Cloudera Runtime 7.2.17

Monitoring Kafka Clusters Using Streams Messaging Manager

Date published: 2019-08-22
Date modified: 2023-06-27

https://docs.cloudera.com/

Legal Notice

Contents

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.

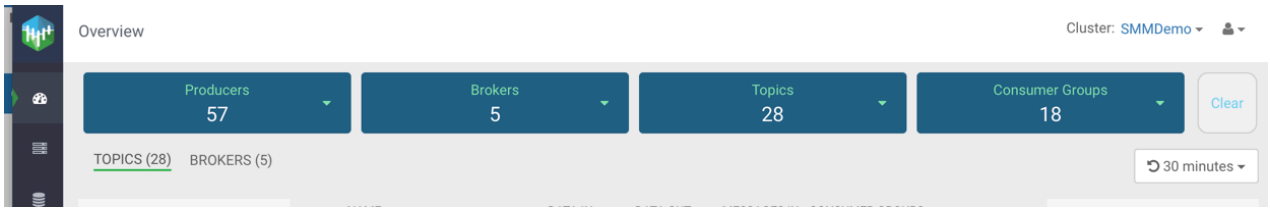

You can click the drop-down arrow in any of the boxes to view a list of Kafka resource. Select one or more Kafka resource to filter your view to just those resource. You can also search for a specific resource. You can click clear at any time to return to the full overview.

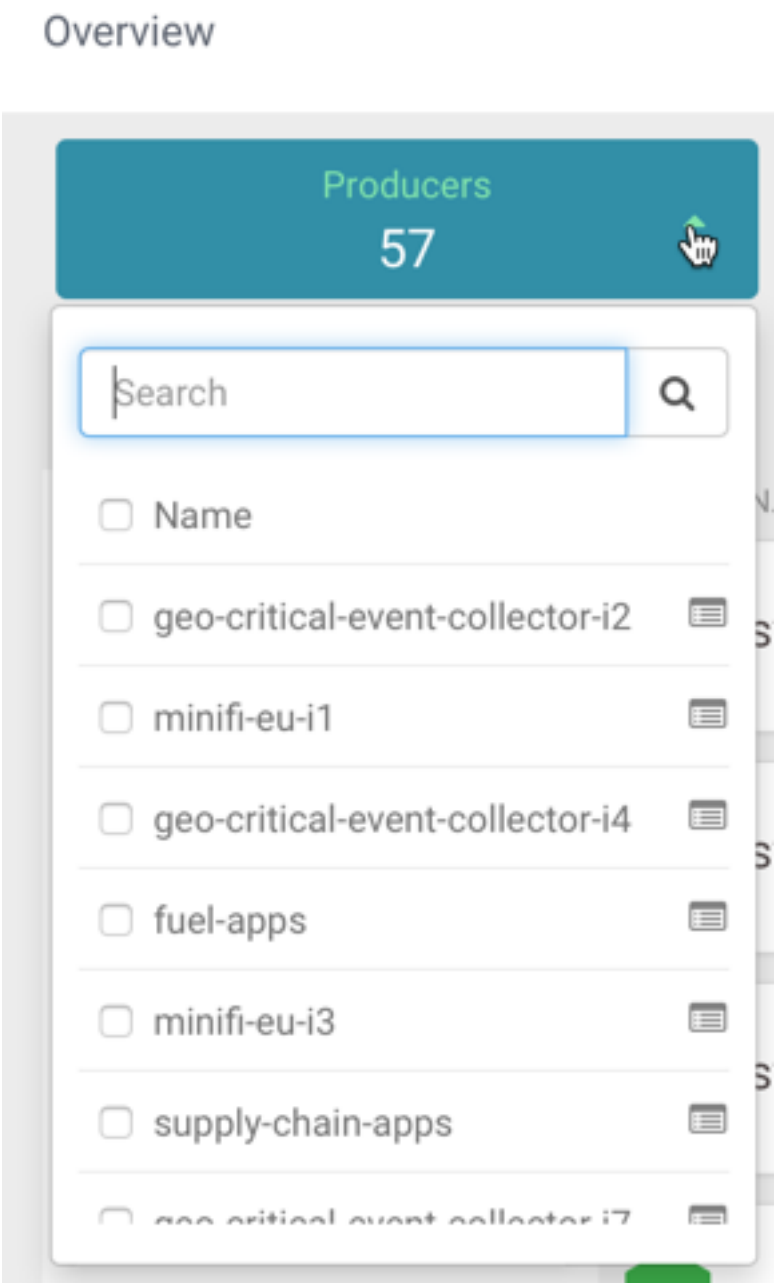

You can select the time period you want to view the metrics for, on the top-right corner of the page. If Cloudera Manager is configured as a metrics backend, the metrics (for example, topic > partition > producermetrics) which are used for time periods larger than 6 hours are calculated asynchronously,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

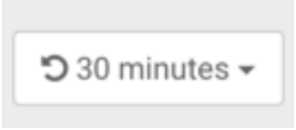

Monitoring Kafka producers

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.

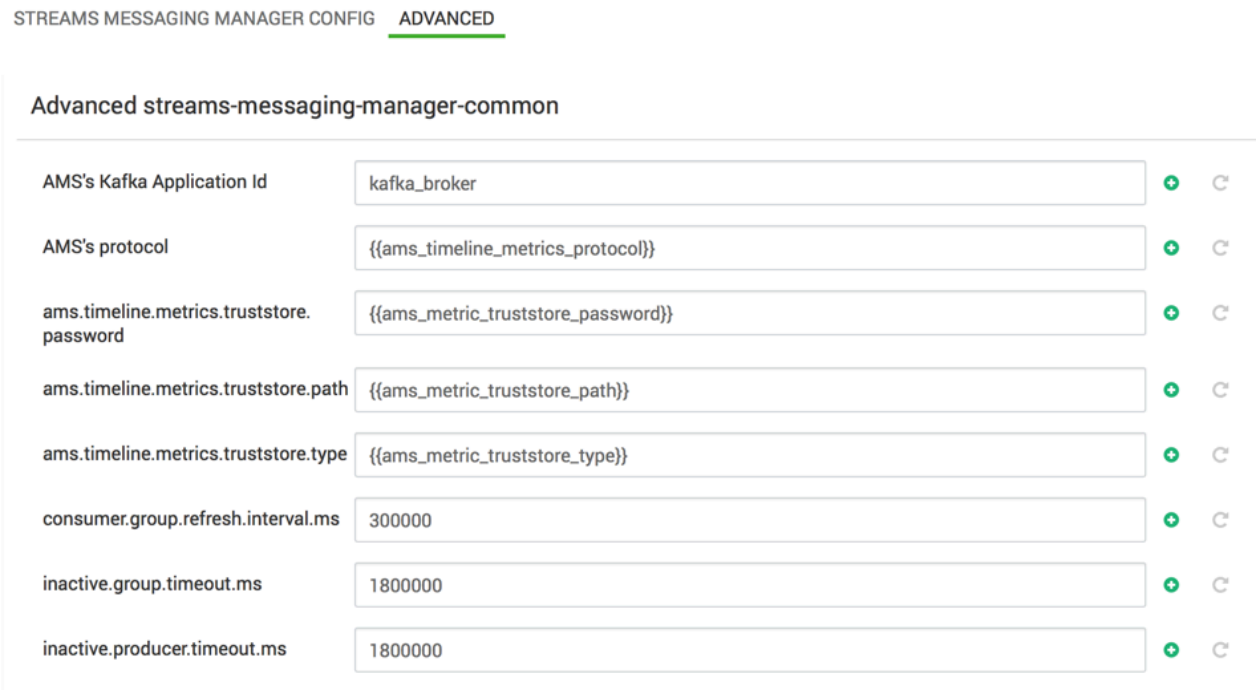

Identifying a producer state

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.

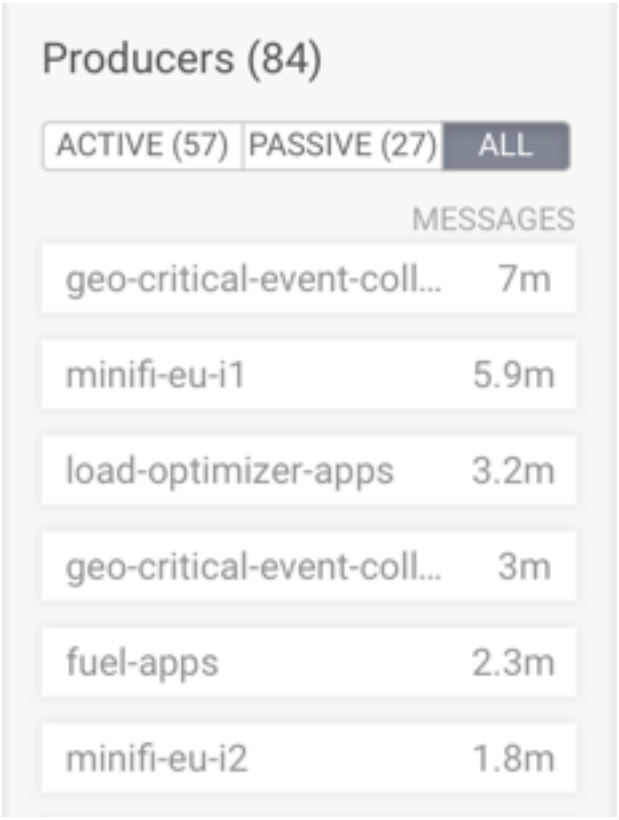

From the Producers page, each producer is listed with the status visible.

