

Streams Replication Manager for HDF and HDP 1.0.0

SRM Installation

Date published: 2019-09-20

Date modified: 2019-09-20

CLOUDERA

<https://docs.cloudera.com/>

Legal Notice

© Cloudera Inc. 2022. 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

Deployment recommendations.....	4
System Requirements.....	5
Installation.....	5
Installing from Packages.....	5
Installing from ZIP or TAR Archives.....	6

Deployment recommendations

Learn about the recommended ways to deploy SRM in your system.

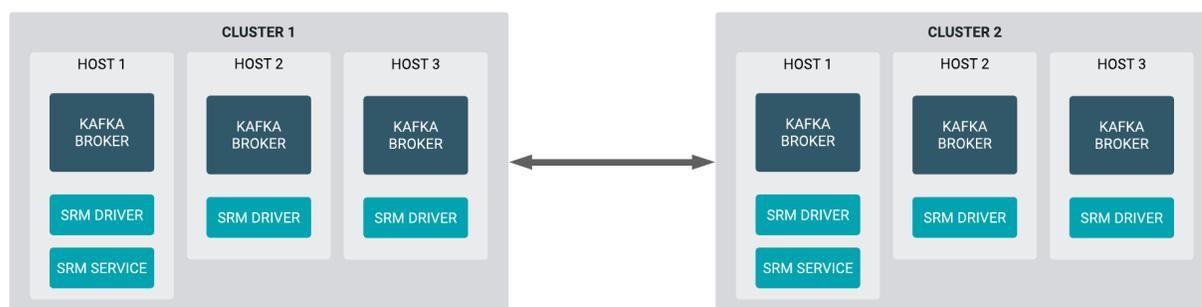
SRM can be deployed in two different ways:

- Co-located with existing Kafka brokers.
- On dedicated machines outside of existing Kafka brokers.

Co-located deployment

Use a co-located deployment approach when your existing Kafka Brokers are not heavily utilized. For this deployment approach, configure each instance of SRM to only write to the cluster that it is running on. This avoids scenarios in which SRM clients deployed on Cluster 1 would write to brokers in Cluster 2. Depending on the available network bandwidth between the two clusters, this approach can increase the replication performance.

Figure 1: Co-located deployment

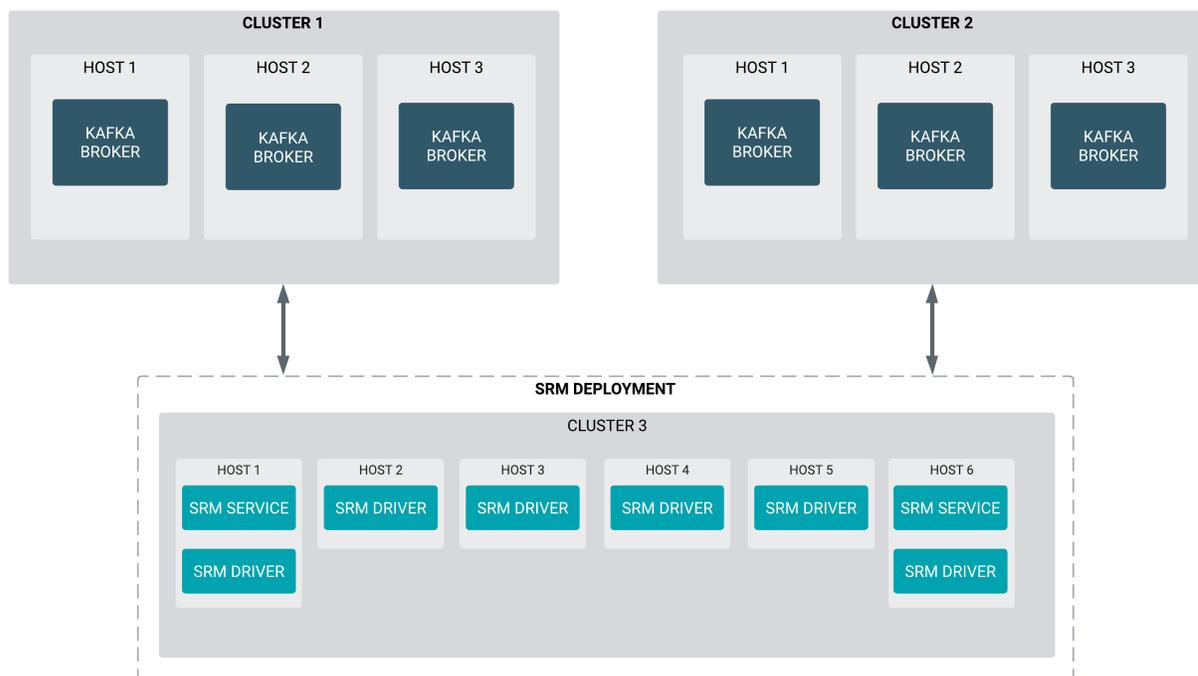


Dedicated deployment

SRM mainly consumes memory and network bandwidth on the hosts it is installed on. If your existing Kafka Brokers do not have a lot of headroom left for these resources, consider deploying SRM on dedicated machines or expanding your Kafka deployments.

