

Monitoring Kafka Clusters Using Streams Messaging Manager 7.2.15

# Monitoring Kafka Clusters Using Streams Messaging Manager

Date published: 2019-09-20

Date modified: 2022-05-12

The Cloudera logo is displayed in a bold, orange, sans-serif font. The word "CLOUDERA" is written in all caps, with the letter 'E' in "CLouDERA" featuring a horizontal line through its center.

<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

<b>Monitoring Kafka clusters.....</b>	<b>4</b>
<b>Monitoring Kafka producers.....</b>	<b>5</b>
<b>Monitoring Kafka topics.....</b>	<b>8</b>
<b>Monitoring Kafka brokers.....</b>	<b>8</b>
<b>Monitoring Kafka consumers.....</b>	<b>10</b>
<b>Monitoring lineage information.....</b>	<b>11</b>

## Monitoring Kafka clusters

The overview page provides you with tools to see a snapshot of the Kafka cluster you are monitoring. After you select the Kafka cluster to monitor, you can see the total number of producers, brokers, topics, and consumer groups in that cluster. You can also monitor producer and consumer metrics.

### Configure Apache Kafka for SMM

After you have installed and configured Apache Kafka, you must set one configuration parameter to enable Kafka and SMM to communicate.

1. Select Kafka from your cluster drop-down, and then select the Configuration tab.
2. Ensure that the Enable Producer Metrics check box is selected.

### Viewing cluster overview information

You can use the Overview tab to review information about your Kafka cluster. This page gives you information about total number of producers, brokers, topics, and consumer groups. It also provides more detailed metrics about producers and consumers.


Review the Producers, Brokers, Topics, and Consumer Groups information at the top of your page to understand how many of each are contained in your Kafka cluster.




The screenshot shows the 'Overview' page for a Kafka cluster named 'SMMDemo'. At the top, there are four blue boxes with white text and drop-down arrows: 'Producers 57', 'Brokers 5', 'Topics 28', and 'Consumer Groups 18'. To the right of these boxes is a 'Clear' button. Below the boxes, there are two links: 'TOPICS (28)' and 'BROKERS (5)'. In the bottom right corner, there is a refresh icon and a '30 minutes' dropdown menu. The page has a dark sidebar on the left with various icons.

You can click the drop-down arrow in any of the boxes to view a list of Kafka resource. Select one or more Kafka resource to filter your view to just those resource. You can also search for a specific resource. You can click clear at any time to return to the full overview.

## Overview



You can select the time period you want to view the metrics for, on the top-right corner of the page. If Cloudera Manager is configured as a metrics backend, the metrics (for example, topic > partition > producermetrics) which are used for time periods larger than 6 hours are calculated asynchronously, and take time to refresh.



## Monitoring Kafka producers

By monitoring Kafka producers, you can track the active and inactive producers in your cluster. You can also change the period of time after which a producer is considered inactive.

## Understanding producer naming conventions

The producers you interact with in Streams Messaging Manager (SMM) are named based on the `client.id` property you added when creating Kafka producers.

## Active vs. passive producers

On the Overview page, producers are referred to as active or passive. Producers are active when they are producing messages over a designated time period.

On the Producers page, passive producers are referred to as inactive.

You can set the period of time after which a producer is considered inactive in the Streams Messaging Manager Configs screen.

1. Select Streams Messaging Manager from the services pane.
2. Click Configs and select Advanced streams-messaging-manager-common from the Advanced tab.
3. Update `inactive.producer.timeout.ms` to change the period of time after which a producer is considered inactive. This value is specified in milliseconds.

STREAMS MESSAGING MANAGER CONFIG ADVANCED

### Advanced streams-messaging-manager-common

AMS's Kafka Application Id	kafka_broker	+	⌂
AMS's protocol	{{ams_timeline_metrics_protocol}}	+	⌂
ams.timeline.metrics.truststore.password	{{ams_metric_truststore_password}}	+	⌂
ams.timeline.metrics.truststore.path	{{ams_metric_truststore_path}}	+	⌂
ams.timeline.metrics.truststore.type	{{ams_metric_truststore_type}}	+	⌂
consumer.group.refresh.interval.ms	300000	+	⌂
inactive.group.timeout.ms	1800000	+	⌂
inactive.producer.timeout.ms	1800000	+	⌂

## Identifying a producer state

There are two ways to identify whether a producer is active or passive.

From the Producer pane in the Overview page, use the Active, Passive, and All tabs to view only active producers, only passive producers, or all of them. This allows you to see the total number of active and passive producers.