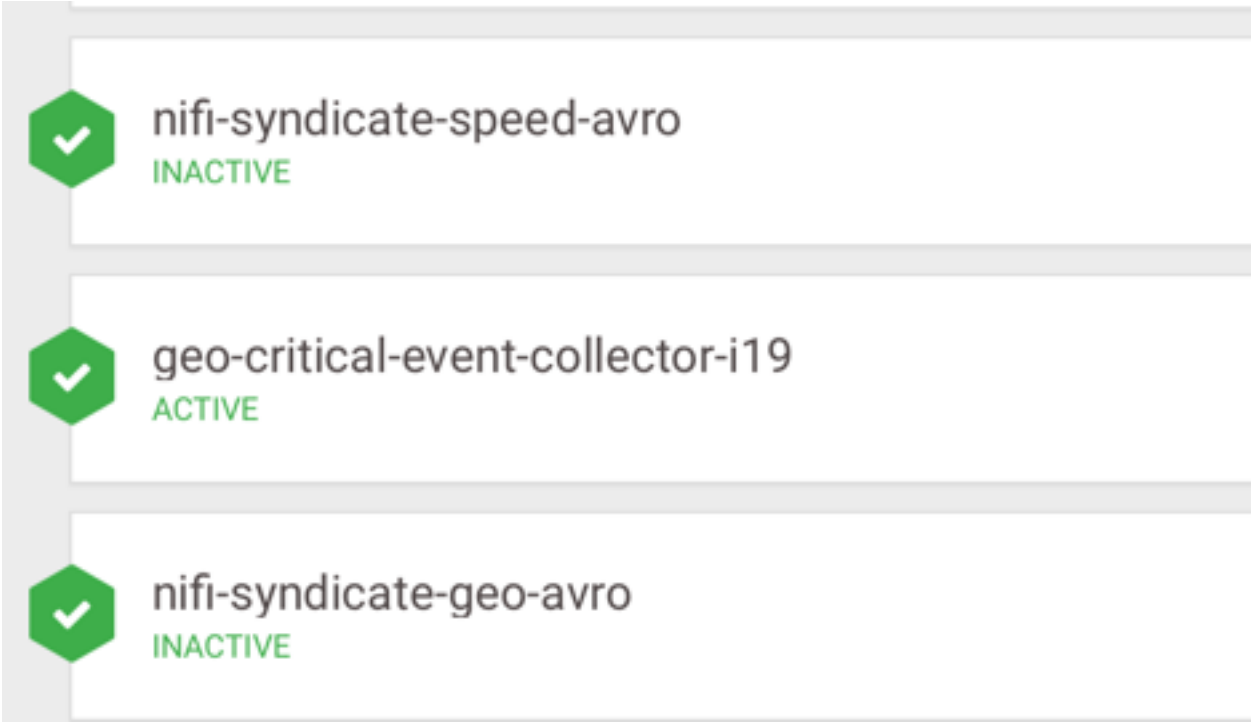

Monitoring Kafka topics

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.

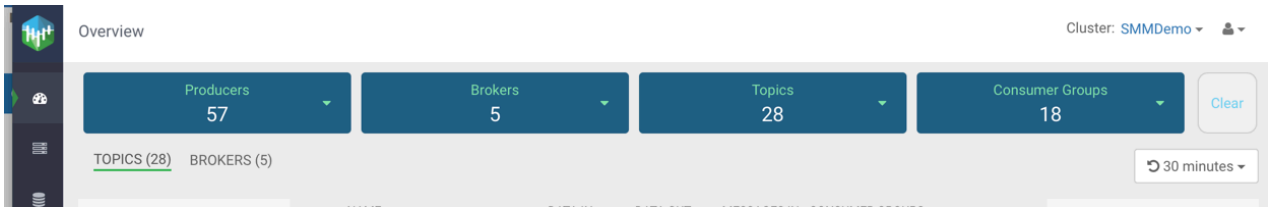

Detailed information about topics

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.

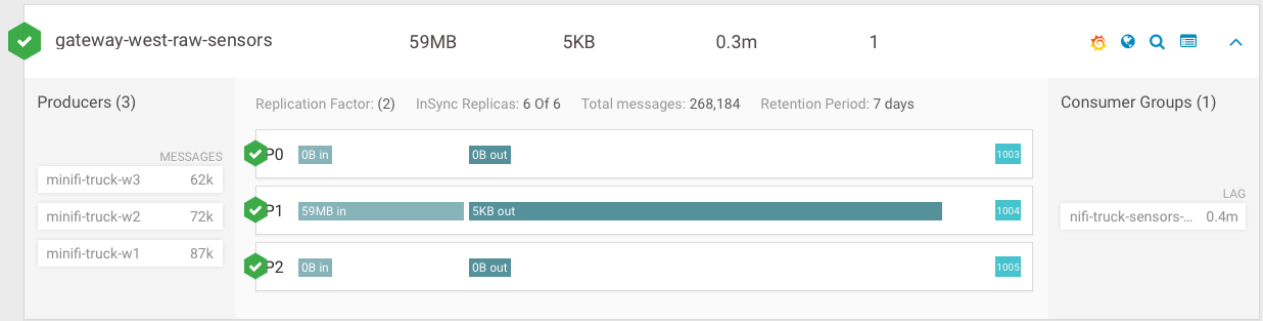

Viewing topic messages using Data Explorer

Data Explorer is a simple Kafka consumer within SMM. It enables you to view the content of a Kafka topic. You can select any Kafka topic and any partition within that topic, and view messages from the selected partition.

You can reach Data Explorer in two ways. One way is from the **Topics** page, and the other is from the **Overview** page. In both pages, you need to either click the magnifier icon, or navigate to the **Topic Details** page and then select the Data Explorer tab. The following steps describe the process:

1. Log in to the SMM UI.
2. From the left navigation pane, click Topics.

3. Identify the topic for which you want the message information. You can either scroll through the list of topics, or use the Search bar to find a topic.

4. Click the Data Explorer icon for that topic.

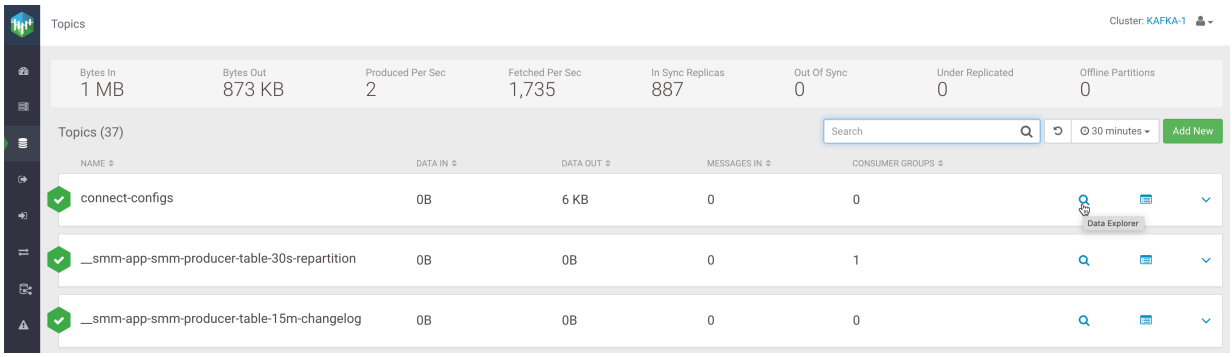

The Data Explorer dialog appears.

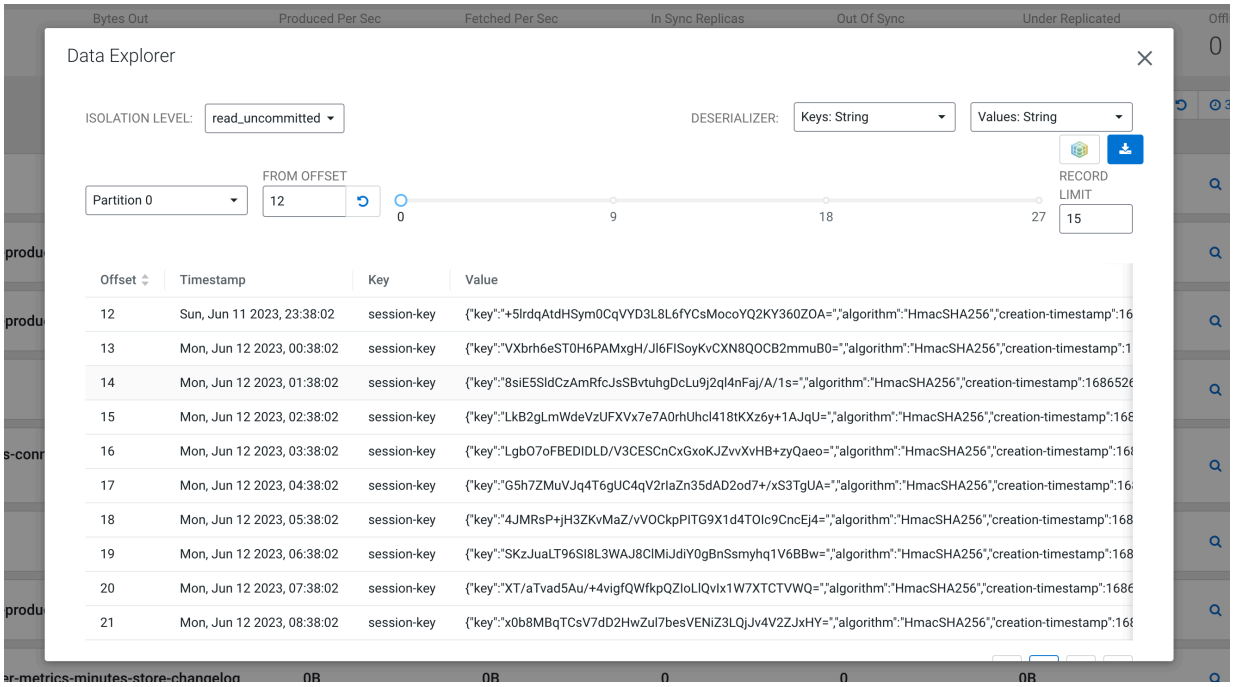

Alternatively, you can click the Profile icon for that topic.

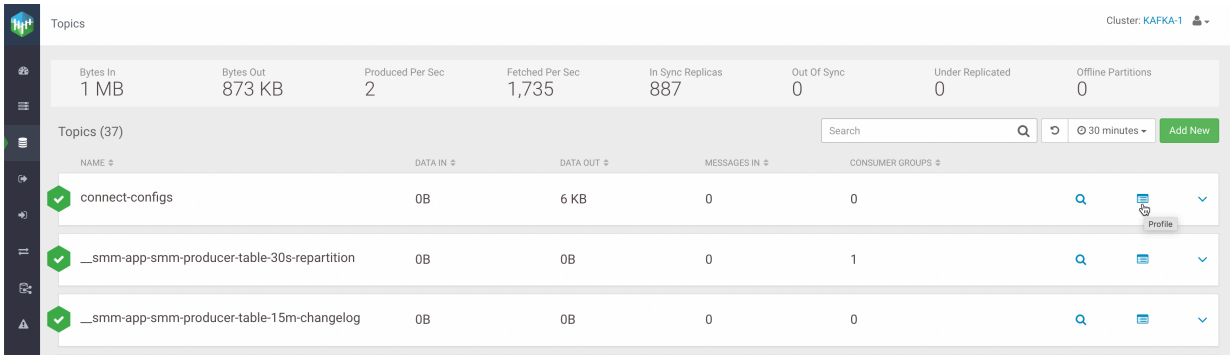

Then go to the Data Explorer tab.

