB

Producers (7) | ACTIVE (2) | PASSIVE (5) | ALL

Consumer Groups (4) | ACTIVE (0) | PASSIVE (4) | ALL

Clicking  Connect, opens the **Connect Cluster** page with the **Connectors** tab open. This is the default page of the Kafka Connect UI. It provides a high level overview of the connectors and the connect cluster.



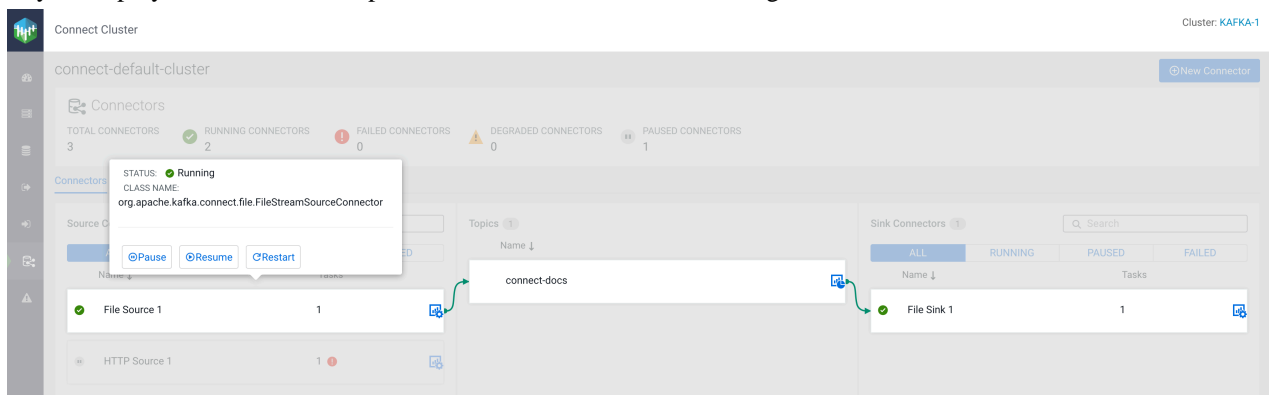
At the top-right corner of the **Connect Cluster** page, you can see the name of the cluster. The  New Connector option below the cluster name starts the **Connector Setup** wizard, which you use to deploy connectors in your cluster. For detailed steps on how to deploy a new connector, see *Deploying and Managing connectors*.

The **Connectors** section provides some basic metrics regarding Kafka Connect including total connectors, running connectors, failed connectors, degraded connectors, and paused connectors. Under the Connectors section you can find two tabs, Connectors and Cluster Profile. The two tabs enable you to monitor and manage your connectors and the connect cluster.

Connectors tab

In the Connectors tab, you can view details of the source connectors, topics, and sink connectors in the cluster. The Source Connectors and Sink Connectors sections show all, running, paused, and failed connectors with connector name and associated task details. Both Source Connectors and Sink Connectors sections contain a Search option that enables you to search for particular connector details. The Topics section shows the Kafka topics where data is read from or written to.

Hovering over a specific connector displays a context menu where you can view the class name and status of the connector. The context menu also provides quick options to pause, resume, and restart the connector. Clicking on any of your deployed connectors or topics visualizes how the data is flowing into or out of the selected entities.



Cluster Profile tab

In the Cluster Profile tab, you can view details of the connect cluster and the workers. The tab includes information regarding your cluster such as the Cloudera Manager URL, the number of workers deployed in the cluster, as well as an uptime counter. Additionally, the **Workers** section lists all worker hosts available in the cluster. This section includes basic worker-level information, such as how many connectors are deployed on a worker, success and failure rates, and more.

Connect Cluster Cluster: KAFKA-1

connect-default-cluster New Connector

Connectors

TOTAL CONNECTORS 3 ● RUNNING CONNECTORS 2 ● FAILED CONNECTORS 0 ▲ DEGRADED CONNECTORS 0 ● PAUSED CONNECTORS 1

Cluster Details

CM URL WORKERS 3 UPTIME 07:10:14

Workers

Search by host

Worker host	Connector Count	Connector Startup Attempts	Connector Startup Failures	Task Count	Task Startup Attempts	Task Startup Failures
+ [host]	1	5	0	1	0	0
+ [host]	1	9	0	1	0	0
+ [host]	1	5	0	1	0	0

If you click the **+** option found next to each host, additional details and metrics are displayed that are specific to the worker host you selected. The information includes a list of connectors assigned to the worker, as well detailed connector, task, and worker rebalance metrics. Using the Search by host field, you can search for worker details by host. Additionally, you can filter the list of connectors deployed on the worker host based on their status.

