

Monitoring Kafka Cluster Replications using Streams Messaging Manager 7.2.6

Monitoring Kafka Cluster Replications Using Streams Messaging Manager

Date published: 2019-09-20

Date modified: 2020-12-11

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 the middle of "UDERA" having a unique design where the top and bottom bars are separated by a horizontal gap.

<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

Introduction to monitoring Kafka cluster replications in SMM.....	4
Configuring SMM for monitoring Kafka cluster replications.....	4
Viewing Kafka cluster replication details.....	5
Searching Kafka cluster replications by source.....	6
Monitoring Kafka cluster replications by quick ranges.....	6
Monitoring status of the clusters to be replicated.....	7
Monitoring topics to be replicated.....	7
Searching by topic name.....	8
Monitoring throughput for cluster replication.....	9
Monitoring replication latency for cluster replication.....	10
Monitoring checkpoint latency for cluster replication.....	11
Monitoring replication throughput and latency by values.....	11

Introduction to monitoring Kafka cluster replications in SMM

You can use Streams Replication Manager (SRM) to implement cross-cluster Kafka topic replication in Streams Messaging Manager (SMM). You need to configure SRM properties in SMM. After you configure SRM in SMM, you can use SMM to monitor all of the Kafka cluster replications writing to the target cluster of the configured SRM Service role.

SRM is an enterprise-grade replication solution that enables fault tolerant, scalable and robust cross-cluster Kafka topic replication. SRM provides the ability to dynamically change configurations and keeps the topic properties in sync across clusters at high performance. SRM also delivers custom extensions that facilitate installation, management and monitoring making SRM a complete replication solution that is built for mission critical workloads. For more information about SRM, see the *Streams Replication Manager for HDF and HDP* document library.

You can view all replication details in SMM. You can monitor status of the Kafka cluster replications, number of topics associated with the replication, throughput, replication latency, and checkpoint latency of Kafka cluster replications in SMM.



Note: You can also create alerts to receive notifications based on the conditions that you configure in the alert policy for monitoring Kafka cluster replications in your system. For more information about creating alerts to monitor Kafka cluster replications, see the *Managing Alert Policies* guide.

Related Information

[Alert Policies Overview](#)

Configuring SMM for monitoring Kafka cluster replications

To use Streams Messaging Manager (SMM) to monitor Kafka cluster replications, you must configure SMM to communicate with Streams Replication Manager (SRM).

About this task

For information on how to install and configure SRM, see *Creating your first Streams Messaging cluster*.

Before you begin


In order for SMM to be able to monitor SRM replications correctly, you must configure the SRM driver to target all clusters that are taking part in the replication. For more information, see *Configuring the driver role target clusters*.

Procedure

1. In Cloudera Manager, select Streams Messaging Manager from the available Clusters.
2. Click Configuration.
3. Enable the STREAMS_REPLICATION_MANAGER Service property.
4. Click Save Changes (CTRL+S).
5. Restart SMM.

Results

You can see the Cluster Replications option in the SMM UI as shown in the following image:



STATUS	SOURCE	TARGET	TOPICS	CONSUMER GROUPS	THROUGHPUT	REPLICATION LATENCY	CHECKPOINT LATENCY
ACTIVE	primary	secondary	7	0	176 B/s	1.4 sec	N/A

Related Information

[Creating your first Streams Messaging cluster](#)

[Configuring the driver role target clusters](#)

Viewing Kafka cluster replication details

After you select a Kafka cluster, Streams Messaging Manager (SMM) displays the details of the cluster replications where the selected cluster is the destination cluster. Learn what details you can view for a Kafka cluster replication.


About this task

You can view the status of the replications, source cluster names, target cluster names, number of topics to be replicated, number of consumer groups, throughput of the replication, replication latency, and checkpoint latency. SMM also displays two graphs, one represents the throughput of the replication and the other displays the replication latency of the replication along with the details of each topic to be replicated.

