

Monitoring Kafka Clusters Using Streams Messaging Manager 7.2.7

# Monitoring Kafka Clusters Using Streams Messaging Manager

Date published: 2019-09-20

Date modified: 2021-01-27

**CLOUDERA**

# 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>9</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.

## 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.



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

The screenshot shows the 'Producers' section of the SMM interface. At the top, a teal header bar displays the title 'Producers' in white and the number '57' in a large font. To the right of the number is a circular icon containing a white hand cursor. Below the header is a search bar with a placeholder 'Search' and a magnifying glass icon. A vertical scroll bar is visible on the right side of the list area. The main content is a list of producer names, each preceded by a small checkbox and a small square icon with a grid pattern. The names listed are: geo-critical-event-collector-i2, minifi-eu-i1, geo-critical-event-collector-i4, fuel-apps, minifi-eu-i3, supply-chain-apps, and geo-critical-event-collector-i7. The list ends with a faint, partially visible name.

## 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.

The screenshot shows the 'STREAMS MESSAGING MANAGER CONFIG' interface with the 'ADVANCED' tab selected. Below it, the section title 'Advanced streams-messaging-manager-common' is displayed. A table lists various configuration properties with their current values and edit icons:

|  |                                    |  |
|--|------------------------------------|--|
| 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.

|  |                                  |          |
|--|----------------------------------|----------|
|  | 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.

The screenshot shows the SMM Overview page for a cluster named 'SMMDemo'. It displays four main metrics: Producers (57), Brokers (5), Topics (28), and Consumer Groups (18). Below these, there are links for 'TOPICS (28)' and 'BROKERS (5)'. A time range selector at the bottom right indicates '30 minutes'.

### 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.

The screenshot shows the SMM Topics page for the topic 'gateway-west-raw-sensors'. It displays the following details:
 

- Producers (3):** minifi-truck-w3 (62k), minifi-truck-w2 (72k), minifi-truck-w1 (87k).
- Replication Factor:** (2)
- InSync Replicas:** 6 Of 6
- Total messages:** 268,184
- Retention Period:** 7 days

 The page also shows three replicas (P0, P1, P2) with their respective message counts and throughput metrics. On the right side, it shows consumer group information for 'nifi-truck-sensors-' with a lag of 0.4m.

## 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.

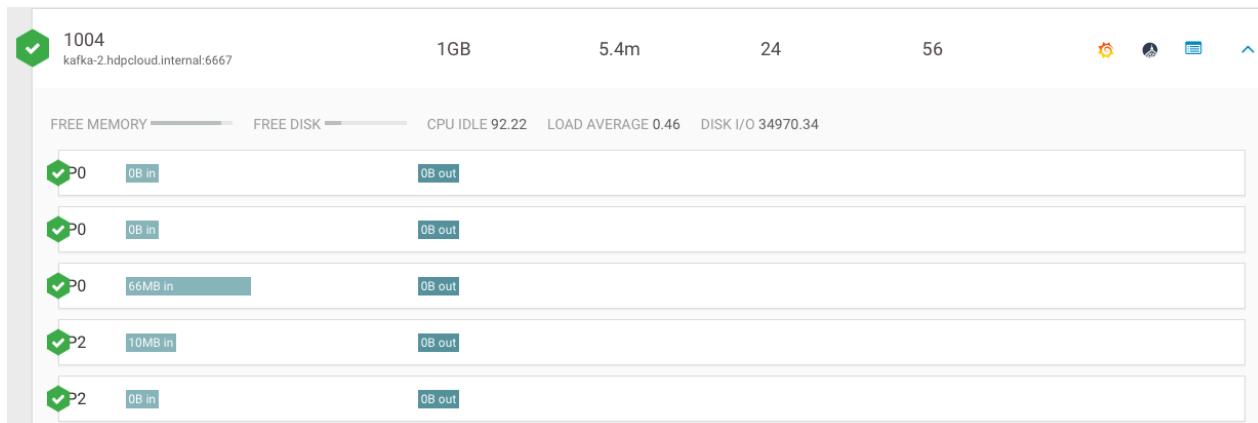
### Detailed broker information

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:

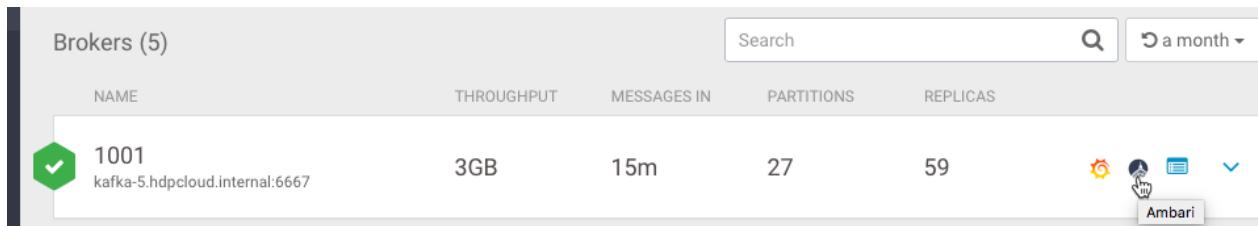
1. From the left navigation pane, click Brokers.
2. Identify the broker about which you want information. You can either scroll through the list of brokers, or use the Search bar at the top left of the page.
3. Click the green hexagon at the left of the broker to view details.



### Viewing additional details about the broker host

You can view additional details about the broker host from Ambari. 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 Ambari icon on the right side of the broker view.



## 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.

The screenshot shows the SMM Overview page with the following statistics:

- Producers:** 84
- Brokers:** 5
- Topics:** 28
- Consumer Groups:** 18

Below these stats, there are two main sections:

- TOPICS (28) BROKERS (5)**: A table listing topics with their names, data in/out, messages in, and consumer groups. One row is expanded to show details for "syndicate-transmission".
- Consumer Groups (18)**: A table listing consumer groups with their names and 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               |

| 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.

The screenshot shows the details for the consumer group "nifi-truck-sensors-west" (ACTIVE). It displays the following information:

- Partitions (3) State: Stable**
- Partition 1004:** gateway-west-r... P0 - 0B in, 0B out
- Partition 1005:** gateway-west-r... P1 - 81MB in, 29KB out
- Partition 1001:** gateway-west-r... P2 - 0B in, 0B out

### 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:

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

