

Cloudera Runtime 7.1.8

Configuring Cruise Control

Date published: 2020-05-04

Date modified: 2022-08-30

CLOUdera

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

Legal Notice

© Cloudera Inc. 2023. 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

Adding Cruise Control as a service.....	4
Configuring capacity estimations and goals.....	4
Configuring Metrics Reporter in Cruise Control.....	5
Enabling self-healing in Cruise Control.....	6
Changing the Anomaly Notifier Class value to self-healing.....	6
Enabling self-healing for all or individual anomaly types.....	6
Adding self-healing goals to Cruise Control in Cloudera Manager.....	8

Adding 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.
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 goal settings 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.

For more information about self-healing and goal violation, see the [How Cruise Control self-healing works](#) section.

You can find the capacity estimate and goal configurations at the following location in Cloudera Manager:

1. Go to your cluster in Cloudera Manager.
2. Select Cloudera Manager from the services.
3. Select Cruise Control from the list of Services.
4. Click Configuration.
5. Select Main from the Filters.

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 supported goals
<code>hard.goals</code>	List of goals that any optimization proposal must fulfill
<code>self.healing.goals</code>	List of goals to be used for self-healing relevant anomalies
<code>anomaly.detection.goals</code>	List of goals that the anomaly detector should detect if they are violated
<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

Configuring Metrics Reporter in Cruise Control

You can choose between using the default Cruise Control Metrics Reporter or using the Cloudera Manager Metrics Reporter for fetching metrics in Cruise Control. Cloudera recommends using the Cloudera Manager solution with light installation, and the default solution with heavy installations of Kafka deployments.

Procedure

1. Access Cloudera Manager for the Cruise Control configurations.
 - a) Go to your cluster in Cloudera Manager.
 - b) Select Cruise Control from the list of Services.
 - c) Click on Configuration tab.
2. Search for Metrics Reporter.
3. Select CM metrics reporter or Cruise Control metrics reporter based on your requirements.
4. Click Save changes.
5. Click on Action > Restart next to the Cruise Control service name to restart Cruise Control.

Enabling self-healing in Cruise Control

To enable Cruise Control self-healing, you need to set the Anomaly Notifier Class, enable the self-healing for the anomaly types using the REST API and add self-healing goals in Cloudera Manager.

Changing the Anomaly Notifier Class value to self-healing

You need to change the default anomaly notifier class to self-healing to enable the internal notification in Cruise Control that triggers the automatic self-healing.

Procedure

1. Go to your cluster in Cloudera Manager.
2. Select Cruise Control from the list of Services.
3. Click on Configuration tab.
4. Search for the Cruise Control Server Advanced Configuration Snippet (Safety Valve) for `cruisecontrol.properties` setting.
5. Add the following parameter to the Safety Valve field:

```
anomaly.notifier.class=com.linkedin.kafka.cruisecontrol.detector.notifier.SelfHealingNotifier
```

6. Click Save changes.
7. Click on Action > Restart next to the Cruise Control service name to restart Cruise Control.

Enabling self-healing for all or individual anomaly types

Self-healing is disabled for Cruise Control by default. You can enable self-healing in Cloudera Manager using the `cruisecontrol.properties` configuration, or with a curl POST request and the corresponding anomaly type.

Enabling self-healing in Cloudera Manager

1. Go to your cluster in Cloudera Manager.
2. Select Cruise Control from the list of Services.
3. Click on Configuration tab.
4. Search for the Cruise Control Server Advanced Configuration Snippet (Safety Valve) for `cruisecontrol.properties` setting.

- Choose to enable self-healing for all or only specific anomaly types, and add the corresponding parameter to the Safety Valve field based on your requirements.

For all anomaly types	For specific anomaly types
self.healing.enabled=true	self.healing.broker.failure.enabled=true
	self.healing.goal.violation.enabled=true
	self.healing.disk.failure.enabled=true
	self.healing.topic.anomaly.enabled=true
	self.healing.slow.broker.removal.enabled=true

- Click Save changes.
- Click on Action > Restart next to the Cruise Control service name to restart Cruise Control.

Enabling self-healing using REST API

- Open a command line tool.
- Use ssh and connect to your cluster running Cruise Control.

```
ssh root@<your_hostname>
```

You will be prompted to provide your password.

- Enable self-healing for the required anomaly types using the following POST command:

```
POST /kafkacruisecontrol/admin?enable_self_healing_for=[anomaly_type]
```

The following parameters must be used for anomaly_type:

- GOAL_VIOLATION
- BROKER_FAILURE
- METRIC_ANOMALY
- DISK_FAILURE
- TOPIC_ANOMALY



Note: In case you do not want to enable self-healing for certain anomaly types, you can disable them by using the following command:

```
POST /kafkacruisecontrol/admin?disable_self_healing_for=[anomaly_type]
```

- Check which anomalies are currently in use, and which are detected with the following GET command:

```
GET /kafkacruisecontrol/state
```

When reviewing the state of Cruise Control, you can check the status of Anomaly Detector at the following parameters:

- selfHealingEnabled - Anomaly type for which self-healing is enabled
- selfHealingDisabled - Anomaly type for which self healing is disabled
- recentGoalViolations - Recently detected goal violations
- recentBrokerFailures - Recently detected broker failures
- recentDiskFailures - Recently detected disk failures
- recentMetricAnomalies - Recently detected metric anomalies

Adding self-healing goals to Cruise Control in Cloudera Manager

As self-healing is enabled by default for Cruise Control, you only need to specify the actions Cruise Control should take when detecting anomaly types by providing self-healing goals in Cloudera Manager.

Procedure

1. Access Cloudera Manager for the Cruise Control configurations.
 - a) Go to your cluster in Cloudera Manager.
 - b) Select Cruise Control from the list of Services.
 - c) Click on Configuration tab.
2. Search for Self-Healing Goals.
3. Add the required self-healing goals to the corresponding field.
4. Click Save changes.
5. Click on Action > Restart next to the Cruise Control service name to restart Cruise Control.