Connect Cluster Cluster: KAFKA-1

Cluster Details

CM URL WORKERS 3 UPTIME 07:15:14

Workers

Search by host

Worker host	Connector Count	Connector Startup Attempts	Connector Startup Failures	Task Count	Task Startup Attempts	Task Startup Failures
+ [host]	1	5	0	1	0	0
+ [host]	1	9	0	1	0	0
+ [host]	1	5	0	1	0	0

CONNECTOR PROFILE

CONNECTOR STARTUP ATTEMPTS: 5 (100% SUCCESSFUL)

TASK METRICS: 0 TASK STARTUP ATTEMPTS (0% SUCCESSFUL)

WORKER REBALANCE METRICS: 30 COMPLETED REBALANCES TOTAL, 33 EPOCH, CURRENTLY REBALANCING: False


CONNECTORS assigned to worker1

Name	Tasks	Type
File Source 1	1	source

REBALANCE AVG TIME: NA, REBALANCE MAX TIME: NA, TIME SINCE LAST REBALANCE: 1225804 ms

The Connector Profile

After you deploy a connector, you can monitor the details of the connector, modify its configuration, and carry out various management actions.

To monitor or manage a specific connector, click  (Connector Profile) next to the connector's name. Clicking the icon redirects you to the **Connector Profile** page of the selected connector. The **Connector Profile** page enables you

to monitor the details of the connector, monitor and manage its tasks, and edit its configuration. The page consists of two tabs **Connector Profile** and **Connector Settings**.

Connector Profile tab

The Connector Profile tab enables you to monitor details of the connector and its tasks. Additionally, a few basic management options are also available.

The screenshot displays the 'Connector Profile' page for 'File Sink 1' in a Cloudera Runtime environment. The page is organized into several sections:

- Connector Profile Summary:** Shows CLASSNAME, ASSIGNED WORKER, STATUS (RUNNING), TOTAL TASKS (1), RUNNING TASKS (1), FAILED TASKS (0), and PAUSED TASKS (0).
- Tasks Section:** Includes a search bar, a 'Restart' button, and a table of tasks with columns for Status, Worker ID, Task ID, Put Batch Avg Time, Sink Record Send Rate, and Partition Count.
- Running Ratio:** A gauge chart showing 100% RUNNING and 0% PAUSED.
- Offset Commits:** A gauge chart showing 100% SUCCESSFUL COMMITS and 0% FAILED COMMITS.
- Additional Sink Record Metrics:** Shows SINK RECORD LAG MAX as NA.
- Additional Sink Task Metrics:** Shows MAX NUMBER OF RECORDS PER BATCH as 4, AVG NUMBER OF RECORDS PER BATCH as 1.9444444444444444, OFFSET COMMIT MAX TIME as 4 ms, and OFFSET COMMIT AVERAGE TIME as 4 ms.


This tab is organized into multiple sections. The various sections enable you to do the following:

- The **Connector Profile** section provides you with details regarding the Classname, Assigned Worker, Status, Total Tasks, Running Tasks, Failed Tasks, and Paused Tasks.
- In the **Tasks** section, you can view and monitor Status, Worker ID, Task ID, and various other details regarding connector tasks.

Clicking **+** next to a task displays detailed information and metrics about the selected task. In addition to viewing status and metrics, the **Tasks** section also allows you to restart a particular task. This can be done by selecting the task you want to restart and clicking the Restart option found within the **Tasks** section. The Search by host option enables you to search for particular task details by host.

- Using the buttons in the top right-hand corner you can pause, resume, restart, and delete the connector, or deploy a new connector.

Connector Settings tab

The Connector Settings tab enables you to review and edit the configuration of the connector. By default editing the connector configuration is disabled. To enable editing, click  Edit in the bottom left-hand corner. In addition to reconfiguring the connector, you can pause, resume, restart, and delete the connector, or deploy a new connector with the buttons in the top right-hand corner of the page.



Tip: The **Connector Settings** tab is identical in functionality to the **Connector Configuration** page in the **Connector Setup** wizard. For detailed information on the various options and features available related to configuration, see *Connector configuration features in SMM*.

