

Configuring Cruise Control

Date published: 2020-05-04

Date modified: 2020-10-14



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

Add Cruise Control as a service.....	4
Configuring capacity estimations and goals.....	4

Add Cruise Control as a service

You need to use the Add Service wizard in Cloudera Manager to have Cruise Control service on your cluster. After selecting the host you want to add the Cruise Control role to, you need to review the default configurations. These configurations can also be set later after adding Cruise Control as a service.

About this task



Note: It is recommended to add Cruise Control on the same CDP cluster as Kafka.

Procedure

1. Open Cloudera Manager.
2. Select the drop-down menu to the right of your cluster on the Home screen.
3. Select Add Service.
The Add Service wizard starts.
4. Select Cruise Control as the type of service from the list of services, then click Continue.
5. Assign roles to Cruise Control, then click Continue.
6. Check the default and suggested settings for configuration parameters on the Review Changes page, and set if needed.
 - a) Required: Set the Kafka Client Security Protocol as it is set for Kafka.



Important: If the Security Protocol is not the same as it is set in Kafka, the Cruise Control service will not start and the cluster will not operate properly.

7. Click Continue and wait until the first run of the Cruise Control service is completed.
8. Click Continue and then Finish.

Results

You have added Cruise Control as a service in Cloudera Manager.

Configuring capacity estimations and goals

Cruise Control rebalancing works using capacity estimations and goals. You need to configure the capacity estimates based on your resources, and set the goals for Cruise Control to achieve the Kafka partition rebalancing that meets your requirements.

For the rebalancing, you need to provide the capacity values of your resources. These values are used for specifying the rebalancing criteria for your deployment. The following capacity values must be set:

Capacity	Description
capacity.default.cpu	100 by default
capacity.default.network-in	Given by the internet provider
capacity.default.network-out	



Note: For the capacity estimates, the disk capacity value is also needed. However, Cruise Control automatically retrieves the disk capacity value from the `kafka_log_directory_disk_total_space` Kafka metric.

The optimizers in Cruise Control use the network incoming and outgoing capacities to define a boundary for optimization. The capacity estimates are generated and read by Cruise Control. A `capacity.json` file is generated when Cruise Control is started. When a new broker is added, Cruise Control uses the default broker capacity values. However, in case disk related goals are used, Cruise Control must be restarted to load the actual disk capacity metrics of the new broker.

After setting the capacity estimates, you can provide different goals that define the optimization proposals given by Cruise Control. `Default.goals` are used to pre-compute optimization proposals that can be applied regardless of any anomalies. These default.goals on a healthy cluster can optimize resource utilization. Supported goals are also available to assist the optimized rebalancing process. When these goals are fulfilled, the rebalancing is successful. When the goals are violated, self-healing can be used and rebalancing must be carried out.

Cruise Control has an anomaly detection feature where goal violations can also be set. The `anomaly.detection.goals` configuration defines when the goals are not met, thus causing a violation. These anomalies can be fixed by the proposal generated from the `self.healing.goals` configuration. In case there is no self-healing goal specified, Cruise Control uses the `default.goals` setting. Hard goals can also be set to guarantee the fulfilment of any optimization or self-healing process. You can set the anomaly-detection goals, self-healing goals, and hard goals in Cloudera Manager under the Cruise Control Server Advanced Configuration Snippet (Safety Valve) for `cruisecontrol.properties` setting.

You can find the capacity estimate and goal configurations at Cloudera Manager>Cruise Control>Configuration>Main .

The following table lists all the configurations that are needed to configure Cruise Control specifically to your environment. For the list of goals, see the [upstream Cruise Control documentation](#).



Note: The security settings are not listed in the table below.

Configuration	Description
<code>num.metric.fetchers</code>	Parallel threads for fetching metrics from the Cloudera Manager database
<code>default.goals</code>	List of default goals
<code>goals</code>	List of subset goals
<code>partition.metric.sample.store.topic</code>	Storing Cruise Control metrics
<code>broker.metric.sample.store.topic</code>	Storing Cruise Control metrics
<code>partition.metrics.window.ms</code>	Time window size for partition metrics
<code>broker.metrics.window.ms</code>	Time window size for broker metrics
<code>num.partition.metrics.windows</code>	Number of stored partition windows
<code>num.broker.metrics.windows</code>	Number of stored broker windows