

Monitoring Kafka Clusters Using Streams Messaging Manager

Date published: 2019-05-31

Date modified: 2019-05-31



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

- Monitoring Clusters.....4**
- Monitoring Producers..... 5**
- Monitoring Topics.....8**
- Monitoring Brokers..... 10**
- Monitoring Consumers.....12**

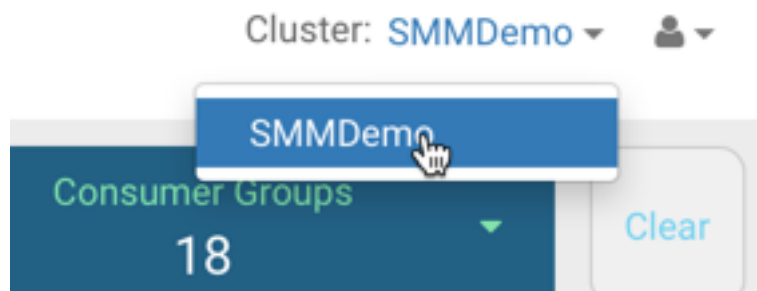
Monitoring Clusters

The Overview page provides you with tools to see a snapshot of the Kafka cluster you are monitoring.

Select the Cluster to Monitor

You can use SMM to monitor all the Kafka clusters in your environment. You can view information about Kafka cluster health and activity by selecting the cluster about which you want information.

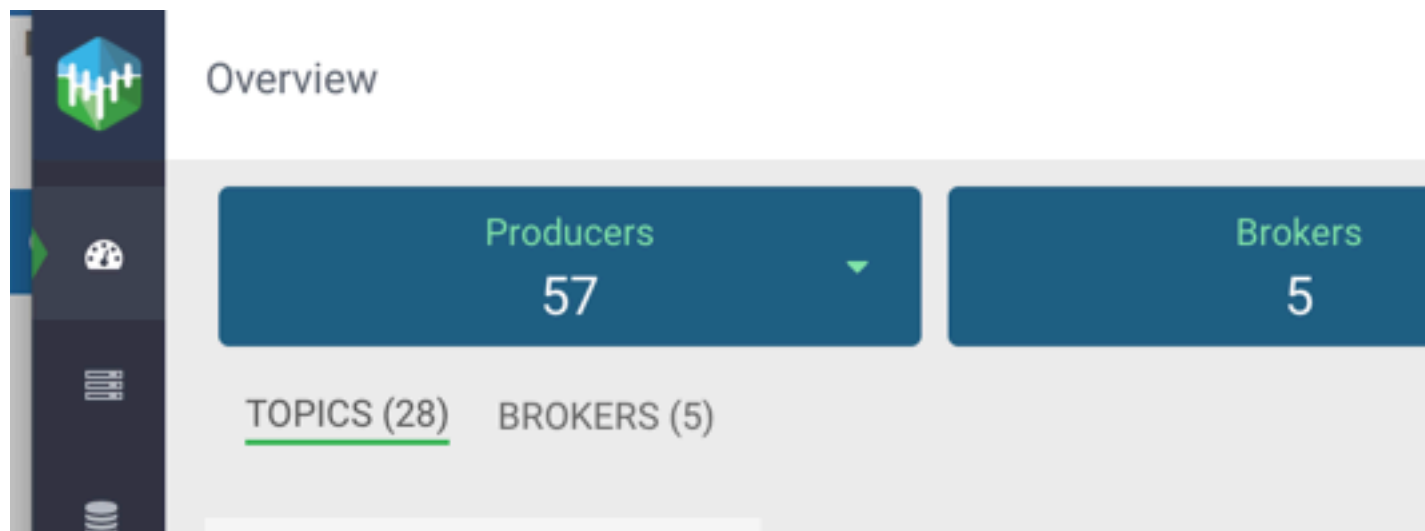
1. From the left navigation pane, select Overview.
2. Select your cluster from the Cluster drop-down list.



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



Monitoring Producers

Understanding Producer Naming Conventions

The Producers you interact with in 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 consuming 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 Ambari Streams Messaging Manager Configs screen.