The screenshot displays the 'File Sink 1' connector configuration in the SMM UI. At the top, there are control buttons: Pause, Resume, Restart, Delete, and New Connector. Below these is a search bar and a 'Reset Filters' button. The main area shows a table of properties for the connector:

Properties (5 rows)		Actions
File Sink 1		
connector.class	org.apache.kafka.connect.file.FileStreamSinkConnector	ⓘ
file	/tmp/filestream.txt	ⓘ
key.converter	org.apache.kafka.connect.storage.StringConverter	ⓘ
topics	connect-docs	ⓘ
value.converter	org.apache.kafka.connect.storage.StringConverter	ⓘ

An 'Edit' button is located at the bottom left of the configuration area.

Related Information

[Kafka Connect Setup](#)

[Connectors](#)

[Deploying and managing Kafka Connect connectors in SMM](#)

[Connector configuration features in SMM](#)

Deploying and managing Kafka Connect connectors in SMM

Learn how to use the SMM UI to deploy new Kafka Connect connectors using the Connector Setup wizard, as well as how to reconfigure, pause, resume, restart, and delete existing connectors. Additionally, learn about the various features and options SMM provides for editing connector configurations.

Deploying a Kafka Connect connector in SMM



Kafka Connect connectors are deployed in SMM using the Connector Setup wizard. Learn how to deploy a new connector using the wizard.

Before you begin

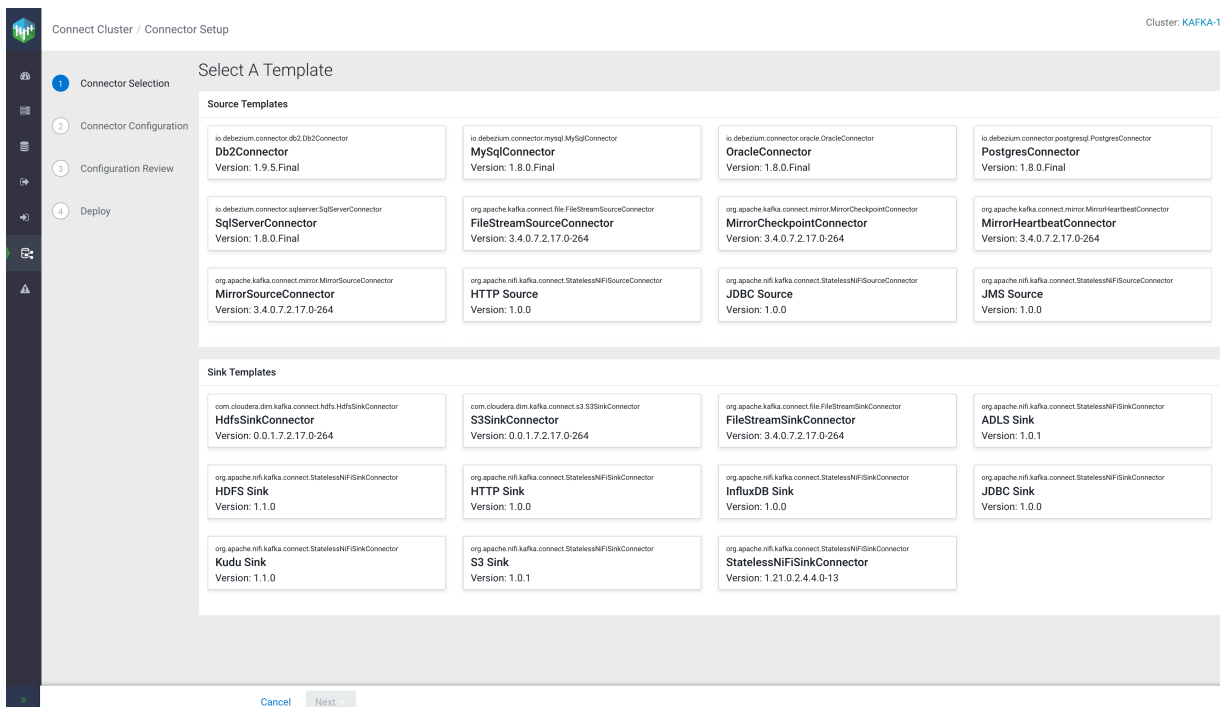
- By default, you can use the Connector Setup wizard to deploy any connector that is shipped with Cloudera Runtime. The only exceptions are the FileStream example connectors (FileStreamSourceConnector and FileStreamSinkConnector). Although these connectors are part of the Kafka distribution and are shipped with Cloudera Runtime, they are not available by default in the Connector Setup wizard. This is because neither connector is considered production ready. If you want to deploy an instance of these connectors, you must install them first. For more information, see [Installing Kafka Connect connectors](#).
- Third-party connectors can be deployed using the wizard, but must be installed first. For more information, see [Installing Kafka Connect connectors](#).