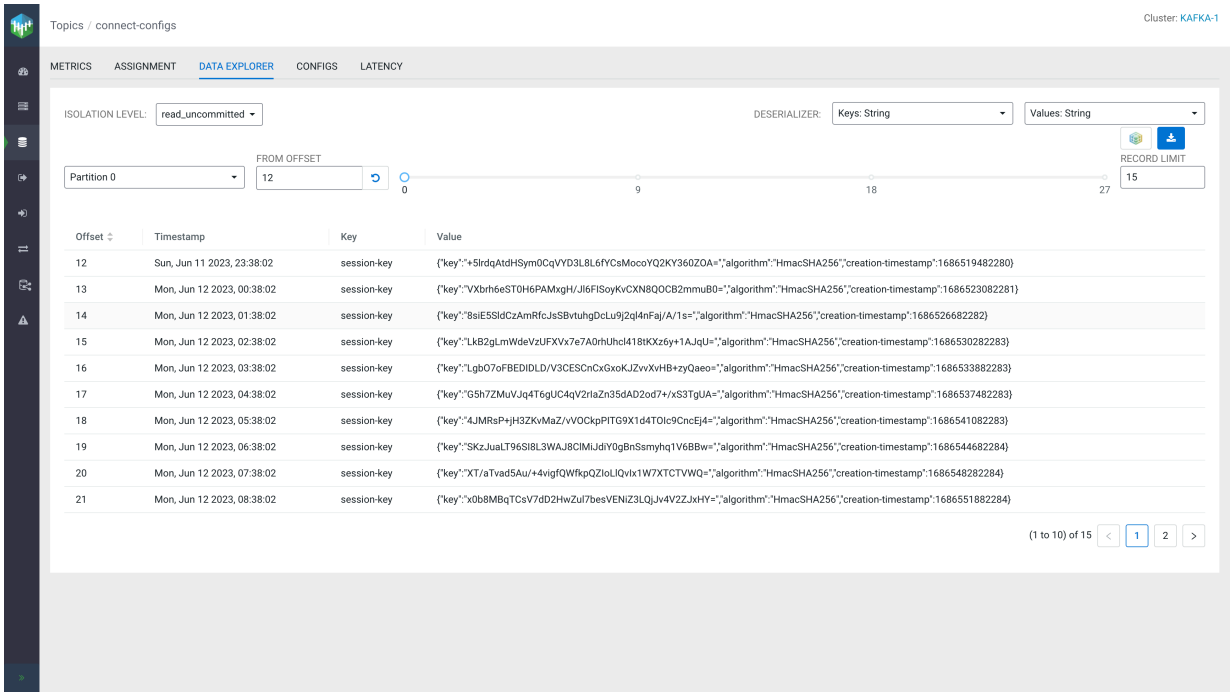

5. Select any of the following modes in the Isolation Level option:
   - read_committed
   - read_uncommitted

   The isolation level specifies whether uncommitted transactional messages should be read. By default, it is set to read_uncommitted.
6. Select the deserializer types for the Keys and Values options.

   For example, if you select Avro, SMM uses the schema that can be found in Schema Registry to deserialize the messages.
7. Select a Partition.

   The Kafka topic must have partitions to select from.
8. Select a value for the From Offset field.

   You can also use the selection bar to select an offset value. The maximum value is the offset of the last message.

   Click ↻ to refresh the partition offset range and to fetch the latest messages.
9. Select a Record Limit.

   The record limit value is the number of messages that are fetched starting from the message offset number selected in the From Offset field.
10. To see long messages, click show more beside a message.

    The message opens in a dialog or a new tab based on the size of the message.
11. Click the Schema Registry icon to go to the related page in the Schema Registry UI.

Increasing topic partition

You can increase the number of partitions of a topic

Warning: Increasing the partition numbers can impact the use of keys on messages. The data that is already written is not redistributed using the new partition, but remains on the partition on which it was located.

The following steps describe the process:

1. Log in to the SMM UI.
2. From the left navigation pane, click Topics.
3. Identify the topic for which you want to increase the topic partition number. You can either scroll through the list of topics, or use the Search bar to find a topic.
4. Click the Profile icon for that topic.
5. Go to the Configs tab.

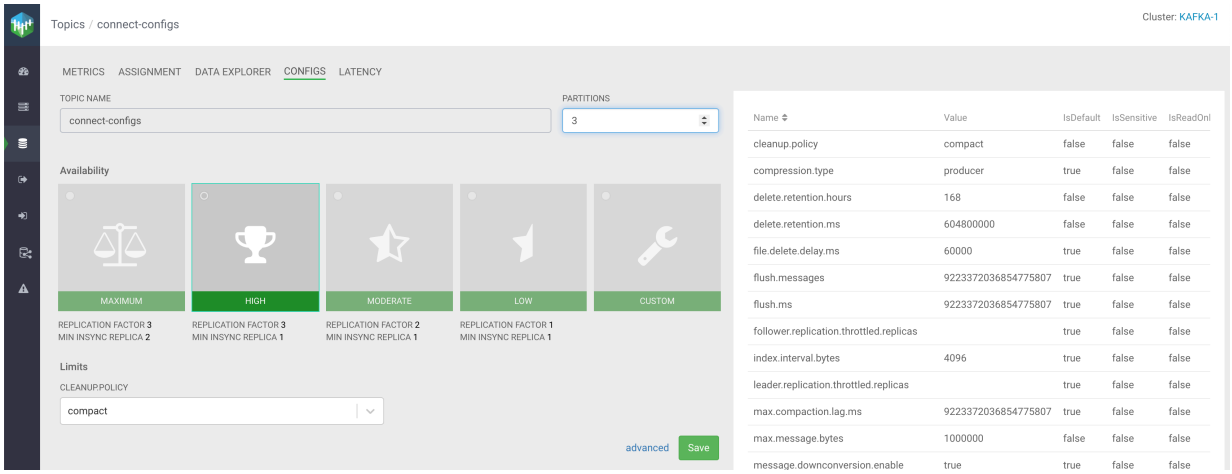

6. Enter the partition number in Partitions or increase it using the arrows.
7. Click Save.

   A dialog appears to ensure the increment.
8. Click Yes.
9. Refresh the page to see the change.

Getting current state of the topic partitions (Tech Preview)

The experimental Assignment tab, on the topic details page, shows the current state of the topic. It shows some topic-level statistics and the replica assignment of all partitions. If rack awareness is used in the Kafka cluster, the replica assignment is shown in a rack-based view. If the rack IDs follow the format of multi-level rack IDs, the rack IDs are rendered as a hierarchy. For more information on rack awareness, see *Kafka rack awareness*.

To go to the Assignment tab, click the Profile icon for a topic from the Overview or Topics page, and then click the Assignment tab. You can view the following statistics there:

- Number of offline partitions

  Shows the number of offline partitions in the topic. A partition is offline if it does not have a leader. Partitions can become offline if all their in-sync replicas are offline.
- Number of under-min-ISR partitions

  Shows the number of under-min-ISR partitions in the topic. A partition is in an under-min-ISR state if the number of in-sync replicas is lower than the value set in the min.insync.replicas property of the topic. The minimum in-sync replicas configuration defines how many replicas must acknowledge a produced message before the produce request is considered successful.

- Number of under-replicated partitions

  Shows the number of under-replicated partitions in the topic. A partition is under-replicated if it has at least one out-of-sync replica.
- Number of unevenly distributed partitions

  This appears if rack awareness is being used.

  Important: This statistics is based on standard rack awareness, and does not respect the multi-level rack IDs. This means that a partition might be shown as evenly distributed while in terms of multi-level rack awareness it should be considered unevenly distributed.

  Shows the number of unevenly distributed partitions in the topic. A partition is unevenly distributed if the difference between the maximum and minimum number of replicas in a rack is greater than one. This typically suggests that the partition does not meet expected durability guarantees or that it causes uneven load on the cluster. If a partition is unevenly distributed, try reassigning them. In most cases, unevenly distributed partitions become evenly distributed across the racks after reassignment.
- Number of unused racks

  This appears if rack awareness is being used.

  Shows the number of racks which are currently not used by this topic. A rack is unused if the topic has no replicas residing in that rack. This typically suggests that the partition does not meet expected durability guarantees because it is not using all available racks (physical locations) to store replicas of the data. This is expected and does not cause issues for non-critical topics that have an intentionally low replication factor.

In the replica assignment table, replicas are shown with different colors:

- Leader replicas are green
- In-sync replicas are blue
- Out-of-sync replicas are red

In addition, offline partitions and out-of-sync partitions are red. Under-replicated partitions, unevenly distributed racks, and unused racks are orange. When a specific partition or rack has one of these issues, a warning sign appears next to it.

Leader and in-sync replicas also act as links to the corresponding broker details page.

In the table header and in the first column, warning icons are shown if the specific column or row is affected by one of the issues listed in the topic statistics.

The following image shows the Assignment tab without racks:

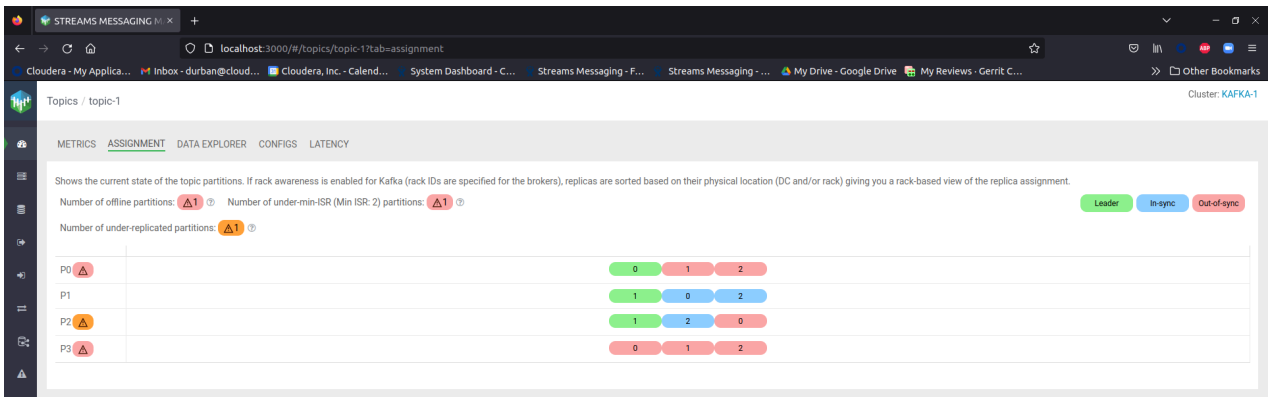

The following image shows the Assignment tab with racks:

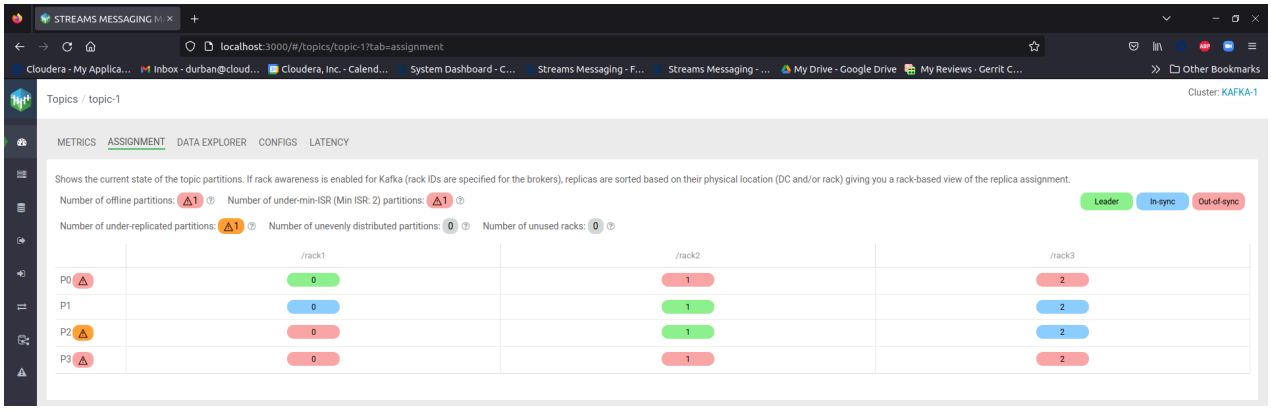

The following image shows the Assignment tab with racks, but with one broker metadata unknown, which adds an extra Unknown column to the table:

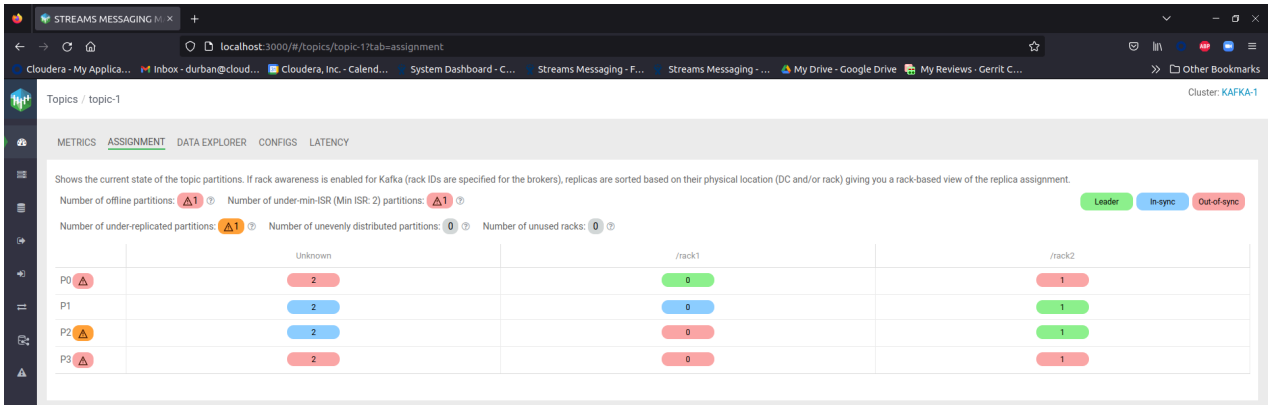

Related Information

Kafka rack awareness

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.

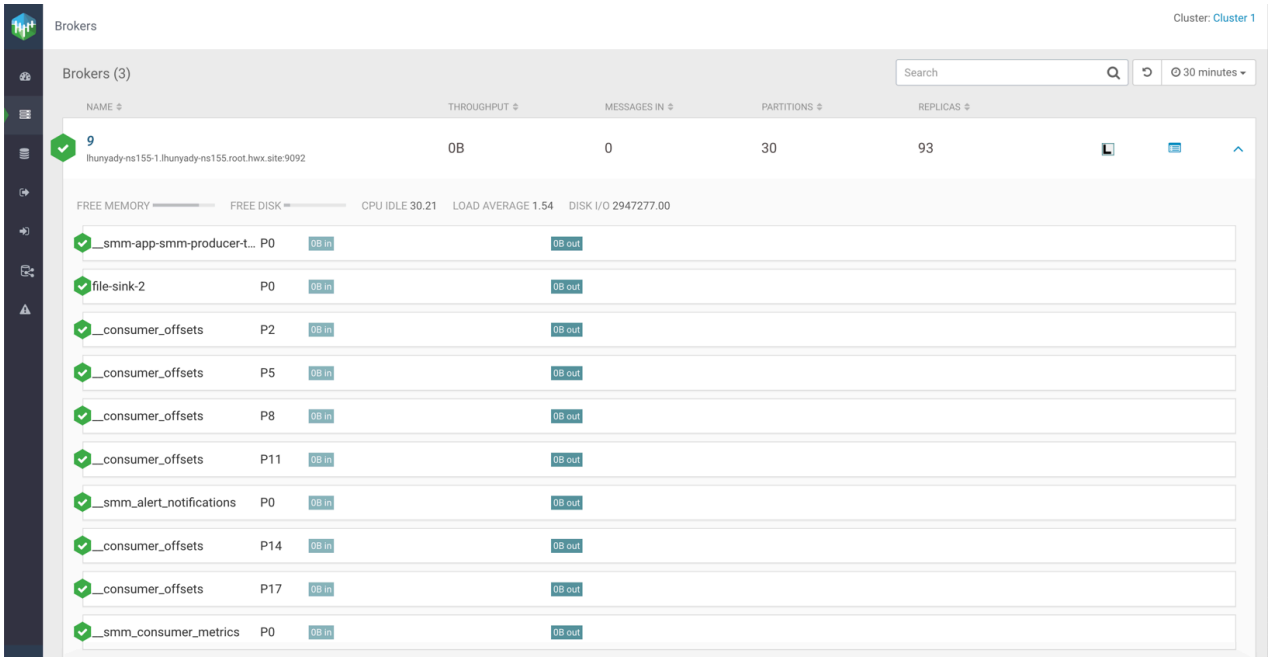

Viewing additional details about the broker host

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

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.

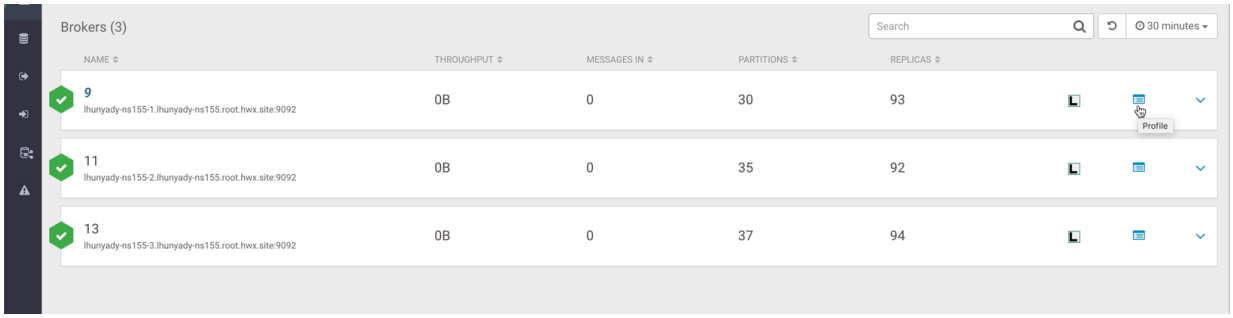

4. Click the Cloudera Manager icon on the right side of the header.

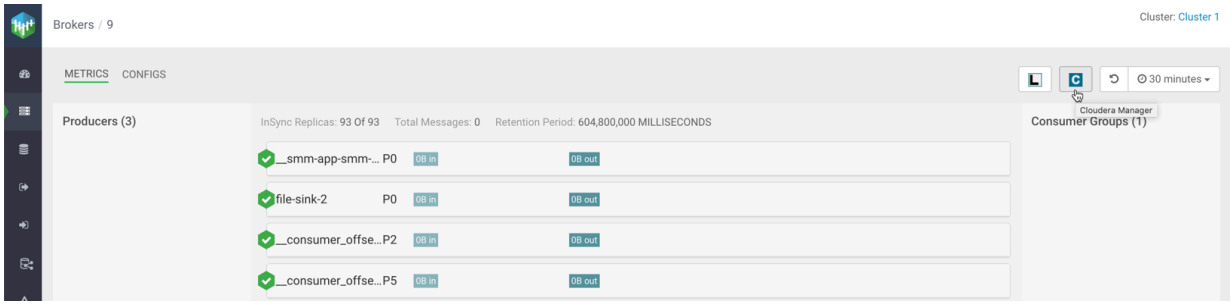

You can track additional details about the broker host in the Metrics and Configs tabs.

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.

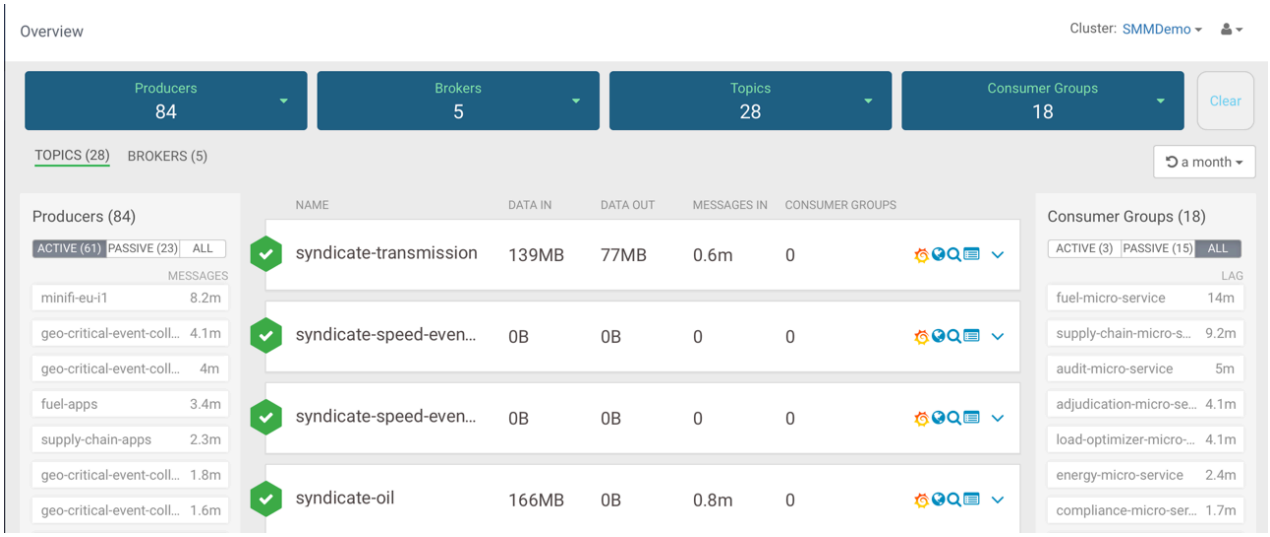

Viewing details about a consumer group

To access detailed consumer group information:

1. From the left navigation pane, click Consumer Group.
2. 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.
3. 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.

