

Cloudera Runtime 7.2.12

Planning for Streams Replication Manager

Date published: 2021-09-20

Date modified: 2021-09-20

The Cloudera logo is displayed in a bold, orange, sans-serif font. The word "CLOUDERA" is written in all caps, with a stylized 'E' that has a horizontal bar extending to the right.

<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

Streams Replication Manager requirements.....	4
Streams Replication Manager recommended deployment architecture.....	4

Streams Replication Manager requirements

When planning and designing your Streams Replication Manager (SRM) deployment that has multiple installations of SRM, there are a number of practices that you must follow. Otherwise, you might encounter issues when you are replicating data with SRM.

Kafka cluster aliases must be uniform across clusters

When you have multiple SRM installations, the cluster aliases must be defined the same way in each installation's configuration. SRM utilizes the cluster alias to detect replication cycles and as a result avoid infinite replication loops.

The heartbeats topic must be present on all source clusters

To be able to properly monitor and discover replications, the heartbeats topic must be present on all source clusters. The heartbeats topic is created by the SRM Driver when a target cluster is set with the Streams Replication Manager Driver Target Cluster property. This means that to ensure that the heartbeats topic is created, and has data written into it periodically, at least a single SRM Driver cluster must target each source cluster.

Driver targets must be managed correctly

The Streams Replication Manager Cluster alias and the Streams Replication Manager Driver Target Cluster properties must be managed correctly when you have multiple SRM installations. If a target is specified for the replication in the Streams Replication Manager Driver Target Cluster property, a Connect worker is created for each possible cluster pair based on the aliases present in Streams Replication Manager Cluster alias. As a result of this, it is possible for different SRM Driver clusters to merge through their Connect workers. To avoid this, ensure that different SRM Driver installations target different clusters.

Service targets must be managed correctly

The Streams Replication Manager Cluster alias and the Streams Replication Manager Service Target Cluster properties must be managed correctly when you have multiple SRM installations. If a target is specified for monitoring in the Streams Replication Manager Service Target Cluster property, a Kafka Streams application instance is created for each cluster present in the Streams Replication Manager Service Target Cluster property. As a result of this, it is possible for different SRM Service clusters to merge through their Kafka Streams application instances. To avoid this, ensure that different SRM Service installations target different clusters.

Streams Replication Manager recommended deployment architecture

Learn about pull mode, which is the Cloudera recommended deployment architecture for Streams Replication Manager.

While Streams Replication Manager (SRM) can be deployed in many different ways, the Cloudera recommended setup is pull mode. Pull mode refers to an SRM deployment where data replication happens by pulling data from remote source clusters, rather than pushing data into remote target clusters. An SRM deployment that is in pull mode conforms to the following:

- A unique SRM installation is deployed on all clusters that host a target Kafka cluster.
In other words, each target Kafka cluster in the deployment has a co-located SRM installation.
- SRM Drivers only execute replications that target their co-located Kafka cluster.
- SRM Services only target their co-located Kafka cluster.

The reason why pull mode is recommended is because this is the deployment type that was found to provide the highest amount of resilience against various timeout and network instability issues. For example:

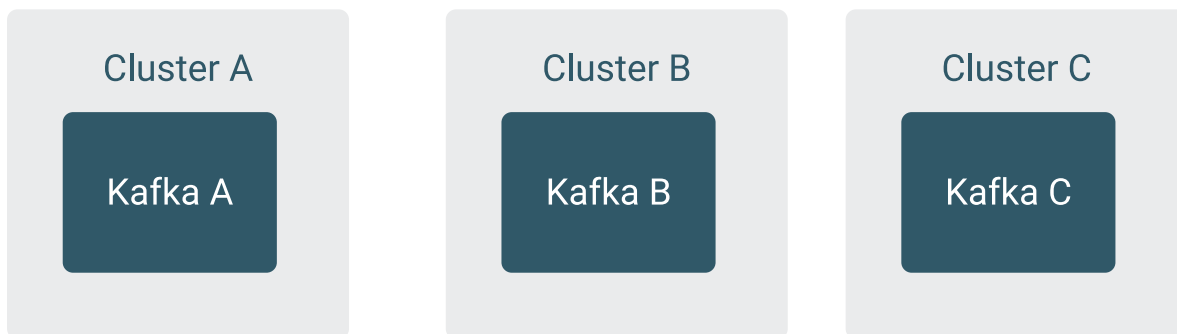
- SRM Drivers run Connect workers which are coordinating through the target cluster of the replication flow. This means that Connect workers are more closely tied to the target cluster than the source cluster. Having the SRM Drivers and the target Kafka cluster closely located to each other minimizes group membership and rebalance timeout issues. Additionally, this also minimizes the network instability on the producer side of the SRM Drivers which reduces data duplication in the target cluster.
- SRM Services also coordinate through the target cluster. This makes them more sensitive to timeout and network partition issues tied to the target cluster. In addition, the SRM Service reads from and writes to the target Kafka through the Kafka Streams application. Having the SRM Services and the target Kafka clusters closely located to each other can minimize timeout and network partition issues. Additionally, for Kafka clusters in the cloud, hosting the SRM Service in the same data center as the Kafka cluster can help with keeping cloud costs to a minimum.
- In a situation where there is a network partition, or one of the clusters in the replication is unavailable, it is preferable to let the target cluster pull the data when the connection is finally established, rather than the source cluster trying to push data indefinitely.