- Before deploying any connector, Cloudera recommends that you review the documentation of the connector. For more information regarding each connector shipped in CDP, see [Connectors](#).

Procedure

1. Click  **Connect** in the navigation sidebar.
2. Click the  **New Connector** option.

This option is available on all Kafka Connect related pages of the SMM UI. Clicking this option starts the **Connector Setup** wizard and redirects you to the **Select A Template** page. This page contains all connectors available for deployment in the form of selectable cards.



Each card includes the following information about the connector:

- The connector's fully qualified class name.
- The connector's display name. If no display name is available, the card includes the unqualified classname.
- The version of the connector.

By default, the page includes all connectors shipped with Cloudera Runtime. Third-party connectors that you install manually are also visible on this page following installation.

3. Select a connector from the **Source Templates** or **Sink Templates** section.

The Connector Configuration page appears.

Cluster: KAFKA-1

Connect Cluster / Connector Setup

Connector Selection

Connector Configuration

Configuration Review

Deploy

Search

Reset Filters

Validate Actions

Enter Connector Name

connector.class	org.apache.nifi.kafka.connect.StatelessNIFISourceConnector	
extensions.directory	/tmp/nifi-stateless-extensions	
key.converter	org.apache.kafka.connect.storage.StringConverter	
krb5.file	/etc/krb5.conf	
meta.smm.predefined.flow.name	HTTP Source	
meta.smm.predefined.flow.version	1.0.0	
nexus.url	https://repository.cloudera.com/artifactory/repo	
parameter.HTTP Source Parameters:Authorized Issuer DN Pattern	*	
parameter.HTTP Source Parameters:Authorized Subject DN Pattern	*	
parameter.HTTP Source Parameters:Base Path	contentListener	
parameter.HTTP Source Parameters:Client Authentication	NONE	
parameter.HTTP Source Parameters:Keystore Filename	Value	
parameter.HTTP Source Parameters:Keystore Key Password	Value	
parameter.HTTP Source Parameters:Keystore Password	Value	
parameter.HTTP Source Parameters:Keystore Type	Value	
parameter.HTTP Source Parameters:Listening Port	Value	
parameter.HTTP Source Parameters:Truststore Filename	Value	

Cancel Back Next

Most connectors shipped with Cloudera Runtime come with a default configuration template to ease configuration. If a template is available for a specific connector, the property keys and values are automatically populated when you select the connector. The properties and values included in the templates depend on the selected connector. In general, the templates include all mandatory properties that are required for successful deployment. However, most connectors also have a number of additional properties that might not be part of the template. As a result, Cloudera recommends that you always review the documentation for the specific connector that you want to deploy.

4. Configure the properties of the connector.

The Connector Configuration page includes various features and options that you can use to configure your connector. Each row in the **Properties** section represents the key and value of a property. Using the different buttons and other options available, you can add, delete, import, or otherwise modify the configuration. For more details on the configuration features available on this page, see *Connector configuration features in SMM*.

5. Click Validate after you are done configuring the connector.

Validating the configuration is mandatory when you deploy or modify a connector. If SMM finds any errors in the configuration, the properties that contain errors are highlighted in red, and an error message with the details regarding the configuration issue is displayed. Resolve any errors until validation passes.

6. Click Next.

The Next option is disabled until validation passes. After you click Next, the **Configuration Review** page appears.

Connect Cluster / Connector Setup Cluster: KAFKA-1

Configuration Review
Please take a moment to review your configuration.

Search Reset Filters

Properties (23 rows) Actions

Test HTTP Source	
connector.class	org.apache.nifi.kafka.connect.StatelessNIFISourceConnector
extensions.directory	/tmp/nifi-stateless-extensions
key.converter	org.apache.kafka.connect.storage.StringConverter
krb5.file	/etc/krb5.conf
meta.smm.predefined.flow.name	HTTP Source
meta.smm.predefined.flow.version	1.0.0
nexus.url	https://repository.cloudera.com/artifactory/repo
parameter.HTTP Source Parameters:Authorized Issuer DN Pattern	*
parameter.HTTP Source Parameters:Authorized Subject DN Pattern	*
parameter.HTTP Source Parameters:Base Path	Value
parameter.HTTP Source Parameters:Client Authentication	NONE
parameter.HTTP Source Parameters:Keystore Filename	Value
parameter.HTTP Source Parameters:Keystore Key Password	Value
parameter.HTTP Source Parameters:Keystore Password	Value
parameter.HTTP Source Parameters:Keystore Type	Value

Cancel ← Back Deploy →

7. Review your connector configuration.

The **Configuration Review** page allows you to review the connector configuration. Use this page to ensure that all properties are correctly configured for your use case.

