

Creating your First Streams Messaging Cluster in CDP Public Cloud

Date published: 2019-12-16

Date modified: 2021-06-14

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 three horizontal bars.

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

Creating your first Streams Messaging cluster.....	4
Meet the prerequisites.....	4
Create your cluster.....	4
Give users access to your cluster.....	6
 Next steps.....	 6

Creating your first Streams Messaging cluster

You can create a managed and secured Streams Messaging cluster in a CDP Public Cloud environment by using one of the prescriptive streams messaging cluster definitions and following these steps.

Meet the prerequisites

Before you start creating your first Streams Messaging Data Hub Cluster, you need to ensure that you have set up the environment properly and have all the necessary accesses to use CDP Public Cloud. Use this checklist to make sure that you meet all the requirements before you start creating your first cluster.

- You have CDP login credentials.
- You have an available CDP environment.
- You have a running Data Lake.
- You have a CDP username and the predefined resource role of this user is EnvironmentAdmin.
- Your CDP user is synchronized to the CDP Public Cloud environment.
- Your CDP user has the correct permissions set up in Ranger allowing access to Kafka, Streams Messaging Manager and Schema Registry.



Important: Make sure that the Runtime version of the Data Lake cluster matches the Runtime version of the Data Hub cluster that you are about to create. If these versions do not match, you may encounter warnings and/or errors.

Related Information

[Getting started as a user](#)

[AWS environments](#)

[Azure environments](#)

[GCP environments](#)

[Data lakes](#)

Create your cluster

Once you have met the prerequisites, you are ready to create your Streams Messaging cluster in CDP Public Cloud using a default cluster definition. For this, you must select the cluster definition that matches your cloud provider for the environment and further choose from the light and heavy duty options available.

About this task

Once you have met the prerequisites, you are ready to create your Streams Messaging cluster using a default cluster definition.

Procedure

1. Log into the CDP web interface.
2. Navigate to Management Console Environments, and select the environment where you would like to create a cluster.

3. Click Create Data Hub.

The following page is displayed:

4. Select Cluster Definition.

5. Select the appropriate streams messaging cluster definition from the Cluster Definition dropdown depending on your operational objectives.

There are two definition options available:

- Streams Messaging Heavy Duty for AWS
- Streams Messaging Light Duty for AWS
- Streams Messaging Heavy Duty for Azure
- Streams Messaging Light Duty for Azure
- Streams Messaging Heavy Duty for GCP
- Streams Messaging Light Duty for GCP

For more information on templates, see [Planning your Streams Messaging deployment](#).

The list of services is automatically shown below the selected cluster definition name.


6. Give the cluster a name and add any tags you might need.

You can define tags that will be applied to your cluster-related resources on your cloud provider account. For more information about tags, see [Tags](#).

7. Optional: Use the Advanced Options section to customize the infrastructure settings.


8. Optional: Configure Heavy Duty clusters for SRM.

SRM is not provisioned by default when using the Streams Messaging Heavy Duty definition. If you are using this definition and want to have SRM deployed in the cluster, set the Instance Count of the Srm nodes host group to at least 1.

- Go to Advanced Options Hardware and Storage.
- Locate the Srm nodes host group and click  to edit host group details.
- Set the Instance Count to at least 1.



Important: If you want to deploy SRM in high availability mode, set the instance count to 2 or higher.

- Click  to save your changes.

9. Click Provision Cluster.

Results

You will be redirected to the Data Hub cluster dashboard, and a new tile representing your cluster will appear at the top of the page.

Related Information

[Create a cluster from a definition on AWS](#)

[Create a cluster from a definition on Azure](#)

[Create a cluster from a definition on GCP](#)

[Tags](#)

[Streams Messaging deployment scenarios](#)

Give users access to your cluster

As an EnvironmentAdmin, you need provide the users access to your environment and to the new Streams Messaging cluster by assigning user roles and adding users to Ranger policies.

About this task

The cluster you have created using the Streams Messaging cluster definition is kerberized and secured with SSL. Users can access cluster UIs and endpoints through a secure gateway powered by Apache Knox. Before you can begin working with Kafka, Schema Registry, Streams Messaging Manager, and Streams Replication Manager, you must give users access to the Streams Messaging cluster components.

Procedure

1. Assign the EnvironmentUser role to the users to grant access to the CDP environment and the Streams Messaging cluster.
2. Sync the updated user permissions to the Data Lake (Ranger) using ActionsSynchronize Users to FreeIPA.
This synchronizes the user to the FreeIPA identity management system to enable SSO.
3. Add the user to the appropriate pre-defined Ranger policies.
These policies are specified within the Ranger instance that provides authorization to the Kafka service in your Streams Messaging cluster.
If needed, you can also create a custom Ranger access policy and add the user to it.

Related Information

[Enable authorization in Kafka with Ranger](#)

[Configure the resource-based Ranger service used for authorization](#)

Next steps

After you have created your Streams Messaging Data Hub cluster, you are ready to start working with Kafka workloads in CDP Public Cloud. You can gather information about Kafka and Streams Messaging details in other parts of the CDP Public Cloud documentation.

Related Information

[Connecting Kafka clients to Data Hub provisioned clusters](#)

[Apache Kafka](#)

[Data ingest use cases in Cloudera Data Flow for Data Hub](#)