

Cloudera Runtime 7.2.17

Managing Cruise Control

Date published: 2019-08-22

Date modified: 2023-06-27

CLOUdera

<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

Rebalancing with Cruise Control.....	4
Cruise Control REST API endpoints.....	4

Rebalancing with Cruise Control

In CDP Public Cloud, when you scale your Kafka nodes that are running in your Data Hub clusters, Cruise Control is responsible for rebalancing the partitions based on the goals previously set.

When using the scaling functionality of Data Hub, you can set the number of Kafka nodes that you want to add or you want to remove, and the partition rebalancing is automatically completed in the background with Cruise Control. You do not need to manually initiate the rebalancing process after managing the Kafka brokers, as the Resize action in Data Hub does this automatically.

For more information, see the [Scaling Kafka broker nodes in a Streams Messaging cluster](#) documentation.

Related Information

[Scaling Kafka broker nodes in a Streams Messaging cluster](#)

Cruise Control REST API endpoints

You can configure and check the status of Cruise Control using the REST API with GET and POST endpoints. You must be aware of the required action or query, and submit a curl command in the command line to use the Cruise Control REST API.

There are two types of REST endpoints for Cruise Control: GET and POST.

For GET

With GET endpoints, you can query information about the rebalancing process, the status of Cruise Control and Kafka brokers. The GET endpoints are read-only operations, and do not have any impact on Cruise Control, Kafka, or the rebalancing process.

For POST

You can use the POST endpoints to change the rebalancing process, modify the number of Kafka brokers, configure Cruise Control, and customize certain sampling and proposal tasks.

Prerequisite for CDP Public Cloud

The ADMIN, USER and VIEWER levels define which endpoint a user can call when managing Cruise Control. After provisioning your Streams Messaging cluster, you must set an access level for user authentication in Cloudera Manager to be able to call the Cruise Control API endpoint.

1. Navigate to Management Console > Environments, and select the environment where you have created your cluster.
2. Select Cloudera Manager from the services.
3. Select Clusters > Cruise Control.
4. Click Configuration.
5. Search for one of the following access levels based on which level you want to add the user:

Security property		Description
ADMIN Level Users	auth_admins	The list of ADMIN level users to have access to all endpoints. The ADMIN level also contains the permissions of the USER and VIEWER level.

Security property		Description
USER Level Users	auth_users	The list of USER level users to have access to all the GET endpoints except bootstrap and train. The USER level also contains the permissions of the VIEWER level.
VIEWER Level Users	auth_viewers	The list of VIEWER level users to have access to the most lightweight kafka_cluster_state, user_tasks and review_board endpoints.

- Click Save Changes.
- Restart the Cruise Control service.

Using the Cruise Control REST API endpoints

In the command line, you need to create the POST or GET commands with curl, and also include the Cruise Control hostname, port, and the required endpoint.

For GET

```
curl -X -u <workload_username>:<workload_password> GET "https://[***MY-DATA-HUB-CLUSTER.COM***]/cdp-proxy-api/cruise-control/kafkacruisecontrol/kafka_cluster_state"
```

For POST

```
curl -X -u <workload_username>:<workload_password> POST "https://[***MY-DATA-HUB-CLUSTER.COM***]/cdp-proxy-api/cruise-control/kafkacruisecontrol/add_broker?brokerid=24"
```

You can find the API base URL for Cruise Control on the Endpoints tab of the **Cluster Details** page: Management Console > Data Hub Clusters > Your Cluster > Endpoints.

The following table summarizes the available operations for GET and POST endpoints. The description column only provides a brief summary of the endpoints. For the detailed description of the Cruise Control REST API, see the [official REST API documentation](#).



Note: Most of the POST actions has a dry-run mode. In dry-run mode, the actions only generate the proposals and estimated results, and not trigger the execution of the proposals. To avoid accidentally triggering the data movement, all of the POST actions are in dry run mode by default. To let Cruise Control actually move the data, you need to explicitly set `dryrun=false`.

For GET

Endpoint	Action	Description
/kafkacruisecontrol/bootstrap	You can start the startup process of Cruise Control and reload all data from Kafka.	This feature can only be used in development mode. The bootstrap endpoint clears the internal state. You can run the bootstrap endpoint in the following modes: <ul style="list-style-type: none"> range - an interval is specified since - a starting point is specified recent - most recent metric samples are used
/kafkacruisecontrol/kafka_cluster_state	You can query partition and replica state.	This includes information for each broker and each partition.
/kafkacruisecontrol/load	You can query the current cluster load.	The Load Monitor needs to be in RUNNING state. This includes the load-per-broker and load-per-host information. Only shows valid partitions with enough metric samples.