Most configuration options are disabled on this page. However, you can use the search and filtering options to search for properties. Additionally, you can export the connector configuration for later use with **Actions Export**. If you find any errors, go back to the **Connector Configuration** page, make any necessary changes, and revalidate the configuration before continuing.


8. Click Deploy to deploy the connector.

After you click Deploy, the **Connector Deployment Status** modal is displayed that shows the status of the deployment.



Note: Connector deployment is not instantaneous. You might need to wait a few seconds for deployment to finish.


Connector Deployment Status

 Deploying...

Your connector is being deployed, you can visit the Connector Overview Page until it is deployed.

[Connector Overview Page](#)

Connector Deployment Status

 **Deployment Successful!**

Your connector was successfully deployed and is now available

[View Connector Profile](#)

What to do next

After the connector is deployed, you can monitor the connector on the Connect Cluster page or the connector's profile page. For more information, see *The Kafka Connect UI*.

Related Information

[Connectors](#)

[Streams Messaging Reference](#)


[The Kafka Connect UI](#)






[Connector configuration features in SMM](#)


Pausing, resuming, restarting, and deleting a Kafka Connect connector in SMM

Learn how to pause, resume, restart, or delete an existing Kafka Connect Connector in SMM.




Procedure

1. Click  **Connect** in the navigation sidebar.
2. Identify the connector you want to manage.

3. Click  (Profile) beside the connector.
4. Click  Pause,  Resume,  Restart, or  Delete.

If you select  Delete, you must confirm the action in a pop-up window.



Tip: The  Pause,  Resume, and  Restart options are also available if you hover over the connector on the **Connect Cluster** page.

Related Information

[Connectors](#)




[Streams Messaging Reference](#)

[The Kafka Connect UI](#)

Reconfiguring Kafka Connect connectors in SMM

Learn how to edit the configuration of a running Kafka Connect connector using the SMM UI.

Procedure

1. Click  **Connect** in the navigation sidebar.
2. Identify the connector you want to modify.
3. Click  (Profile) beside the connector.
4. Go to the Connector Settings tab.
5. Click  Edit at the bottom-left corner of the page.

All properties except connector name become editable.

The screenshot displays the SMM UI for a Kafka Connect connector named 'Test HTTP Source'. The interface includes a navigation sidebar on the left, a top header with 'Cluster: KAFKA-1', and a main content area. The main content area has a search bar, a 'Reset Filters' button, and a 'Validate' button. Below these is a table of properties for the connector. The table has columns for the property name, its value, and a set of action icons (Pause, Resume, Restart, Delete, and New Connector). The properties listed are:

Property Name	Value	Actions
connector.class	org.apache.nifi.kafka.connect.StatelessNIFISourceConnector	⏸ ⏪ ⏩ 🗑️ ➕
extensions.directory	/tmp/nifi-stateless-extensions	⏸ ⏪ ⏩ 🗑️ ➕
key.converter	org.apache.kafka.connect.storage.StringConverter	⏸ ⏪ ⏩ 🗑️ ➕
krb5.file	/etc/krb5.conf	⏸ ⏪ ⏩ 🗑️ ➕
meta.smm.predefined.flow.name	HTTP Source	⏸ ⏪ ⏩ 🗑️ ➕
meta.smm.predefined.flow.version	1.0.0	⏸ ⏪ ⏩ 🗑️ ➕
nexus.url	https://repository.cloudera.com/artifactory/repo	⏸ ⏪ ⏩ 🗑️ ➕
parameter.HTTP Source Parameters:Authorized Issuer DN Pattern	.*	⏸ ⏪ ⏩ 🗑️ ➕
parameter.HTTP Source Parameters:Authorized Subject DN Pattern	.*	⏸ ⏪ ⏩ 🗑️ ➕
parameter.HTTP Source Parameters:Base Path	Value	⏸ ⏪ ⏩ 🗑️ ➕
parameter.HTTP Source Parameters:Client Authentication	NONE	⏸ ⏪ ⏩ 🗑️ ➕
parameter.HTTP Source Parameters:Keystore Filename	Value	⏸ ⏪ ⏩ 🗑️ ➕
parameter.HTTP Source Parameters:Keystore Key Password	⏸ ⏪ ⏩ 🗑️ ➕
parameter.HTTP Source Parameters:Keystore Password	⏸ ⏪ ⏩ 🗑️ ➕
parameter.HTTP Source Parameters:Keystore Type	Value	⏸ ⏪ ⏩ 🗑️ ➕