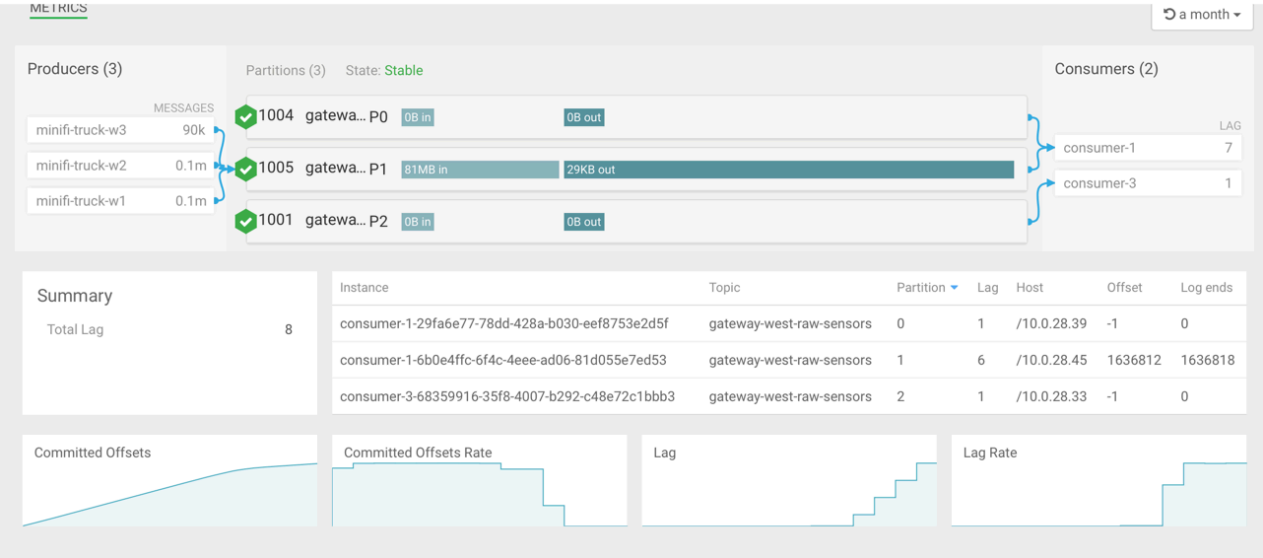

Resetting consumer offset

To reset the offset of a consumer group, perform the following steps:

1. From the left navigation pane, click Consumer Group.
2. Choose the consumer group for which you want to reset the offset, and click the Profile icon.

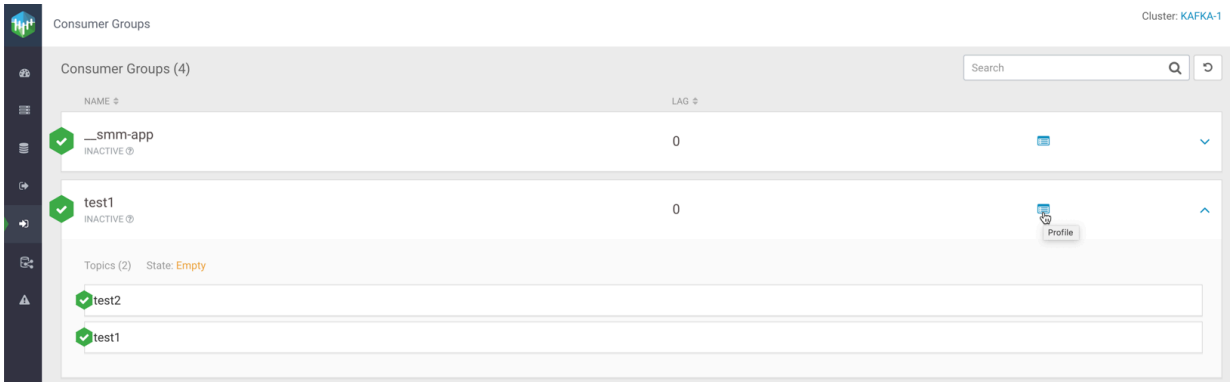

Note: Resetting offsets is only possible for those groups whose state is Empty or Dead. Attempting to reset offset of any other group results in an error.

3. In the Metrics page, click Actions Reset offset .

   The Reset offsets for consumer group dialog appears.

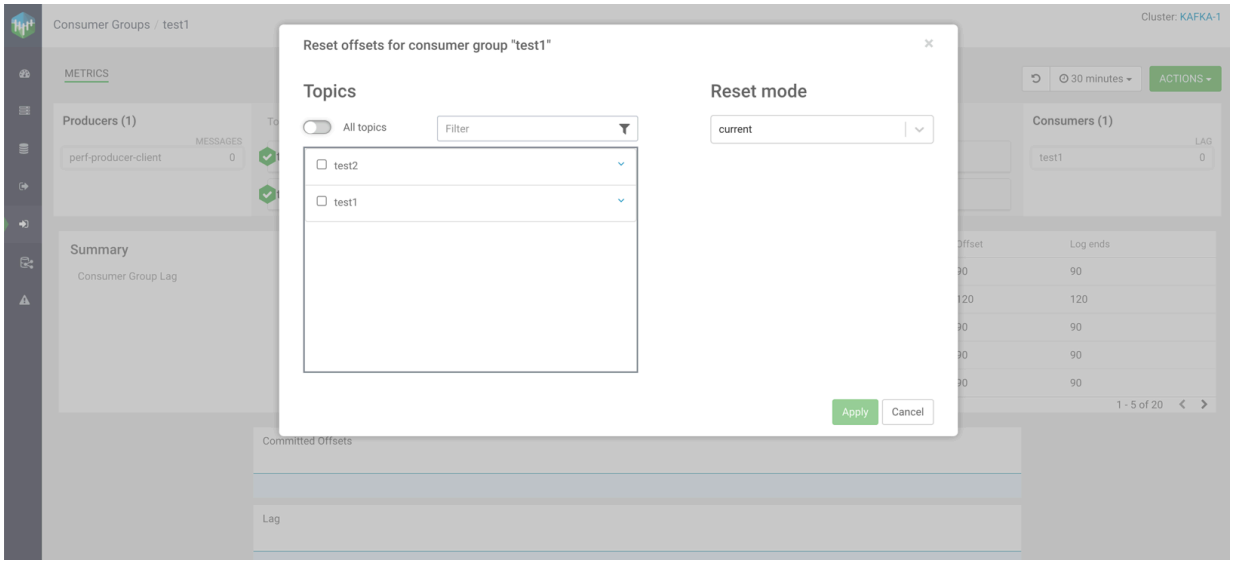

4. Select the topic and partition(s) you want to reset.

   You can use the arrow beside each topic to display the partitions of that topic, and select partition(s) as required. You can also select all topics and all related partitions by selecting the All topics option. You can use the Filter option to find a specific topic.

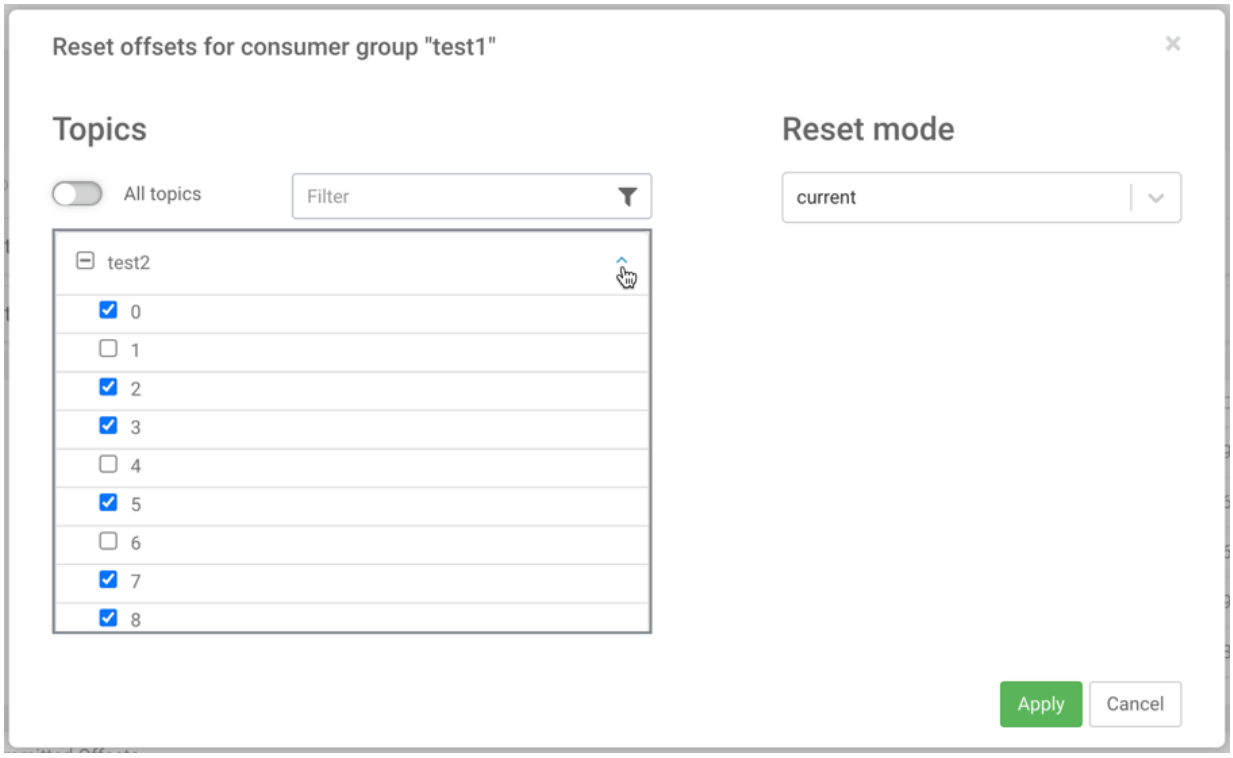

5. Select the reset option from the Reset mode.

   The available options are as shown in the following image:

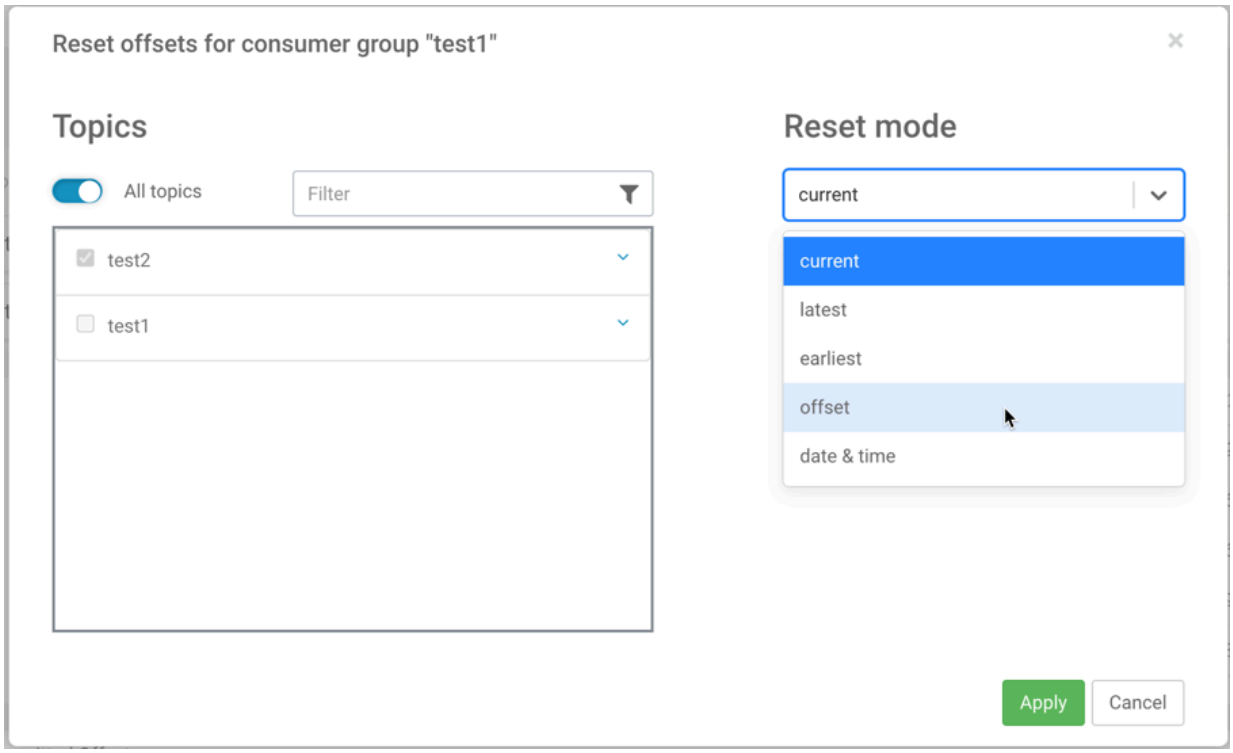

6. Click Apply.

Monitoring log size information

The SMM UI shows log size related information about brokers, topics, and partitions. Furthermore, warning messages appear when log directory related errors happen.

The feature uses a cache mechanism. The following cache attributes are set by default. You can modify these attributes in the SMM configuration page of the Cloudera Manager UI:

- Kafka Logdir Cache Expiration ms

  Sets cache expiration time.
- Kafka Logdir Cache Maximum Size

  Sets maximum cache size.
- Kafka Logdir Call Timeout ms

  Sets timeout of the Kafka query.

The following image shows the properties in Cloudera Manager:

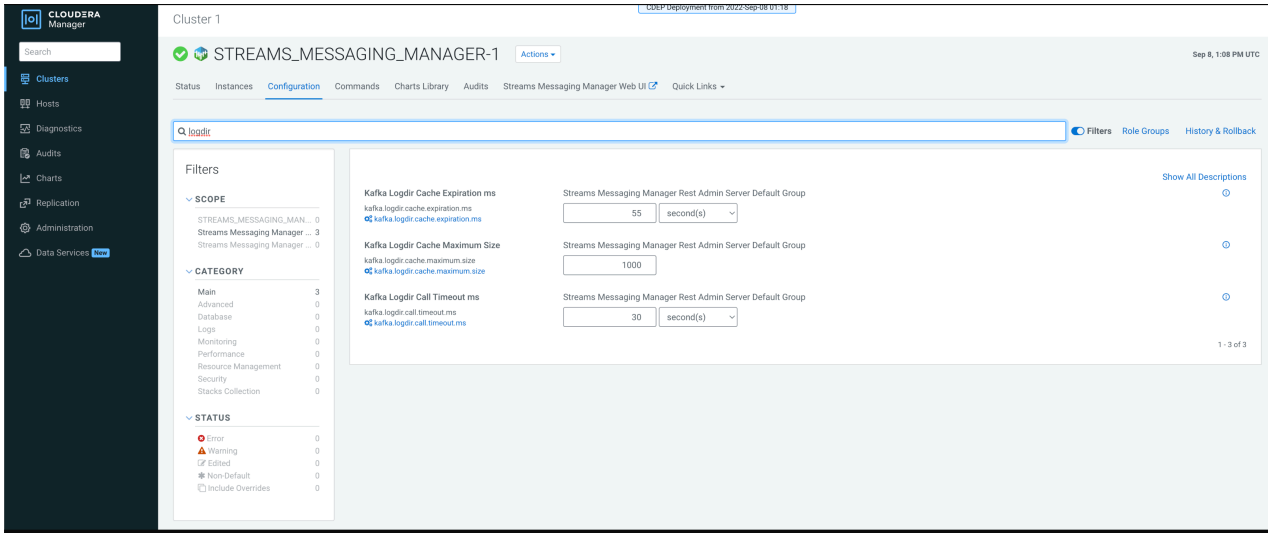

Log size details for topics

On the Overview page of the SMM UI, the Topics tab contains a sortable column called Current Log Size which shows the data with the measure unit. This log size is calculated by the leader partitions; so, the follower partitions are not included. The column contains a tooltip which is a question-mark symbol. The topic related log-size improvements or information are also available on the Topics page.

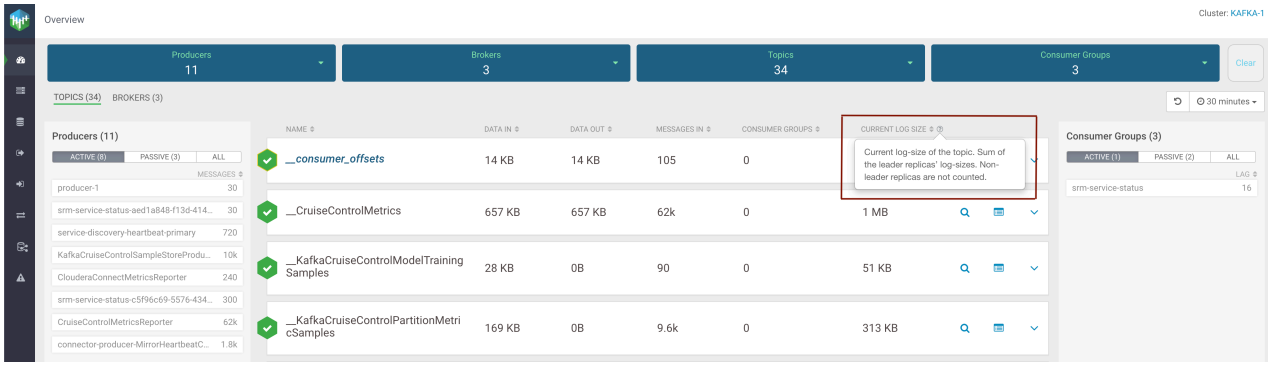

The Topics tab on the Overview page represents topics and you can select one of them. In this case the topic's leader partitions are viewable and here you can find the leader partition log-size. If you click a partition, a pop-up window appears, where log-size information is displayed.

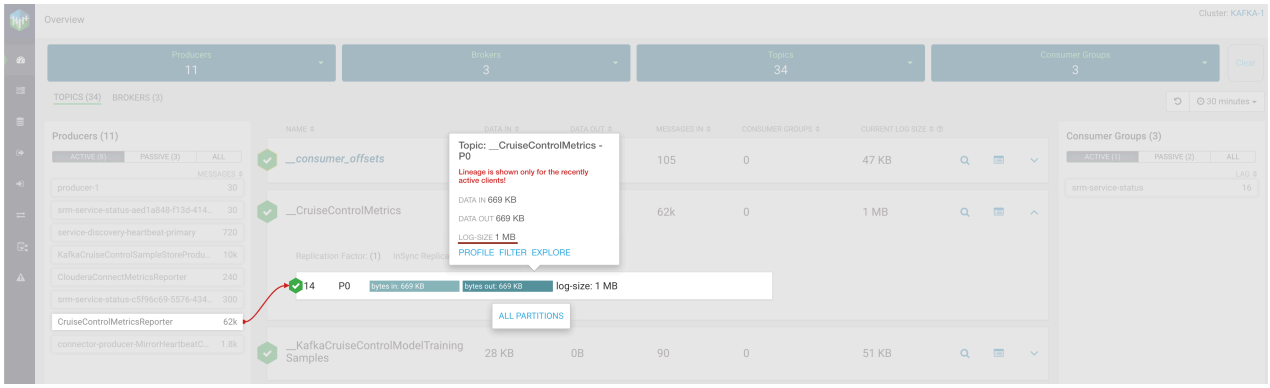

If you click Profile for a topic, log-size of replica-leaders appears, hosted by the broker.

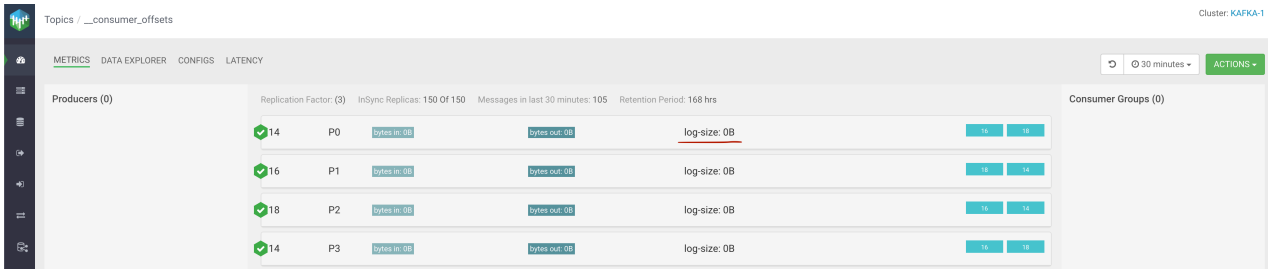

Log size details for brokers

On the Overview page, the Brokers tab contains two sortable columns called Log Size and Remaining Storage. The log-size is calculated for all the partitions including the follower partitions. The remaining storage column contains the available capacity in all of the configured log directories. The broker related log-size improvements or information are also available on the Brokers page.

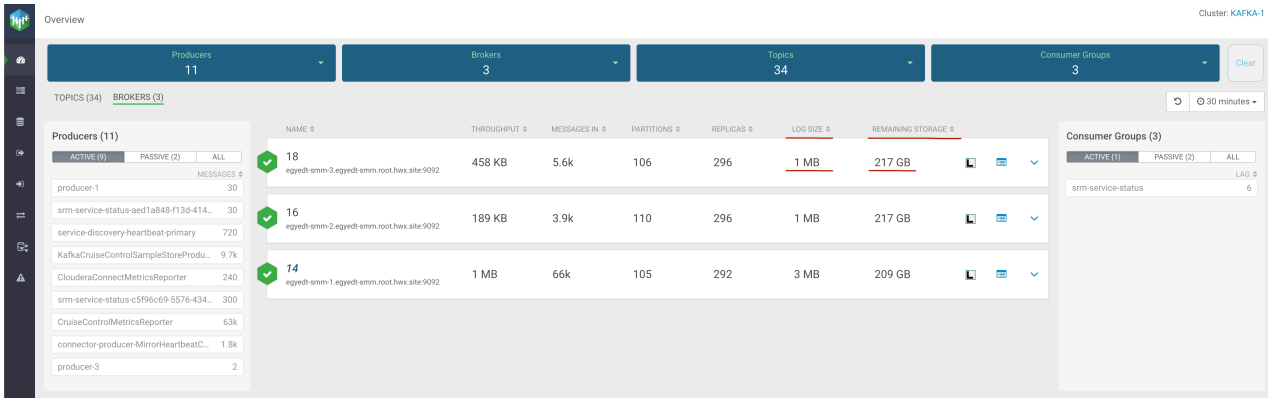

If you click the arrow icon for a broker, the details of that broker are displayed. The Leader Partition Log-Size shows the sum of leader partitions' log-size. The All Partitions Log-Size field's value is equal to the sortable column's value called Log Size (naming is different to make the sortable column names as short as possible). The Remaining Storage sortable column and field is the same (value and name).