Pull mode deployment example



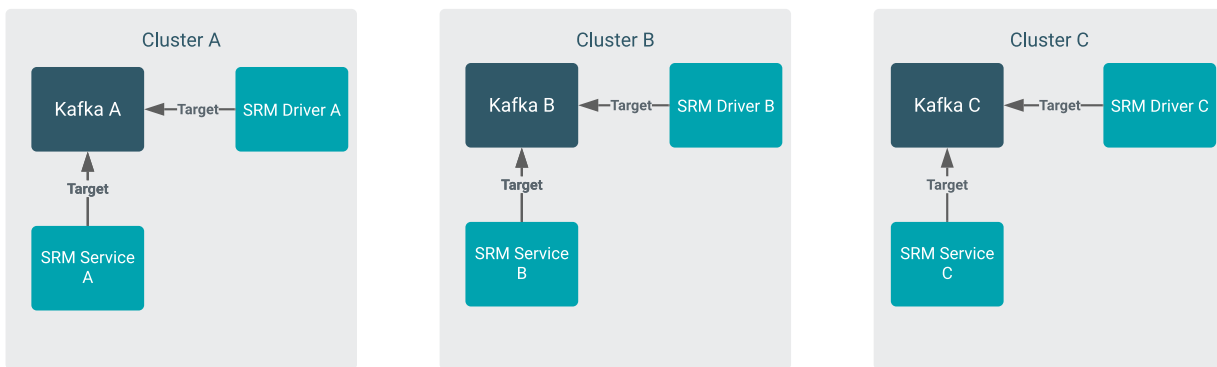
Note: The following simple example does not go into the detail regarding the exact configuration steps required to set up SRM or your clusters. It is meant to provide you with a general overview of how a pull mode deployment looks like.

Consider a simple deployment that has three clusters. Cluster A, B, and C. Each of them has a Kafka cluster.

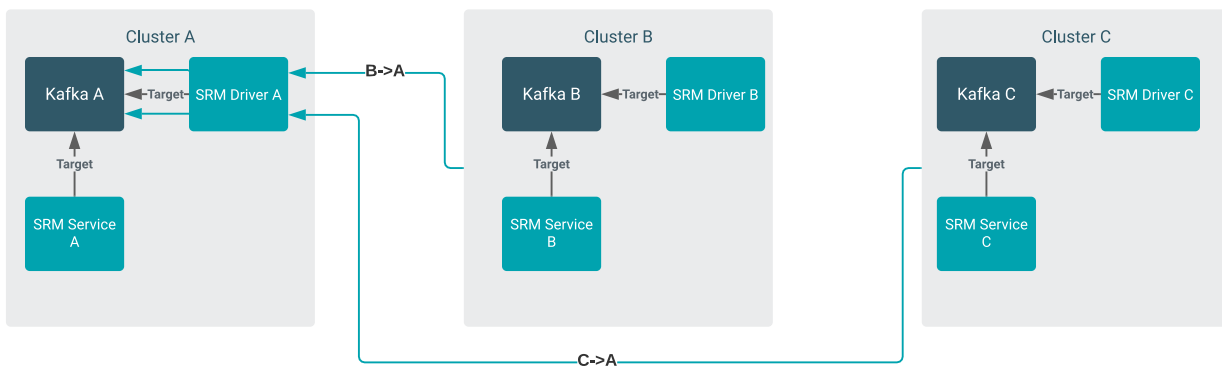


To achieve a pull mode deployment, you must deploy SRM (both Service and Driver) on all three clusters. The Services and Drivers must target their co-located Kafka cluster.

To be more precise, in the case of Driver targets, what you must ensure is that each source cluster is targeted by at least a single Driver. This is required so that a heartbeats topic is created in that cluster. This is ensured in this example because each Kafka has a co-located Driver targeting it. However, if you have a unidirectional replication setup, where the source Kafka cluster does not have a co-located SRM installation, you must ensure that one of your Drivers is targeting that Kafka cluster.



Once setup is complete, you can start configuring your replications. For example, assume that you want to have data replicated from Cluster B and C to cluster A. In this case, you need to enable two separate replications, B->A and C->A. To achieve pull mode, the two replications must be executed by SRM Driver B.



Any number of replications can be set up between the clusters, but you must always ensure that each Driver is only executing the replications targeting their co-located cluster. For example, assume that in addition to the replications set up previously, you also want to set up replication to Cluster B from Cluster A and Cluster C. In this case, the deployment would change as follows.