At the bottom of the page, there are 'Cancel' and 'Update' buttons.

6. Configure the properties of the connector.

The Connector Settings page includes various features and options that you can use to configure your connector. Each row in the **Properties** section represents the key and value of a property. Using the different buttons and other options available, you can add, delete, import, or otherwise modify the configuration. For more details on the configuration features available on this page, see *Connector configuration features in SMM*.

7. Click Validate after you are done configuring the connector.

Validating the configuration is mandatory when you deploy or modify a connector. If SMM finds any errors in the configuration, the properties that contain errors are highlighted in red, and an error message with the details regarding the configuration issue is displayed. Resolve any errors until validation passes.

8. Click Update .

Related Information

[Connectors](#)

[Streams Messaging Reference](#)

[The Kafka Connect UI](#)

[Connector configuration features in SMM](#)

Connector configuration features in SMM



Learn about the various configuration features and options that you can use when configuring Kafka Connect connectors in SMM.

The SMM UI includes two pages where you configure Kafka Connect connector properties. The **Connector Configuration** step of the **Connector Setup** wizard and the Connector Profile Connector Settings tab. Both pages include various features and options that are designed to help you with connector configuration. The following sections go over each of the configuration options available on the UI.



Note: The only difference between **Connector Configuration** and **Connector Settings** is that the former is used to configure new connectors, the latter is used to reconfigure already running connectors. Otherwise, the two pages are identical in functionality.

Configuring properties

Each row on the **Connector Configuration** page represents the key (name) and configuration value of a specific property. You can use the icons next to each property to add or remove properties. Clicking  next to a property opens a context menu that includes additional configuration options. The options available in the  context menu depend on the property. For example, you can edit and reset the connector.class property, but you cannot configure its type.

Connect Cluster / Connector Setup Cluster: KAFKA-1

Connector Selection

Connector Configuration

Configuration Review

Deploy

Search Reset Filters


Properties (23 rows) Validate Actions

Test HTTP Source

connector.class	org.apache.nifi.kafka.connect.StatelessNIFISourceConnector	
extensions.directory	/tmp/nifi-stateless-extensions	
key.converter	org.apache.kafka.connect.storage.StringConverter	
krb5.file	/etc/krb5.conf	
meta.smm.predefined.flow.name	HTTP Source	
meta.smm.predefined.flow.version	1.0.0	
nexus.url	https://repository.cloudera.com/artifactory/repo	
parameter.HTTP Source Parameters:Authorized Issuer DN Pattern	*	
parameter.HTTP Source Parameters:Authorized Subject DN Pattern	*	
parameter.HTTP Source Parameters:Base Path	contentListener	
parameter.HTTP Source Parameters:Client Authentication	NONE	
parameter.HTTP Source Parameters:Keystore Filename	Value	
parameter.HTTP Source Parameters:Keystore Key Password	Value	
parameter.HTTP Source Parameters:Keystore Password	Value	
parameter.HTTP Source Parameters:Keystore Type	Value	
parameter.HTTP Source Parameters:Listening Port	Value	
parameter.HTTP Source Parameters:Truststore Filename	Value	

Cancel ← Back Next →


Adding properties


Click  to add a new property. An empty row appears for the new property. Enter property keys in the left field and property values in the right field. When you select a key field, a list of suggested properties appear that are valid for the connector. The list is automatically filtered when you start typing.

Deleting properties


Click  to delete a property. Use **Actions Remove** to remove all properties from the configuration with the exception of `connector.class`.

Editing properties


Click  **Edit** to edit the configuration value of a property in a separate window. Use this option if the value you add is long, complex, or consists of multiple lines.