1. Select Streams Messaging Manager from the Ambari 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.

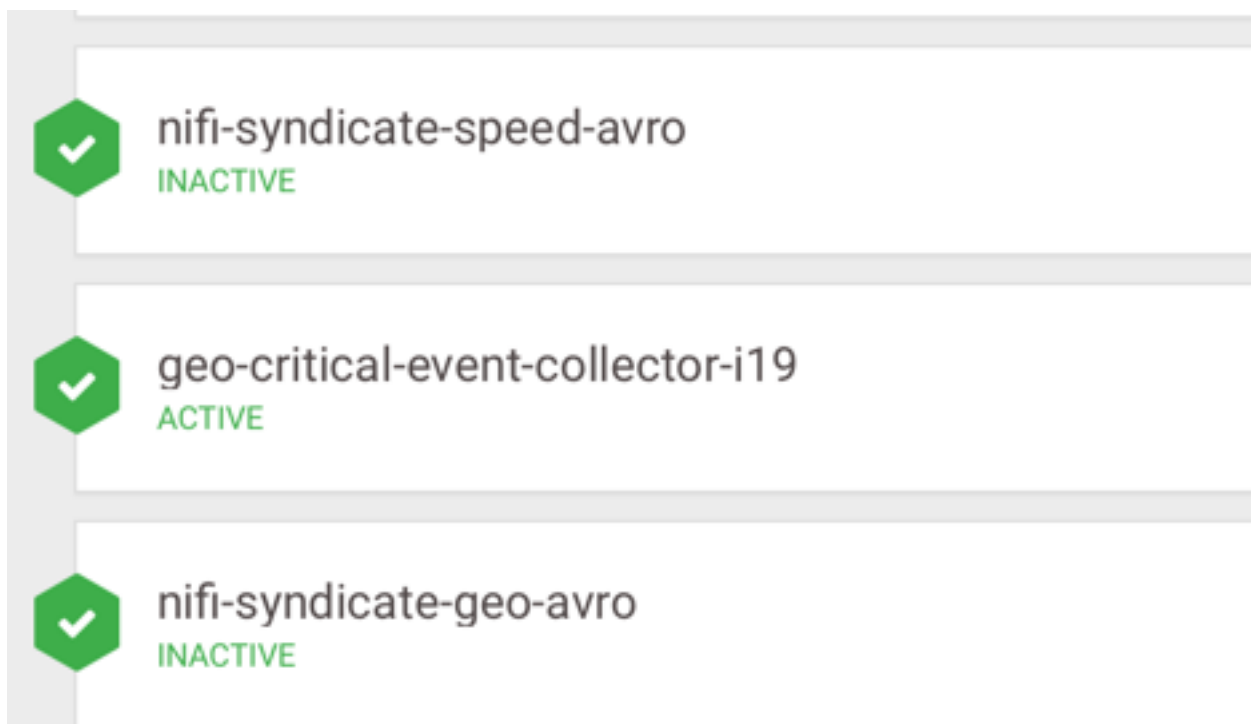
From the Overview page Producer pane, use the Active, Passive, and All tabs to view only active producers, only passive producers, or all of them. This allows you to see the total number of active and passive Producers.



The screenshot shows a panel titled "Producers (84)" with three tabs: "ACTIVE (57)", "PASSIVE (27)", and "ALL". Below the tabs is a "MESSAGES" header. The list of producers and their message counts is as follows:

Producer Name	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.