Perform the following steps to view the details of a cluster replication:

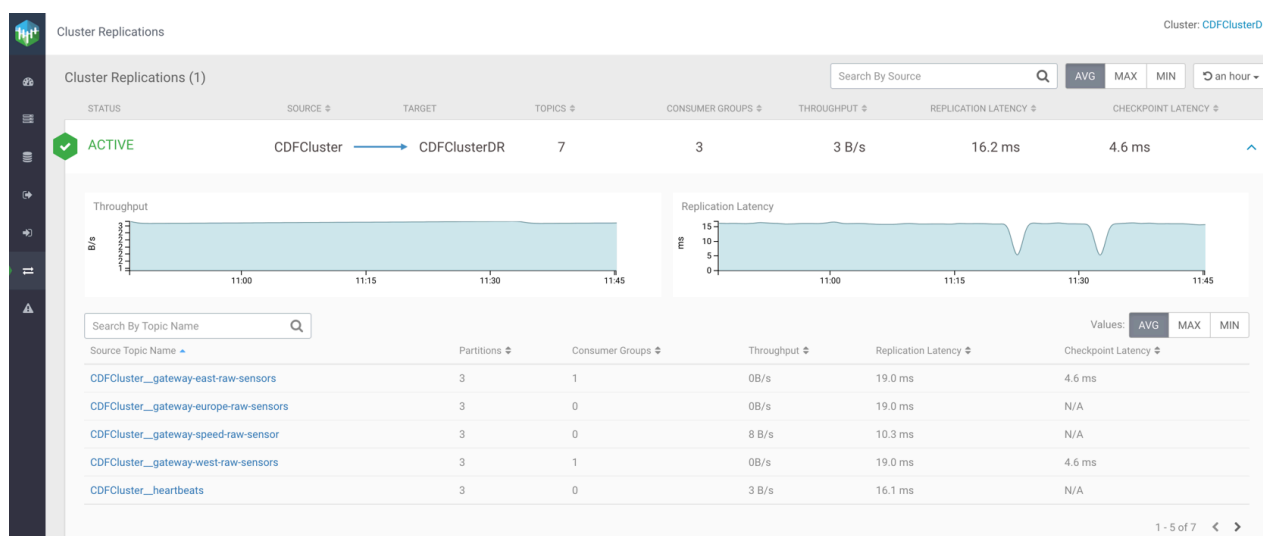
Procedure

In the Cluster Replications page, click a cluster replication or the drop-down icon beside a replication, as shown in the following image:



STATUS	SOURCE	TARGET	TOPICS	CONSUMER GROUPS	THROUGHPUT	REPLICATION LATENCY	CHECKPOINT LATENCY
ACTIVE	CDFCluster	CDFClusterDR	7	3	3 B/s	16.2 ms	4.6 ms

The replication details appear as shown in the following image:

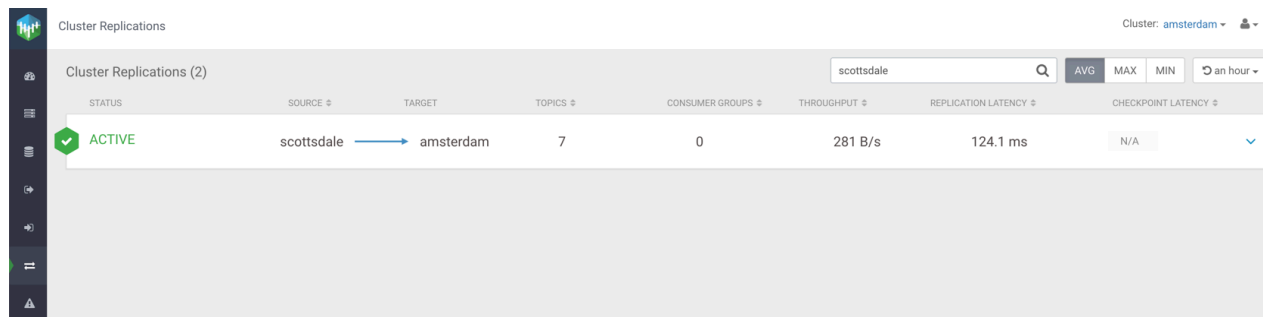


Searching Kafka cluster replications by source

Learn how to filter your Kafka cluster replications by source cluster after you select the Kafka cluster which acts as the target cluster.