If you are deploying a NiFi Stateless Sink or Source connector, the modal that opens with  for the `flow.snapshot` property is unique and includes the **Browse...** and **Save and Enhance** options. **Browse...** enables you to upload a flow definition JSON from your machine, **Save and Enhance** adds the parameters specified in your flow definition to the connector configuration. For more information on `flow.snapshot` configuration, see *Configuring flow.snapshot for Stateless NiFi connectors*.

Resetting properties


You can use  **Reset value** to reset the value of a single property. Alternatively, you can use **Actions Reset** to reset the full configuration. The reset options behave uniquely depending on the configuration page that you use them on.

When deploying new connectors using the **Connector Configuration** page of the wizard:



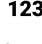

-  **Reset value** resets the value to the default value set in the configuration template. This option is not available for properties that are not part of the template.





- Actions Reset resets all configuration properties and reverts the configuration to the default configuration template.


When editing the configuration of existing connectors on the Connector Profile Connector Settings page:

-  Reset value resets the value to the last value that the connector was deployed with. This option is not available for new properties that you add to the configuration while editing the configuration.
- Actions Reset resets all configuration properties and reverts to the configuration that the connector was last deployed with.

Type configuration

The  String,  Boolean,  Number, and  Password options enable you to set the type of the property. Setting the type of a property changes the property's value field depending on the type you select.

-  String sets the type of the property to string and changes the input field to a single line text box.
-  Number sets the type to number and changes the input field to a text box that only allows negative and positive whole numbers and decimals.
-  Boolean sets the type of the property to boolean and changes the input field to a checkbox.
-  Password sets the type of the property to password.

Selecting  Password hides (locks) the value on the UI and stores it in a secure manner. Once the connector is deployed, properties marked as passwords are encrypted and stored securely in an internal Kafka topic. The actual values only resolve at runtime. These values cannot be retrieved or otherwise read from the configuration of the connector. In most cases, you use this option to hide passwords or other types of sensitive credentials found in connector configurations, however, you can use this option for any property value that your organization considers sensitive.



Tip: The password option is the UI implementation of the Kafka Connect Secrets storage feature. For more information on the feature, and how Kafka Connect handles values you mark as sensitive data, see *Kafka Connect Secrets Storage*.


Property search and filtering


At the top of the page you can find a search bar that you can use to search and filter for connector properties. Search only works for property keys, values are disregarded. When you search for a string, the **Properties** section of the page is automatically filtered and only the relevant properties are shown. The number of matches is displayed next to **Properties**.


In addition to searching, you can filter properties based on their group and importance. To do so, click the filter icon and select the relevant group and importance that you want to search for. To clear all applied filters, click Reset Filters.



Viewing property help

Hovering over  (Help), found next to some property values, displays information about that property. What information is displayed is property dependent, but at minimum you are presented with a description of the property.

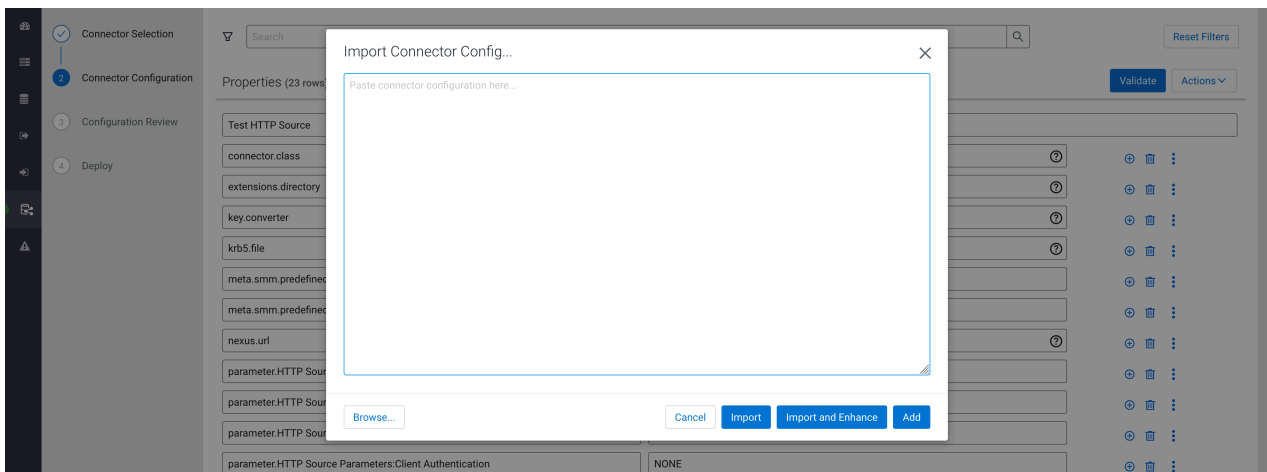
If available, the group, importance, and expected type of a property is also displayed. The  icon is only available for properties that have their metadata (description, type, group, and so on) defined.