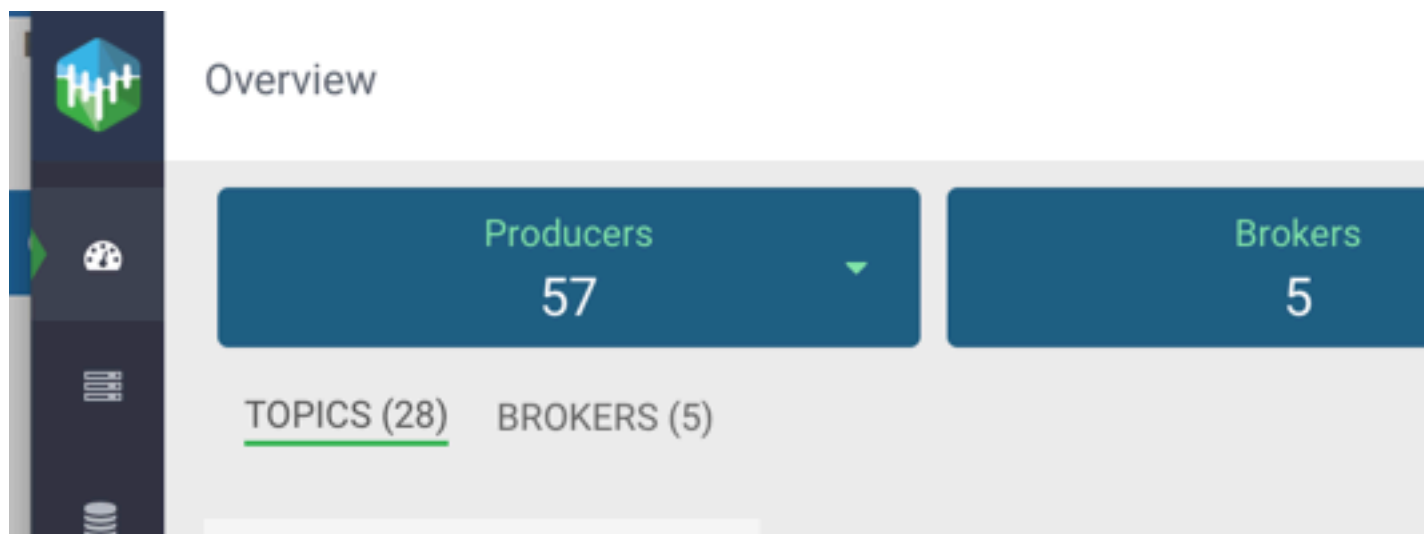
The screenshot shows a list of three producers with their status indicated by a green checkmark icon and the word "INACTIVE" or "ACTIVE" below the producer name:

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

Monitoring Topics

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.