Endpoint	Action	Description
/kafkacruisecontrol/partition_load	You can query partition resource utilization.	The partition load is sorted based on the given resource.
/kafkacruisecontrol/permissions	You can query the list of permissions set for the current user	This endpoint can be used to review the configured permissions of the current user that include ADMIN, USER and VIEWER levels. Permissions are not used when security is enabled for Cruise Control.
/kafkacruisecontrol/proposals	You can query the optimization proposals generated on the workload model.	The proposals can be generated based on either the snapshot window or the valid partitions.
/kafkacruisecontrol/review_board	You can check the review processes of Cruise Control.	This endpoint shows the status of ongoing review processes in Cruise Control.
/kafkacruisecontrol/state	You can query the state of Kafka Cruise Control.	This includes the state of monitor, executor, analyzer, and anomaly detector.
/kafkacruisecontrol/train	You can use the endpoint to train the partition of Cruise Control for CPU estimation.	This endpoint uses a regression algorithm to train the CPU estimation of the Cruise Control partition. This is an experimental feature.
/kafkacruisecontrol/user_tasks	You can query the user request result.	This includes a full list of all the active and completed tasks in Cruise Control.

For POST

Endpoint	Action	Description
/kafkacruisecontrol/add_broker?brokerid=[id1,id2...]	You can add a list of brokers from a Kafka cluster.	Only moves replicas from existing brokers to the newly added brokers. You can choose to throttle the movement of the replica to the newly added broker. The replica movement is throttled on the existing broker. Dry-run mode is activated by default, set dryrun=false to trigger adding brokers.
/kafkacruisecontrol/admin?concurrent_partition_movements_per_broker=[integer]	You can increase or decrease the execution concurrency.	The configuration is only possible with administrator rights.
/kafkacruisecontrol/admin?disable_self_healing_for=[anomaly_type]	You can enable or disable self-healing for Cruise Control.	The configuration is only possible with administrator rights.
/kafkacruisecontrol/admin?drop_recently_removed_brokers=[broker_ids]	You can drop recently removed or demoted Kafka brokers.	The configuration is only possible with administrator rights.
/kafkacruisecontrol/demote_broker?brokerid=[id1, id2...]	You can move all the leader replicas from a list of brokers.	This makes all the replicas on a given broker to be preferred the least, and triggers a preferred leader replica election to migrate the leader replicas off the broker. Dry-run mode is activated by default, set dryrun=false to trigger demoting brokers.
/kafkacruisecontrol/demote_broker?brokerid_and_logdirs=[id1-logdir1, id2-logdir2...]	You can move all the leader replicas from a list of disks.	This makes all the replicas on a given disk to be preferred the least, and triggers a preferred leader replica election to migrate the leader replicas off the disk.
/kafkacruisecontrol/fix_offline_replicas	You can fix offline replicas in a Kafka cluster.	If the specified topic has offline replicas, they are still moved to healthy brokers. Dry-run mode is activated by default, set dryrun=false to trigger fixing offline replicas.
/kafkacruisecontrol/pause_sampling	You can pause an ongoing metrics sampling process.	The cause of the pause is recorded and shows up when querying the state of Cruise Control.

Endpoint	Action	Description
/kafkacruisecontrol/rebalance	You can trigger a workload balance for a Kafka cluster.	Can be based on valid windows or valid partitions. You can specify the goal you want to use. Dry-run mode is activated by default, set <code>dryrun=false</code> to trigger rebalancing.
/kafkacruisecontrol/remove_broker?brokerid=[id1,id2...]	You can remove a list of brokers from a Kafka cluster.	Only moves partitions from the brokers to be removed to the other existing brokers. You can also specify the destination broker. You can choose to throttle the removed broker during the partition movement. Dry-run mode is activated by default, set <code>dryrun=false</code> to trigger removing brokers.
/kafkacruisecontrol/resume_sampling	You can resume a paused metrics load sampling.	The cause of the resumption is recorded and shows up when querying the state of Cruise Control.
/kafkacruisecontrol/review	You can review the pending Cruise Control requests.	You can check the previously submitted requests in Cruise Control.
/kafkacruisecontrol/stop_proposal_execution	You can stop an ongoing proposal execution task.	You can stop a proposal when the execution of a task is still ongoing.
/kafkacruisecontrol/topic_configuration?topic=[topic_regex]&replication_factor=[target_replication_factor]	You can change the replication factor of a topic.	Existing replicas are not moved. You can determine which replica should be deleted and which broker should be assigned to the new replica. Dry-run mode is activated by default, set <code>dryrun=false</code> to trigger changing the replication factor of a topic.