About this task

You can use the Search By Source bar at the top right of the page to search cluster replications by the source cluster name. For example, if the destination cluster is amsterdam, and you enter scottsdale in the Search By Source bar, SMM displays the Kafka cluster replication between scottsdale and amsterdam, as shown in the following image:



Monitoring Kafka cluster replications by quick ranges

Learn how to filter your cluster replications by specific time periods after you select the Kafka cluster which acts as the target cluster.

About this task

You can filter the cluster replications by time range. You can choose any of the following filter values from the drop-down to display the cluster replication details in the selected time range:

- Last 1 Hour
- Last 6 Hours
- Last 24 Hours
- Last 2 Days

The following image shows the Quick Ranges drop-down:

STATUS	SOURCE	TARGET	TOPICS	CONSUMER GROUPS	THROUGHPUT	REPLICATION LATENCY	CHECKPOINT
ACTIVE	springfield	amsterdam	9	0	745 B/s	8.0 ms	N/A
ACTIVE	scottsdale	amsterdam	7	0	281 B/s	125.3 ms	N/A

Monitoring status of the clusters to be replicated

You can monitor the status of Kafka cluster replications in the Status column in the Cluster Replications page. You can see whether a cluster replication is active, inactive, or showing a warning.

The status has three variations:

- Active
Indicates that the cluster replication is running.
- Inactive
Indicates that the cluster replication is not running.
- Warning
Indicates that the cluster replication is facing issues.

If the status of a Kafka cluster replication shows inactive or warning, check the logs, and troubleshoot the replication.

Monitoring topics to be replicated

You can monitor the number of topics associated with a Kafka cluster replication in the Topics column in the Cluster Replications page. You can also monitor details on source topic name, partitions, consumer groups, throughput, replication latency, and checkpoint latency.

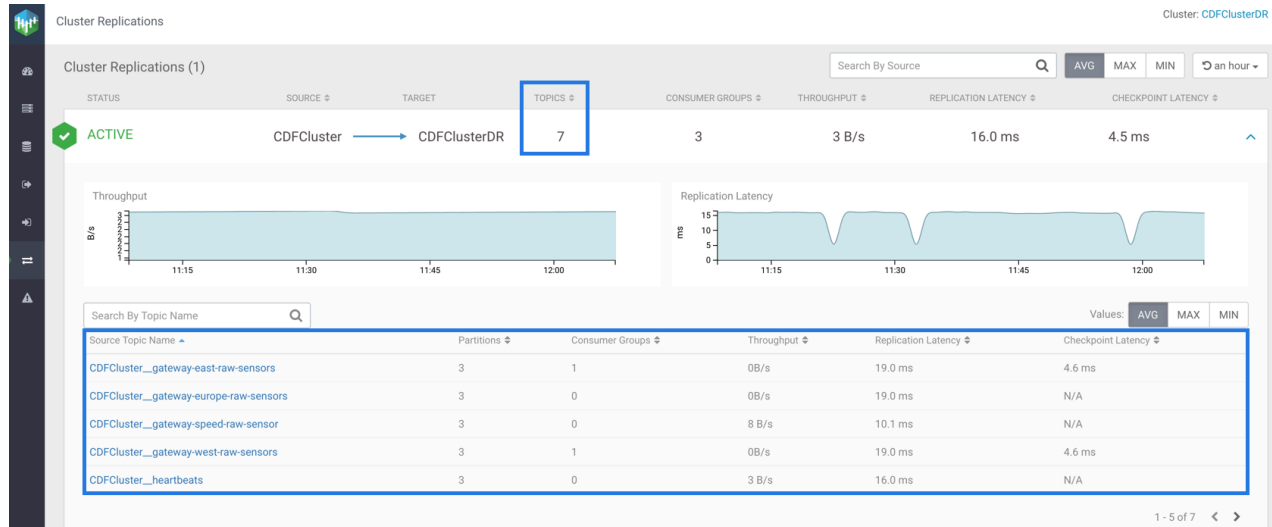
Click on the cluster replication to fetch topic details. SMM displays the following details about the topics:

- Source Topic Name
Name of the topic at the source.
- Partitions
Number of partitions of the topic at the source.
- Consumer Groups
Number of consumer groups consuming data from the topic.
- Throughput
Data replicated between the source cluster and the target cluster per second from a topic. Throughput is measured in bytes per second. By default, SMM displays the average throughput. You can also fetch the maximum or minimum throughputs for a topic by clicking the MAX or MIN button located above the topic details.
- Replication Latency
Amount of time taken for a message of a topic to get replicated from the source cluster to the target cluster. Replication latency is measured in milliseconds. By default, SMM displays the average replication latency. You can also fetch the maximum or minimum replication latency for a topic by clicking the MAX or MIN button located above the topic details.

- Checkpoint Latency

Amount of time taken for a message of a topic to checkpoint on the target cluster after the message is committed on the source cluster. Checkpoint latency is measured in milliseconds. By default, SMM displays the average checkpoint latency. You can also fetch the maximum or minimum checkpoint latency for a topic by clicking the MAX or MIN button located above the topic details.

The following diagram shows details of topics in a cluster replication:

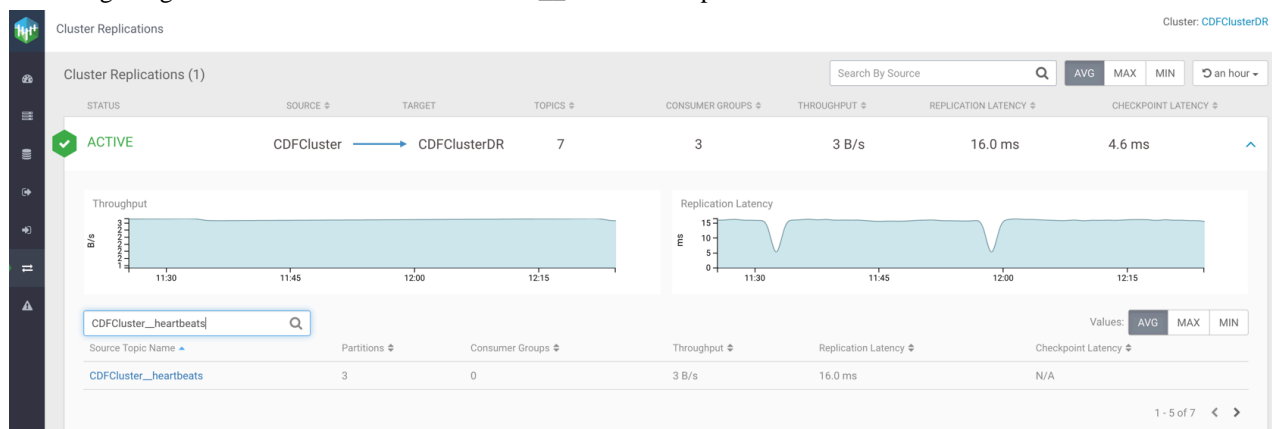


In the image, you can see that the number of topics to be replicated from CDFCluster to CDFClusterDR is 7, and topic details including topic names, number of partitions for the topics at the source cluster, number of consumer groups consuming messages from each topic, throughput, replication latency, and checkpoint latency of each topic.

Searching by topic name

Learn how to search a Kafka topic by name in a Kafka cluster replication and fetch details for that topic. After you find a topic, you can see the topic name, partitions, consumer groups, throughput, replication latency, and checkpoint latency for that topic.