### Producers (84)




ACTIVE (57) PASSIVE (27) ALL

MESSAGES

geo-critical-event-coll...	7m
minifi-eu-i1	5.9m
load-optimizer-apps	3.2m
geo-critical-event-coll...	3m
fuel-apps	2.3m
minifi-eu-i2	1.8m

From the Producers page, each producer is listed with the status visible.

The screenshot shows three producer entries, each with a green hexagonal icon containing a white checkmark. The first entry is 'nifi-syndicate-speed-avro' with the status 'INACTIVE'. The second entry is 'geo-critical-event-collector-i19' with the status 'ACTIVE'. The third entry is 'nifi-syndicate-geo-avro' with the status 'INACTIVE'.


	nifi-syndicate-speed-avro INACTIVE
	geo-critical-event-collector-i19 ACTIVE
	nifi-syndicate-geo-avro INACTIVE

## Monitoring Kafka topics

By monitoring Kafka topics, you can track the total number of topics in your cluster and details about the topics. You can also monitor Grafana metrics for the topics in your cluster.

### Viewing the total number of topics in your cluster

You can see the total number of topics in your Kafka cluster on the Overview page.




### Detailed information about topics

The Topics page contains a number of useful details about your Kafka topics. This page helps you answer the following questions:

- How can I see if the replicas in this topic are in sync?
- How do I see this topic's retention rate?
- How can I see the replication factor for this topic?
- How do I see the producers and consumers that are connected to this topic?
- How do I find the total number of messages going into this topic, over a specified time range?

To access this detailed topic information:

1. From the left navigation pane, click Topics.
2. Identify the topic about which you want information. You can either scroll through the list of topics, or use the Search bar at the top left of the page.
3. Click the green hexagon at the left of the topic to view details.



## Monitoring Kafka brokers

By monitoring Kafka brokers, you can track various details about brokers including the host where the broker is located, disk space used by the broker, throughput, messages coming in, partitions, and replicas.

### Detailed broker information

The Brokers page contains a number of useful details about your Kafka brokers. This page helps you answer the following questions:



- On what host is my broker located?
- Is my broker running out of disk space?

To access detailed broker information:

1. From the left navigation pane, click Brokers.
2. Identify the broker about which you want information. You can either scroll through the list of brokers, or use the Search bar at the top left of the page.
3. Click the green hexagon at the left of the broker to view details.

The screenshot shows the Cloudera Manager interface for monitoring Kafka brokers. The main view is for broker 9, which is highlighted with a green hexagon. The broker's name is 'ihunyady-ns155-1.ihunyady-ns155.root.hwx.site:9092'. It has 0B throughput, 0 messages in, 30 partitions, and 93 replicas. Below the main view, there are several sections: 'FREE MEMORY', 'FREE DISK', 'CPU IDLE 30.21', 'LOAD AVERAGE 1.54', and 'DISK I/O 2947277.00'. A list of topics is shown below, including '\_smm-app-smm-producer-t...', 'file-sink-2', and several '\_consumer\_offsets' topics, each with its partition ID and '0B in'/'0B out' status.

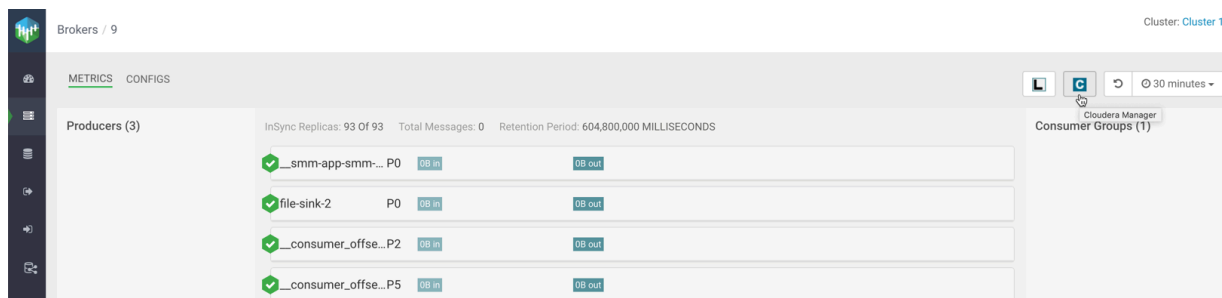
### Viewing additional details about the broker host

You can view additional details about the broker host from Cloudera Manager. To access this information:

1. From the left navigation pane, click Brokers.
2. Identify the broker about which you want information. You can either scroll through the list of brokers, or use the Search bar at the top left of the page.
3. Click the Profile icon on the right side of the broker view.