If you have data centers with multiple Kafka clusters, a dedicated SRM deployment in each data center can help to simplify management of the SRM nodes. In this scenario, configure each instance SRM to write to all Kafka clusters running in the local data center.

Figure 2: Dedicated deployment



System Requirements

SRM can be installed on an HDF or HDP cluster that is running Kafka as a service. For a comprehensive list of supported operating systems, platforms and services, see the Hortonworks Support Matrix.

Related Information

[Hortonworks Support Matrix](#)

Installation

Installing from Packages

Install SRM from packages.

Before you begin

Obtain the SRM software artifacts from the Cloudera Downloads portal. For more information, see [Download from the SRM repository](#).

Procedure

1. Copy the package to the machine you want to install SRM on.
2. Install SRM:
 - a) For RHEL/CentOS:

```
yum install [SRM_PACKAGE_NAME]
```

- b) For Ubuntu/Debian

```
dpkg -i [SRM_PACKAGE_NAME]
```

- c) For SLES

```
zypper install [SRM_PACKAGE_NAME]
```

Replace `[SRM_PACKAGE_NAME]` with the name of the package you have downloaded. For example, `streams-replication-manager-0.0.1.1.0.0.0-42.x86_64.rpm`.

3. Optional: Specify an alternate temporary-file directory for RocksDB.

This step is only required if you have mounted `/tmp` with the `noexec` option. In this case you have to specify an alternate `tmp` directory for RocksDB. You can achieve this by adding the alternate directory path to `srm-env.sh`:

- a) Navigate to the SRM config directory located at `/opt/streams-replication-manager/config`.
- b) Open `srm-env.sh` with an editor of your choice and add the following environment variable:

```
export ROCKSDB_SHARED_LIB_DIR=/root/tmp/srmrocksdb
```



Note: Verify that the specified directory exists and that the users running SRM have read, write and execute permissions. If the directory does not exist, create it and assign the correct permissions.

Results

SRM is installed on the host.

What to do next

Use the command line tools provided with SRM to start the driver and service.

Related Information

[SRM Command Line Tools](#)

Installing from ZIP or TAR Archives

Install SRM from ZIP or TAR archives.

Before you begin

Obtain the SRM software artifacts for your system from the Cloudera Downloads Portal. For more information, see [Download from the SRM repository](#).

Procedure

1. Create a home directory for SRM.

Cloudera recommends `/opt/streams-replication-manager/`.

```
mkdir -p /opt/streams-replication-manager/
```

2. Extract the files from the downloaded archive.

Use the `unzip` or `tar` command or any other archive manager to extract the archives. For example:

```
tar -xzvf [ARCHIVE_PATH] -C /opt/streams-replication-manager/
```

```
unzip -o [ARCHIVE_PATH] -d /opt/streams-replication-manager/
```

Replace `[ARCHIVE_PATH]` with the path to the archive you have downloaded. For example, `/root/streams-replication-manager-0.0.1.1.0.0.0-42.tgz`.

3. Verify that the directories have correct file permissions. If not assign them.

```
chmod -R 755 /opt/streams-replication-manager/
```

4. Set environment variables:

- a) Navigate to the SRM config directory. For example, `/opt/streams-replication-manager/[ARCHIVE_NAME]/config`.

Replace `[ARCHIVE_NAME]` with the full name of the archive that you have downloaded. For example, `streams-replication-manager-0.0.1.1.0.0.0-42`.

- b) Open `srm-env.sh` with an editor of your choice and add the following environment variables:

```
export SRM_HOME=/opt/streams-replication-manager/[ARCHIVE_NAME]
export SRM_BIN_DIR=${SRM_HOME}/bin
export SRM_CONF_DIR=${SRM_HOME}/config
export SRM_LOG_DIR=${SRM_HOME}/logs
export SRM_PID_DIR=${SRM_HOME}/run
```

- c) Optional: Specify an alternate temporary-file directory for RocksDB.

This step is only required if you have mounted `/tmp` with the `noexec` option. In this case you have to specify an alternate `tmp` directory for RocksDB. You can achieve this by adding the following environment variable to `srm-env.sh`:

```
export ROCKSDB_SHARED_LIB_DIR=/root/tmp/srmrocksdb
```



Note: Verify that the specified directory exists and that the users running SRM have read, write and execute permissions. If the directory does not exist, create it and assign the correct permissions.

5. Optional: Add the SRM bin directory to your `PATH` environment variable. For example:

```
export PATH=/opt/streams-replication-manager/[ARCHIVE_NAME]/bin:$PATH
```

Results

SRM is installed on the host.

What to do next

Use the command line tools provided with SRM to start the driver and service.

Related Information

[SRM Command Line Tools](#)