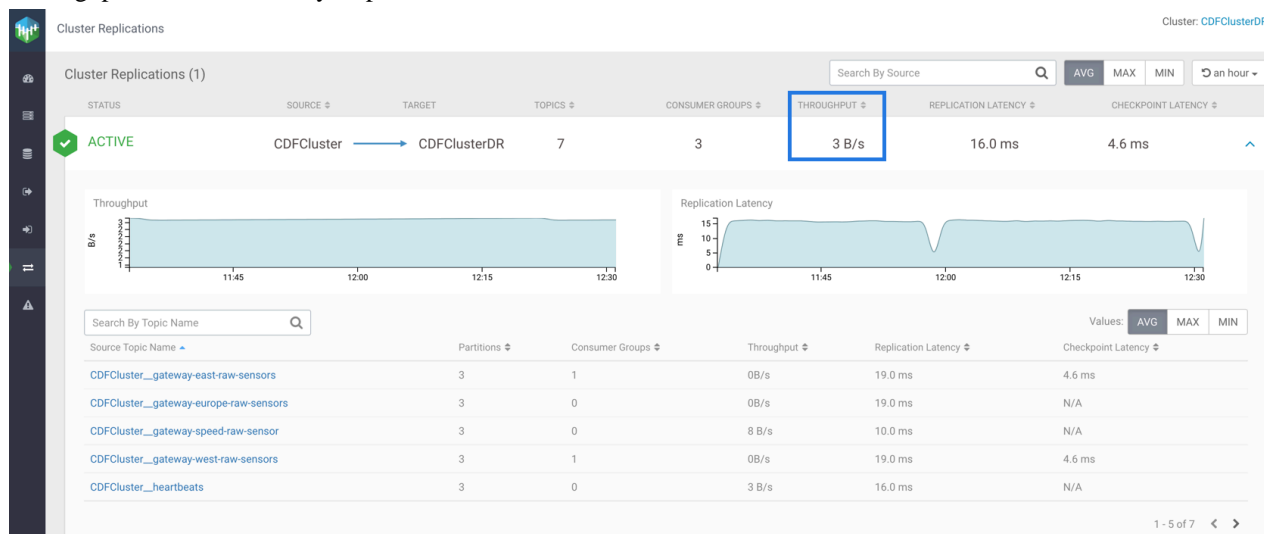
You can use the Search By Topic Name bar to search a Kafka topic by name and get details of that topic. The following image shows details of the CDFCluster__heartbeats topic:



Monitoring throughput for cluster replication

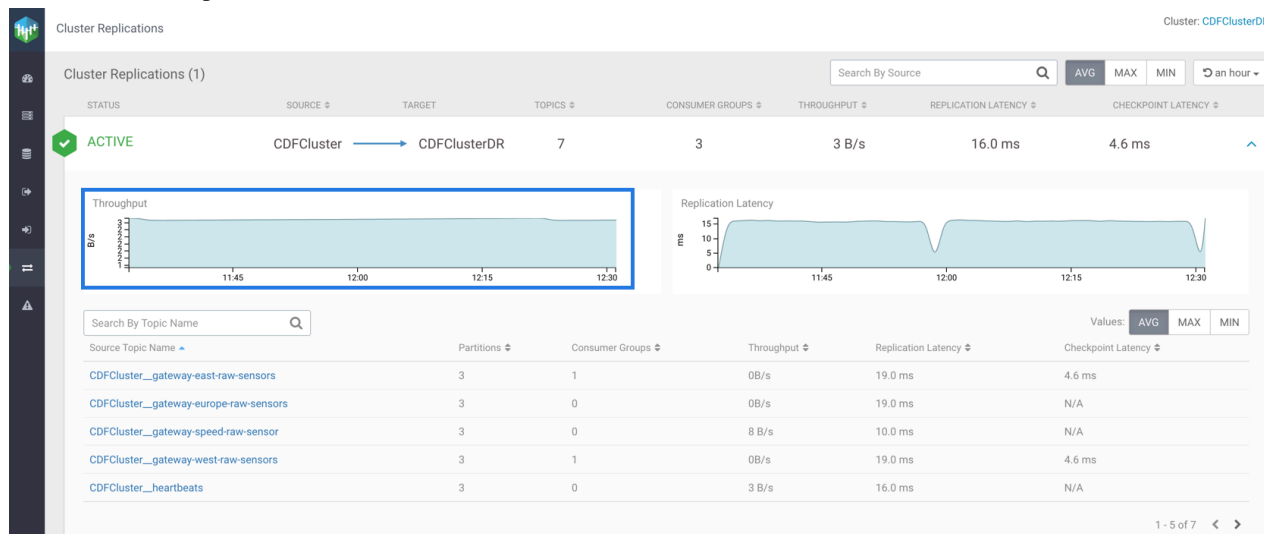
Learn how to monitor the throughput of a Kafka cluster replication in Streams Messaging Manager (SMM). You can monitor average, maximum, and minimum throughput of a cluster replication. You can also monitor the throughput of a cluster replication graphically.