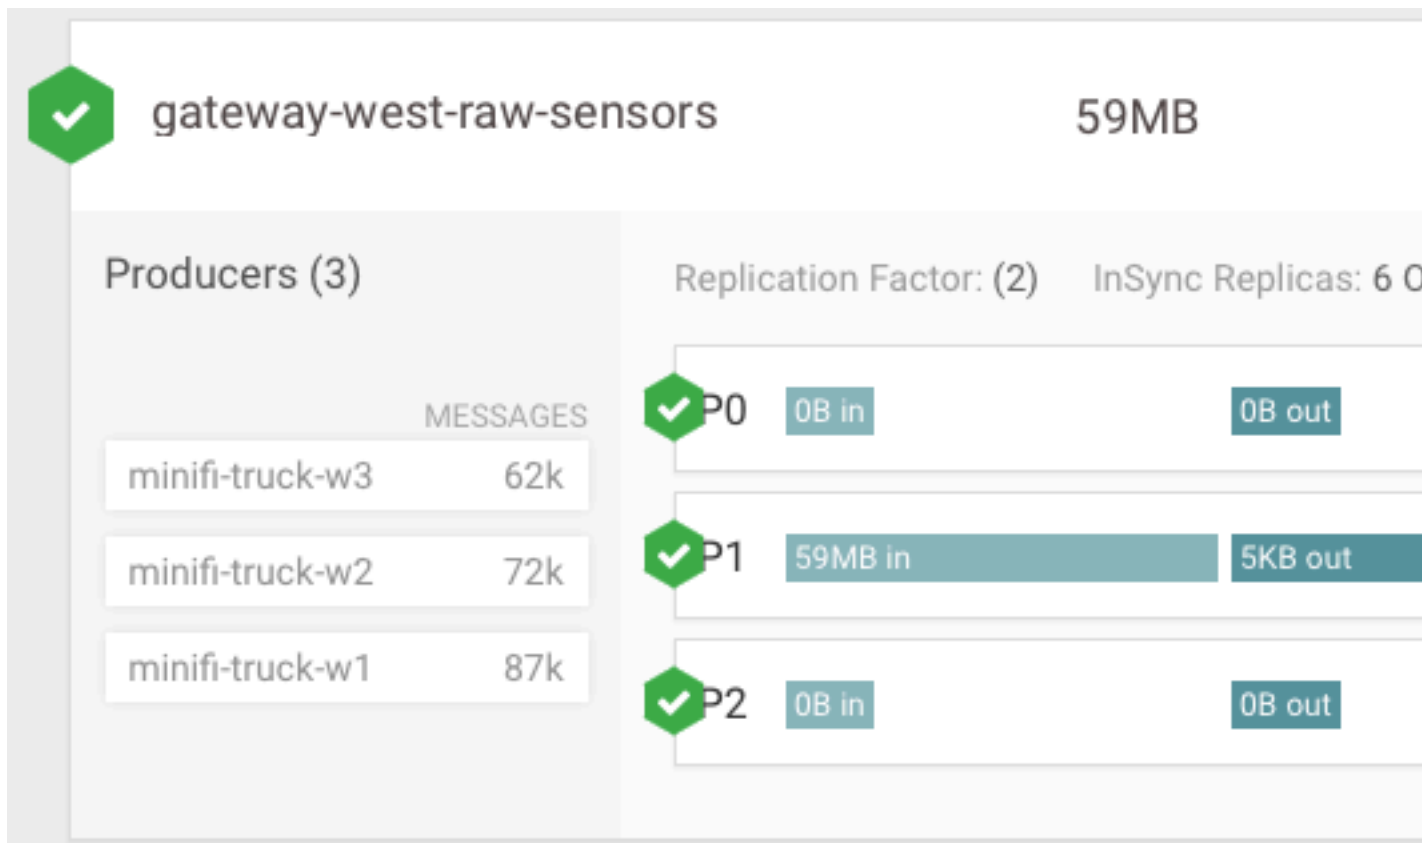
Detailed Information about Topics

The Topic 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:

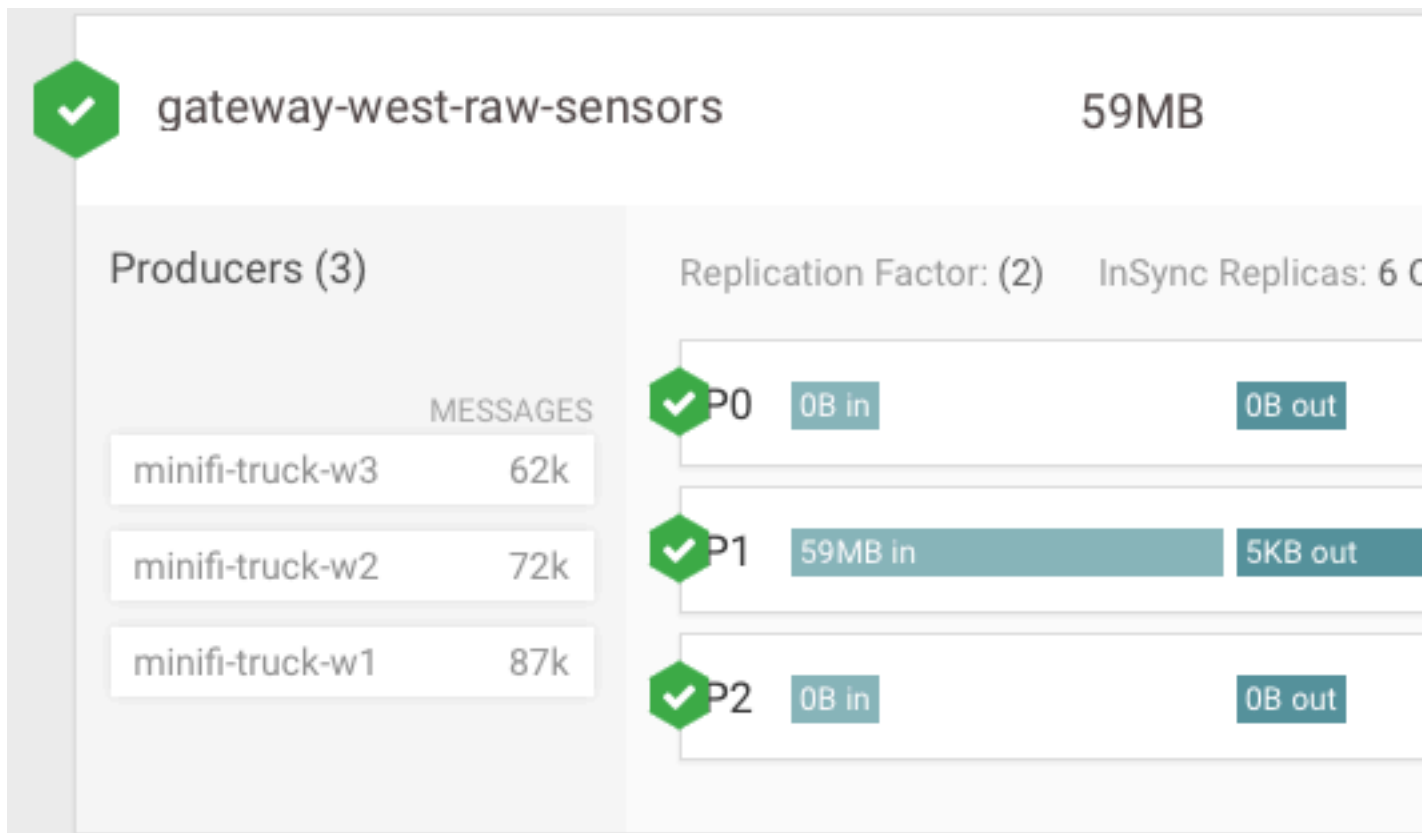
1. From the left navigation pane, click Topics.
2. Identify the Topic about which you want information. You can either scroll through the list of Topics, or use the Search bar at the top left of the page.
3. Click the green hexagon at the left of the Topic to view details.