The screenshot shows the Cloudera Manager interface for monitoring Kafka brokers. The main view is a list of brokers. The list has columns for 'NAME', 'THROUGHPUT', 'MESSAGES IN', 'PARTITIONS', and 'REPLICAS'. Three brokers are visible: broker 9, broker 11, and broker 13. Each broker has a green hexagon icon on the left and a 'Profile' icon on the right. The 'Profile' icon for broker 9 is highlighted with a mouse cursor.

- Click the Cloudera Manager icon on the right side of the header.



## Monitoring Kafka consumers

By monitoring Kafka consumer groups, you can track active and passive consumer groups, or all consumer groups, which use the default internal `__consumer_offsets` topic to store the consumed offset information. You can track additional details about consumer groups. You can also track details including number of consumers and consumer instances included in a group and consumer group lag in the consumer group profile.

Streams Messaging Manager (SMM) displays consumer groups that have offsets stored in Kafka's internal topic `__consumer_offsets`, which is also the default store if the `auto.commit.enable` property is set to `true` for consumers. SMM does not display consumer groups that have offsets stored anywhere else other than this default store.

### Viewing summary information about consumer groups

The Overview page gives you summary information about consumer groups on the right side of the page. You can use the Active, Passive, and All tabs to view consumer groups only in the Active or Passives, or all of the consumer groups, which use the default internal `__consumer_offsets` topic to store the consumed offset information. Use the Lag tab to sort consumer groups based on ascending or descending amounts of lag.

NAME	DATA IN	DATA OUT	MESSAGES IN	CONSUMER GROUPS
syndicate-transmission	139MB	77MB	0.6m	0
syndicate-speed-even...	0B	0B	0	0
syndicate-speed-even...	0B	0B	0	0
syndicate-oil	166MB	0B	0.8m	0

NAME	LAG
fuel-micro-service	14m
supply-chain-micro-s...	9.2m
audit-micro-service	5m
adjudication-micro-se...	4.1m
load-optimizer-micro...	4.1m
energy-micro-service	2.4m
compliance-micro-ser...	1.7m

### Viewing details about a consumer group

To access detailed consumer group information:

- From the left navigation pane, click Consumer Group.
- Identify the consumer group about which you want information. You can either scroll through the list of consumer groups, or use the Search bar at the top left of the page.
- Click the green hexagon at the left of the consumer group to view details.



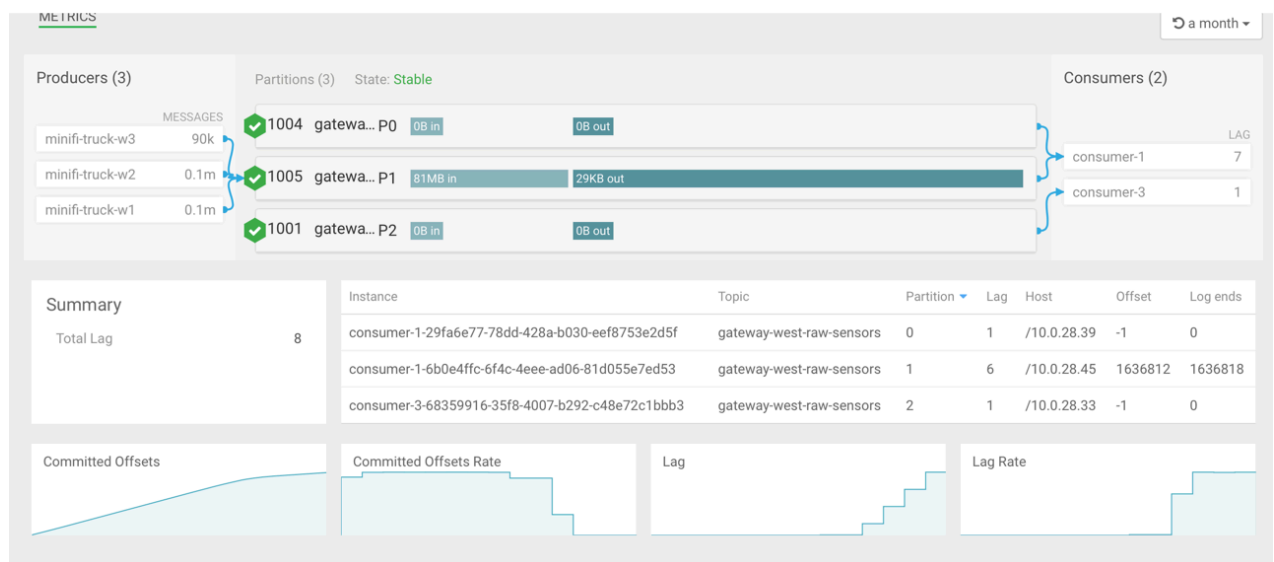
## Viewing the consumer group profile

The Consumer Group profile displays detailed information about each consumer group, including:

- The number of consumers included in the group.
- The number of consumer instances in the group.
- Details about consumer group lag.