Throughput is defined as the data replicated between the source cluster and the destination cluster per second. Throughput is measured in bytes per second.



In the image, you can see the average throughput for the CDFCluster to CDFClusterDR replication is 3 bytes per second. You can fetch the maximum or minimum throughput for the replication by clicking the MAX or MIN button located above the cluster replication.

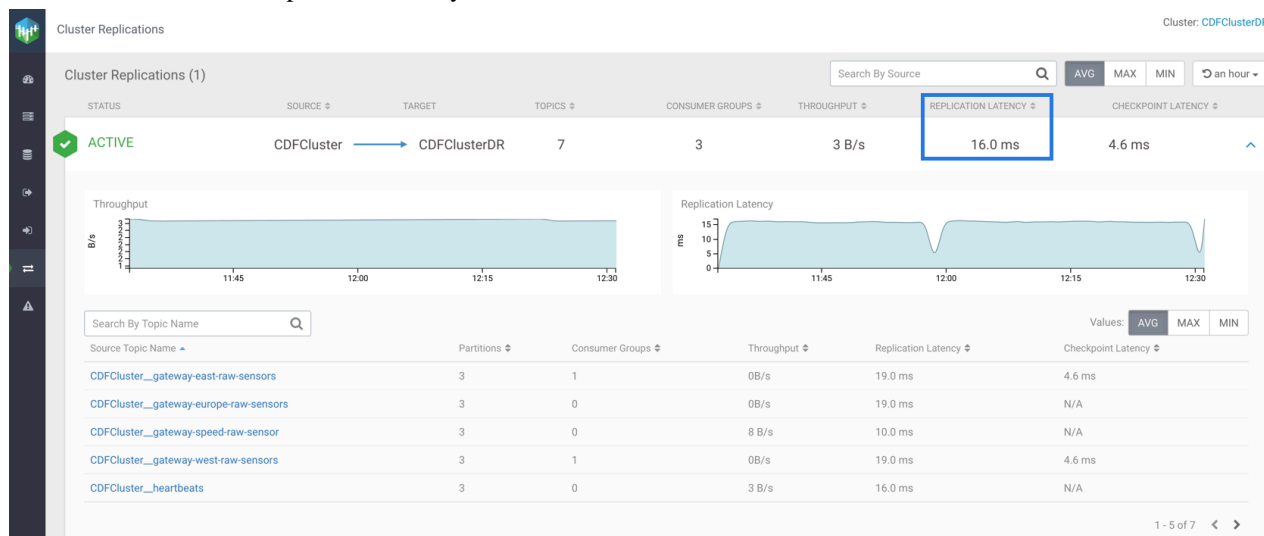
You can also monitor the throughput of a cluster replication graphically. SMM displays the Throughput graph for each cluster in the cluster details. The following image shows the graph for throughput for CDFCluster to CDFClusterDR replication:



Monitoring replication latency for cluster replication


Learn how to monitor the replication latency of a Kafka cluster replication in Streams Messaging Manager (SMM). You can monitor average, maximum, and minimum replication latency of a cluster replication. You can also monitor the replication latency of a cluster replication graphically.

Replication latency is defined as the amount of time taken for a message to get replicated from the source cluster to the destination cluster. Replication latency is measured in milliseconds.



In the image, you can see the average replication latency for the CDFCluster to CDFClusterDR replication is 16.0 milliseconds. You can fetch the maximum or minimum replication latency for the replication by clicking the MAX or MIN button located above the cluster replication.


You can monitor the replication latency of a cluster replication graphically. SMM displays the Replication Latency graphs for each cluster in the cluster details. The following image shows the graph for replication latency for CDFCluster to CDFClusterDR replication:



Monitoring checkpoint latency for cluster replication

Learn how to monitor the checkpoint latency of a Kafka cluster replication in Streams Messaging Manager (SMM). You can monitor average, maximum, and minimum checkpoint latency of a cluster replication. You can also monitor the checkpoint latency of a cluster replication graphically.