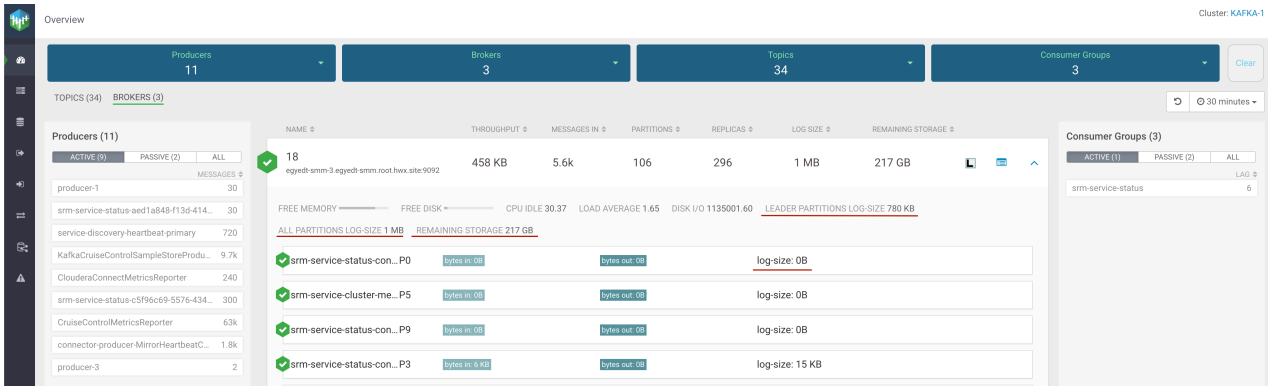

If you click a specific partition, a pop-up window appears where log-size information is displayed.

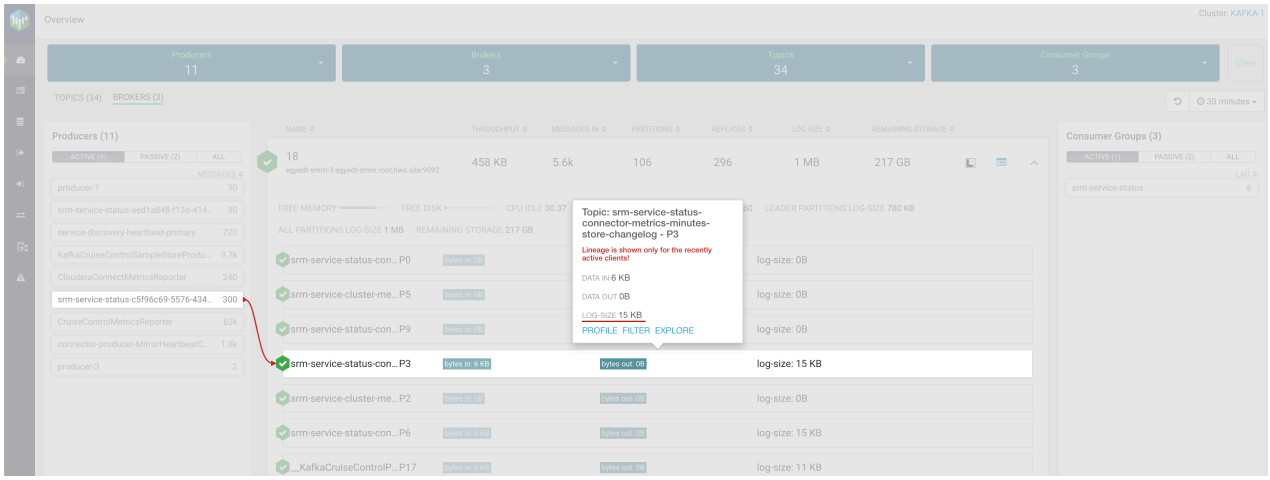

If you click Profile of a broker, then the metrics details for that broker are displayed. The log-size information per partition is available. Furthermore, the sum of leader partitions' log-size, remaining storage size, and all partitions log-size details are also available. The Metrics tab on the Brokers page contains the available storage size per log directory.

Note: There is no advantage of creating more log directories on the same disk, so this should be avoided.

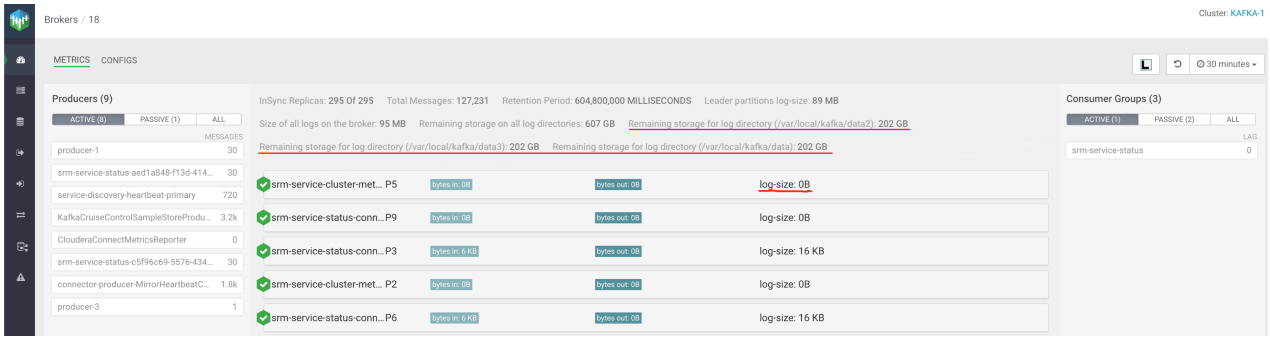

If there is any issue with the query on the log directory information, then a warning message appears. There can be multiple warnings if more than one broker is related. But no new warning message from a specific broker appears until the previous message disappears.

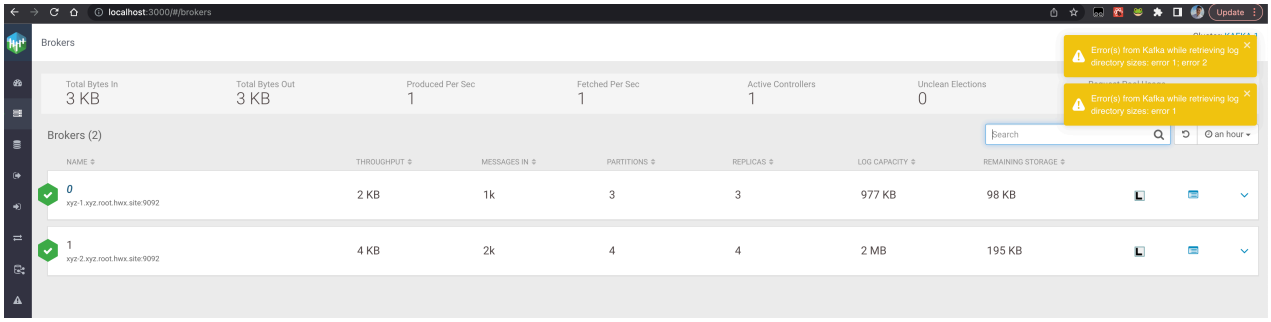

Cloudera Manager chart builder

For historical information, you can use the chart builder provided in Cloudera Manager. The following command and screenshot show Kafka log-size data from the past where the broker ID and topic name with the topic partition are specified.

```
SELECT counter_sum(kafka_bytes_received_by_partition_rate) WHERE serviceName
=KafkaServiceName AND kafkaBrokerId = KafkaBrokerId AND kafkaTopicName = Kaf
kaTopicName AND kafkaPartitionId = kafkapartitionId
```

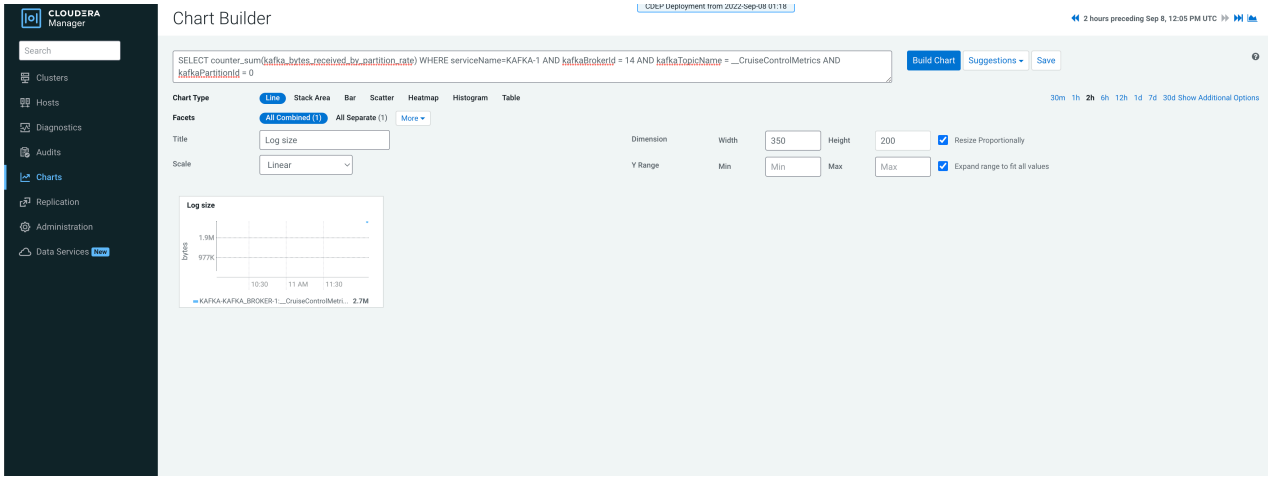

Monitoring lineage information

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:

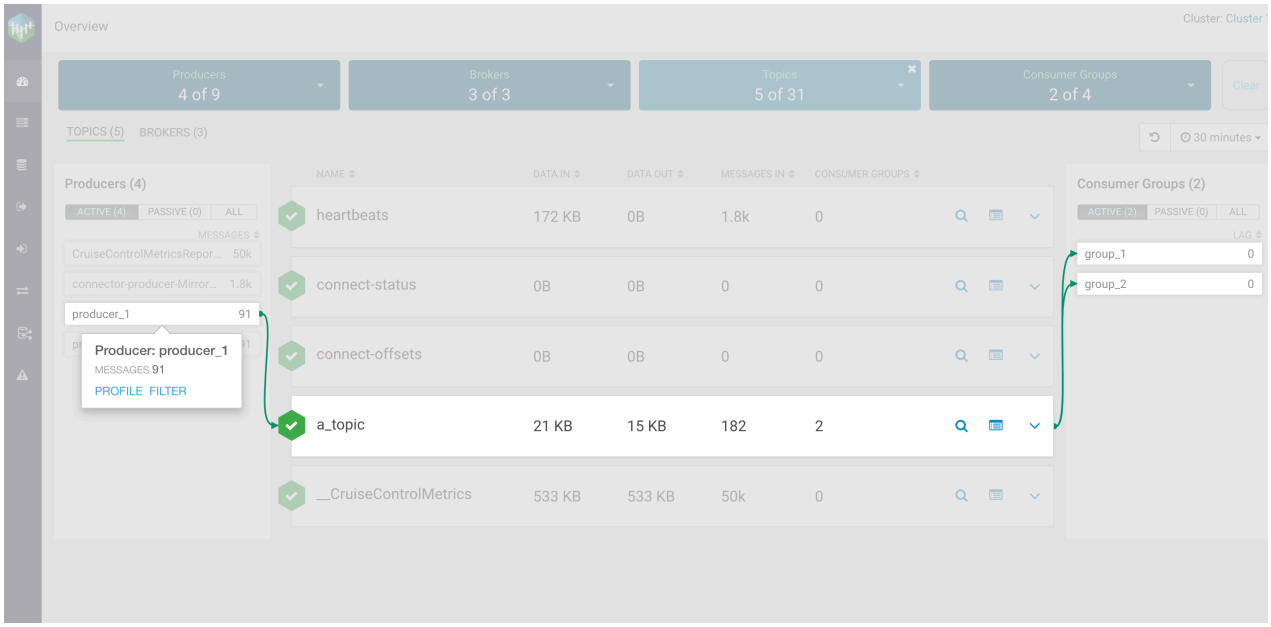

After you click producer_1, you can see that it produces to a topic called a_topic, and that both consumer groups (group_1 and group_2) consume from that topic.

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:

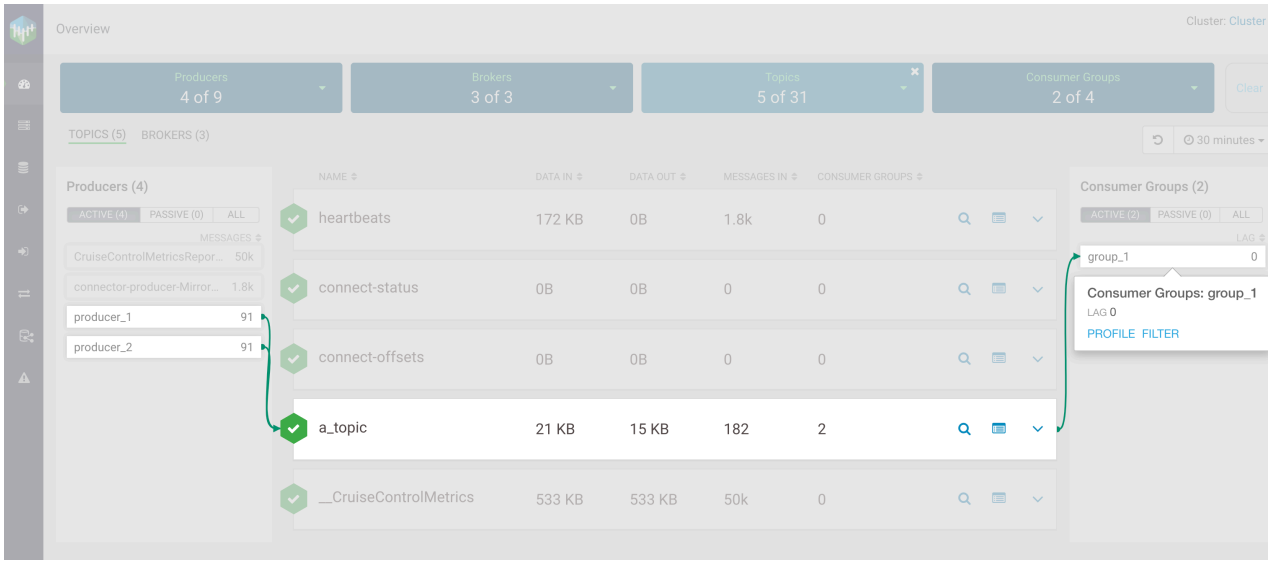

After you click group_1 consumer group, you can see that it consumes from the topic called a_topic, and that two producers produce to that topic (producer_1 and producer_2).

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:

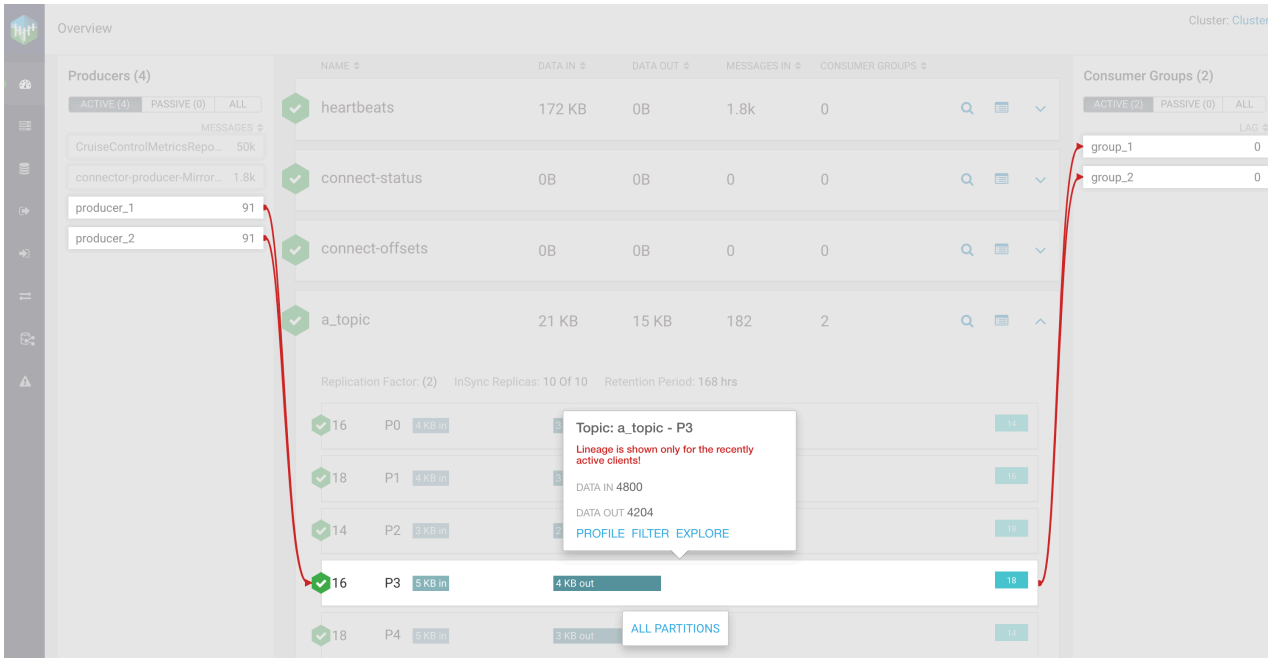

After you click P3 partition in the topic called a_topic, you can see that producer_1 and producer_2 produce to that partition, and group_1 and group_2 consume from it.

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

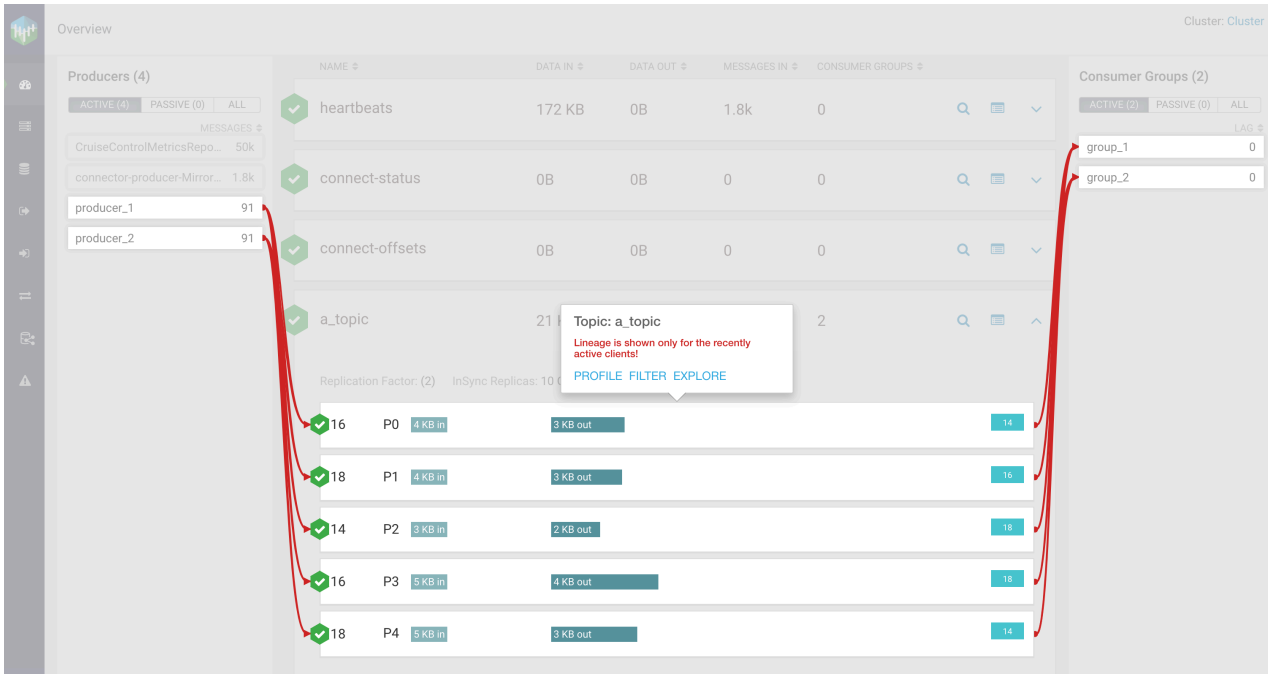

You can also access the lineage information from the experimental endpoints. You can find the endpoints at the *Streams Messaging Manager REST API Reference*.

Related Information

Streams Messaging Manager REST API Reference