Tip: The information displayed in the help context menu is the property metadata defined with `ConfigDef`. If you are developing your own connector and want help information to be available in SMM, ensure that you define property metadata in your connector code using the `ConfigDef` class. Additionally, if you specify a display name for the property in `ConfigDef`, SMM will print the display name above the key field of the property.

Importing configurations

To import a configuration, click **Actions Import**. Clicking **Import** opens the **Connector Config...** modal.



Using the modal you can:

- Type or paste your configuration.
- Use **Browse...** to upload the contents of a file from your machine.

The configuration that you add can be a full configuration that contains all properties necessary for the connector. Alternatively, it can be a partial configuration that contains a select number of properties that you want to add to the configuration. The configuration you include in the text box must be in JSON format. For example:

```
{
  "tasks.max": 1,
  "key.converter": "org.apache.kafka.connect.storage.StringConverter",
  "value.converter": "com.cloudera.dim.kafka.connect.converts.AvroConverter",
  "value.converter.passthrough.enabled": true,
  "value.converter.schema.registry.url": "http://schema-registry:9090/api/v1",
  "topics": "avro_topic",
  .
  .
  .
}
```

```
}

```

After adding properties, you can use the Import, Import and Enhance, or Add options to import the configuration. Each option imports the properties in a different way.

- Import

This option populates the **Properties** section with the keys and values you added in the **Import Connector Config...** modal. This option overwrites all existing properties in the configuration.

- Import and Enhance

This option populates the **Properties** section with the keys and values you added in the **Import Connector Config...** modal. Additionally, properties that are most likely needed for the respective connector are also added.

Import and Enhance is used specifically for the NiFi Stateless Source and NiFi Stateless Sink connectors. When this option is used, SMM parses the dataflow specified in `flow.snapshot`, extracts all parameters that are available in the dataflow, and adds them to the connector configuration. This option overwrites all existing properties in the configuration.


- Add

This option amends the existing configuration with the properties from the **Import Connector Config...** modal. Use this option if you want to batch import additional properties to your configuration. For example, you can use this option to batch add security properties.

Exporting Configurations

To export a configuration, click **Actions** **Export**. Clicking **Export** opens the **Configuration of the connector** modal.

- Click **Copy to Clipboard** to copy the configuration to the clipboard.
- Click **Export to file** to download the configuration as a JSON file.



Validating configurations

Before you can deploy or redeploy a connector, you must validate the configuration. This is done using **Validate**. If SMM finds any errors in the configuration, the properties that contain errors are highlighted in red, and an error message with the details regarding the configuration issue is displayed.

If the configuration is missing a mandatory property, the name of the missing property is displayed in an **Errors** section on the top of the page. You can add missing mandatory properties by clicking **Add missing configurations**.

If the connector configuration is filtered when you run the validation and SMM finds errors with properties that are filtered, an error message is displayed on the top of the page notifying you that not all validation errors are visible. In a case like this, click **Reset Filters** to view all validation errors.

The screenshot displays the Cloudera Runtime Connector Setup interface. The left sidebar shows the navigation menu with steps: Connector Selection, Connector Configuration (active), Configuration Review, and Deploy. The main area shows the configuration for a connector named 'Oracle'. A search bar and 'Reset Filters' button are at the top. Below, a message states 'Configuration is not valid. Reset filters to see all errors.' and 'Errors' are listed. The 'Properties (5 of 11)' section shows several configuration fields with error messages:

- database.dbname**: A value is required.
- database.hostname**: A value is required.
- database.password**: A value is required.
- database.port**: 1521.
- database.server.name**: has invalid format (only the underscore, hyphen, dot and alphanumeric characters are allowed). A value is required.
- database.user**: A value is required.

Buttons for 'Add Missing Configurations', 'Validate', and 'Actions' are visible. At the bottom, there are 'Cancel', 'Back', and 'Next' buttons.

Related Information

[Connectors](#)

[Streams Messaging Reference](#)

[The Kafka Connect UI](#)

[Kafka Connect Secrets Storage](#)

[Configuring flow.snapshot for Stateless NiFi connectors](#)