To access the Consumer Group profile:

1. From the Consumer Group page, select the consumer group for which you want to view the profile.
2. Click the profile icon in the upper right of the Consumer Group tile.



## Monitoring lineage information

Learn how you can visualize the lineage between producers and consumers.

To check which topics a producer is producing to, and which consumers consume from those topics, go to the Overview page and click on a single producer on the Producer pane. For example, click producer\_1, as shown in the following image:

The screenshot shows the 'Overview' page of the Kafka Streams Messaging Manager. At the top, there are four summary cards: Producers (4 of 9), Brokers (3 of 3), Topics (5 of 31), and Consumer Groups (2 of 4). Below these, there are tabs for 'TOPICS (5)' and 'BROKERS (3)'. The main table lists topics with columns for NAME, DATA IN, DATA OUT, MESSAGES IN, and CONSUMER GROUPS. The 'a\_topic' row is highlighted. A tooltip for 'producer\_1' is open, showing 'Producer: producer\_1', 'MESSAGES 91', and a 'PROFILE FILTER' link. A green arrow points from the 'producer\_1' tooltip to the 'a\_topic' row. Another green arrow points from the 'a\_topic' row to the 'Consumer Groups (2)' panel on the right, which shows 'group\_1' and 'group\_2' with a LAG of 0.

After you click `producer_1`, you can see that it produces to a topic called `a_topic`, and that both consumer groups (`group_1` and `group_2`) consume from that topic.

This works the other way round as well. If you click on a single consumer group, you see what topics it consumes from and which producers produce to those topics. For example, click `group_1`, as shown in the following image:

The screenshot shows the 'Overview' page of the Kafka Streams Messaging Manager. At the top, there are four summary cards: Producers (4 of 9), Brokers (3 of 3), Topics (5 of 31), and Consumer Groups (2 of 4). Below these, there are tabs for 'TOPICS (5)' and 'BROKERS (3)'. The main table lists topics with columns for NAME, DATA IN, DATA OUT, MESSAGES IN, and CONSUMER GROUPS. The 'a\_topic' row is highlighted. A tooltip for 'Consumer Groups: group\_1' is open, showing 'Consumer Groups: group\_1', 'LAG 0', and a 'PROFILE FILTER' link. A green arrow points from the 'group\_1' tooltip to the 'a\_topic' row. Another green arrow points from the 'a\_topic' row to the 'Producers (4)' panel on the left, which shows 'producer\_1' and 'producer\_2' with a LAG of 91.

After you click `group_1` consumer group, you can see that it consumes from the topic called `a_topic`, and that two producers produce to that topic (`producer_1` and `producer_2`).

If you are interested in a more detailed view and want to check the lineage information for a single partition, you can do that as well, however, it is important to note that the lineage information is provided exclusively for the last 30 minutes. For example, click `P3`, as shown in the following image:

The screenshot shows the 'Overview' page for a Kafka cluster. On the left, there are 'Producers (4)' listed: 'producer\_1' and 'producer\_2' both with 91 messages. In the center, a table lists topics: 'heartbeats', 'connect-status', 'connect-offsets', and 'a\_topic'. The 'a\_topic' row is expanded to show its partitions: P0, P1, P2, P3, and P4. Partition P3 is selected, and a tooltip shows its details: 'Topic: a\_topic - P3', 'Lineage is shown only for the recently active clients!', 'DATA IN 4800', and 'DATA OUT 4204'. On the right, 'Consumer Groups (2)' are listed: 'group\_1' and 'group\_2', both with 0 lag. Red arrows trace the lineage from 'producer\_1' and 'producer\_2' to 'a\_topic - P3', and from 'a\_topic - P3' to 'group\_1' and 'group\_2'.

After you click P3 partition in the topic called `a_topic`, you can see that `producer_1` and `producer_2` produce to that partition, and `group_1` and `group_2` consume from it.

If you click the All Partitions button, you are shown the lineage information for every partition in a single topic.

This screenshot is similar to the previous one but shows the 'All Partitions' view for the topic 'a\_topic'. The tooltip for the selected partition (P3) now shows 'Topic: a\_topic' and 'Lineage is shown only for the recently active clients!'. Red arrows trace the lineage from 'producer\_1' and 'producer\_2' to all five partitions (P0, P1, P2, P3, P4) of 'a\_topic', and from all five partitions to 'group\_1' and 'group\_2'.

You can also access the lineage information from the experimental endpoints. You can find the endpoints at the [Streams Messaging Manager REST API Reference](#).

### Related Information

[Streams Messaging Manager REST API Reference](#)