Viewing Grafana Metrics for Your Topic

If Grafana is installed as part of your Ambari-managed target cluster, you can access detailed Grafana Metrics for Topics in your Kafka cluster.

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 Grafana icon in the top right of the Topic view.



Monitoring Brokers

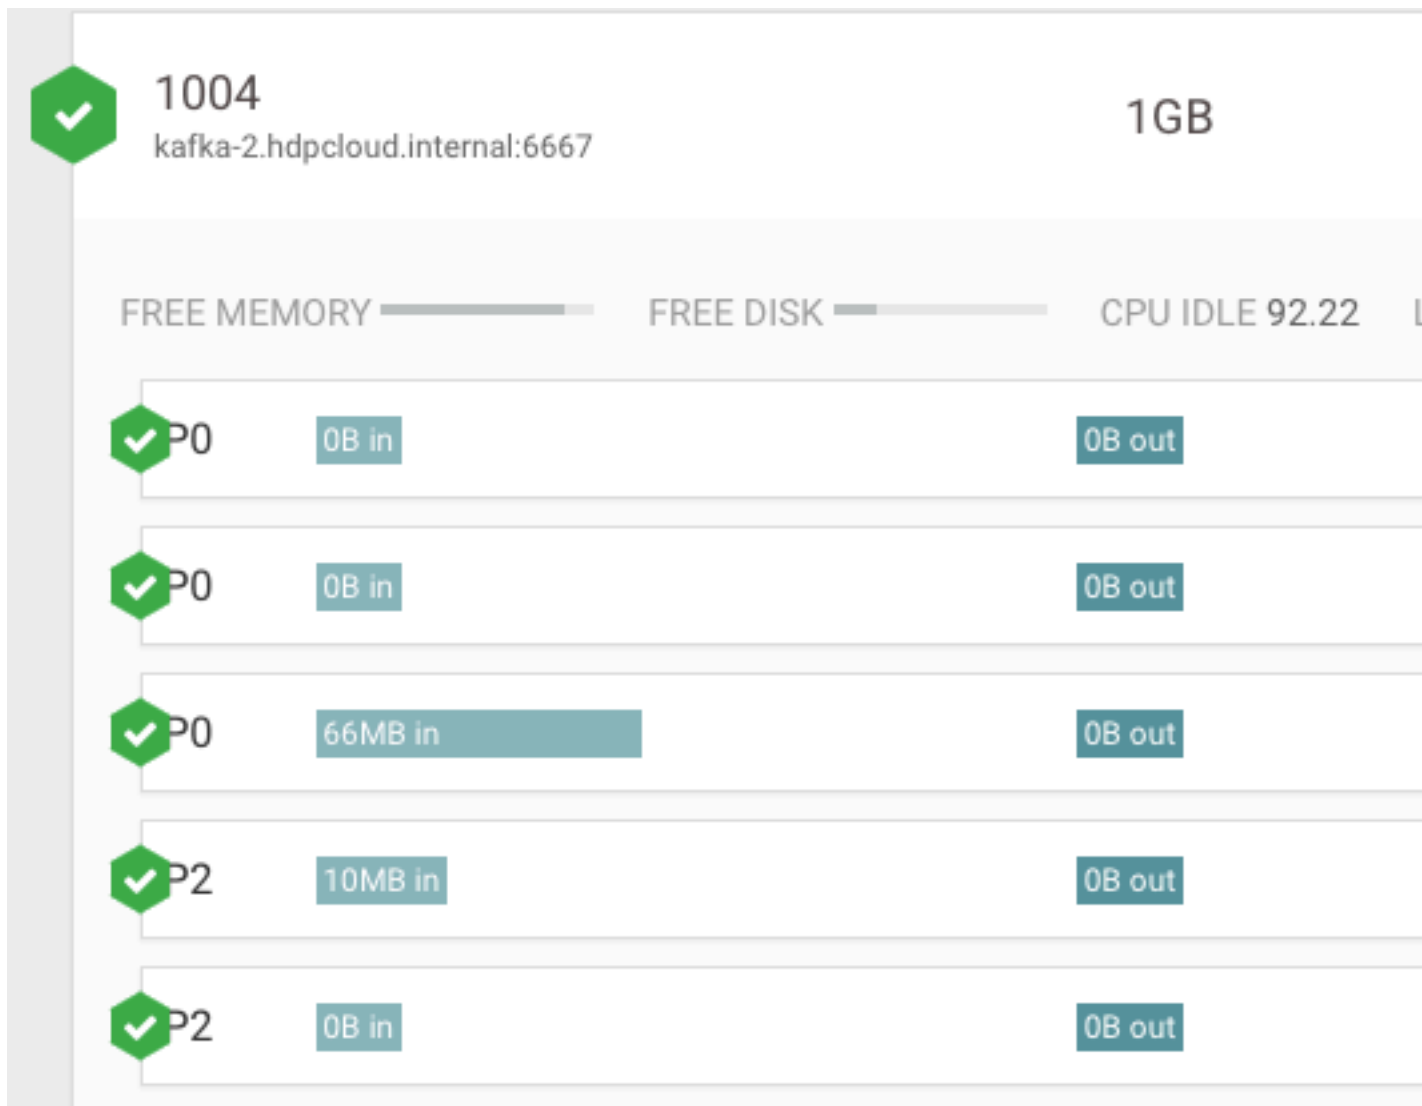
Detailed Broker Information

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.

Brokers (5)

	NAME	THROUGHPUT	MESSAGES I
	1001 kafka-5.hdpcloud.internal:6667	3GB	15m

Monitoring Consumers

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. Use the Lag tab to sort consumer groups based on ascending or descending amounts of lag.

Overview

Producers
84

Brokers
5

TOPICS (28) BROKERS (5)

Producers (84)

ACTIVE (61) PASSIVE (23) ALL

MESSAGES

minifi-eu-i1 8.2m

geo-critical-event-coll... 4.1m

geo-critical-event-coll... 4m

fuel-apps 3.4m

supply-chain-apps 2.3m

geo-critical-event-coll... 1.8m

geo-critical-event-coll... 1.6m

NAME

DATE



syndicate-transmission

13



syndicate-speed-even...

0E



syndicate-speed-even...

0E



syndicate-oil

16

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.



The screenshot displays the profile for the consumer group 'nifi-truck-sensors-west', which is in an 'ACTIVE' state. It shows three partitions: P0, P1, and P2. Partition P1 is currently processing 81MB of data, while P0 and P2 are at 0B.

Partition ID	Consumer Name	Partition	Current State
1004	gateway-west-r...	P0	0B in
1005	gateway-west-r...	P1	81MB in
1001	gateway-west-r...	P2	0B in

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.

METRICS

Producers (3)

Producer	MESSAGES
minifi-truck-w3	90k
minifi-truck-w2	0.1m
minifi-truck-w1	0.1m

Partitions (3) State: Stable

✓ 1004	gawewa... P0	0B in
✓ 1005	gawewa... P1	81MB in
✓ 1001	gawewa... P2	0B in

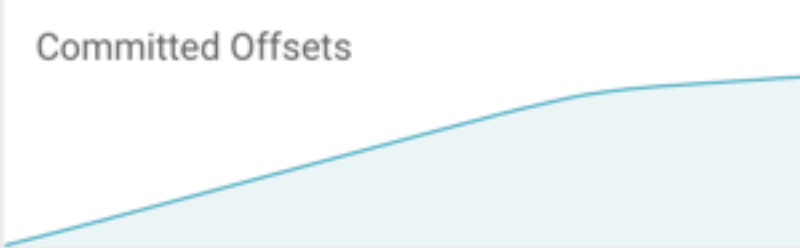
Summary

Total Lag	8
-----------	---

Instance

- consumer-1-29fa6e77-78dd-4...
- consumer-1-6b0e4ffc-6f4c-4e...
- consumer-3-68359916-35f8-4...

Committed Offsets



Committed Offsets Rate