Checkpoint latency is defined as the amount of time taken for a message of a topic to checkpoint on the target cluster after the message is committed on the source cluster. Checkpoint latency is measured in milliseconds.



In the image, you can see the average checkpoint latency for the CDFCluster to CDFClusterDR replication is 4.6 milliseconds. You can fetch the maximum or minimum replication latency for the replication by clicking the MAX or MIN button located above the cluster replication.



Note: If the checkpoint latency for a cluster replication or a topic appears as Not Available, then that means there are no consumer groups defined.

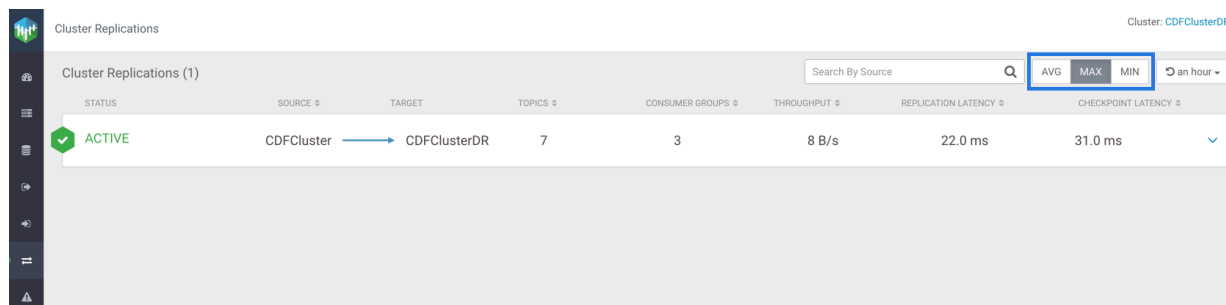
Monitoring replication throughput and latency by values

Learn how to fetch the average, maximum, and minimum values for throughput, replication latency, and checkpoint latency of a Kafka cluster replication in Streams Messaging Manager (SMM).

You can do this at the following levels:

- Cluster Replication Level

Click the AVG, MAX, or MIN buttons, as shown in the following image, to fetch average, maximum, or minimum values of throughput, replication latency, and checkpoint latency for cluster replications.



- Topic Level

Click the AVG, MAX, or MIN buttons, as shown in the following image, to fetch average, maximum, or minimum values of throughput, replication latency, and checkpoint latency for topics.

The screenshot displays the 'Cluster Replications' page in the Kafka Streams Messaging Manager. At the top, it shows the cluster name 'CDFClusterDR' and a search bar. Below this, a summary row indicates the replication status as 'ACTIVE', with 7 topics, 3 consumer groups, a throughput of 3 B/s, a replication latency of 16.0 ms, and a checkpoint latency of 4.6 ms. Two line graphs are provided: 'Throughput' (B/s) and 'Replication Latency' (ms), both showing data from 13:00 to 13:45. A table below the graphs lists replication details for various topics, including 'CDFCluster__gateway-east-raw-sensors', 'CDFCluster__gateway-europe-raw-sensors', 'CDFCluster__gateway-speed-raw-sensor', 'CDFCluster__gateway-west-raw-sensors', and 'CDFCluster__heartbeats'. The table columns are Source Topic Name, Partitions, Consumer Groups, Throughput, Replication Latency, and Checkpoint Latency. A 'Values' dropdown menu is visible, with 'AVG', 'MAX', and 'MIN' options, where 'MIN' is currently selected and highlighted with a red box.

Source Topic Name	Partitions	Consumer Groups	Throughput	Replication Latency	Checkpoint Latency
CDFCluster__gateway-east-raw-sensors	3	1	0B/s	19.0 ms	2.0 ms
CDFCluster__gateway-europe-raw-sensors	3	0	0B/s	19.0 ms	N/A
CDFCluster__gateway-speed-raw-sensor	3	0	4 B/s	8.0 ms	N/A
CDFCluster__gateway-west-raw-sensors	3	1	0B/s	19.0 ms	2.0 ms
CDFCluster__heartbeats	3	0	0B/s	8.0 ms	N/A