Cloudera Runtime 7.3.2

Cruise Control

Date published: 2019-08-22
Date modified: 2026-03-31

CLOUDERA

Legal Notice

Contents

Kafka cluster load balancing using Cruise Control

You can use Cruise Control as a load balancing component in large Kafka installations to automatically balance the partitions based on specific conditions for your deployment. The elements in the Cruise Control architecture are responsible for different parts of the rebalancing process that uses Kafka metrics and optimization goals.

The following illustration shows the architecture of Cruise Control.
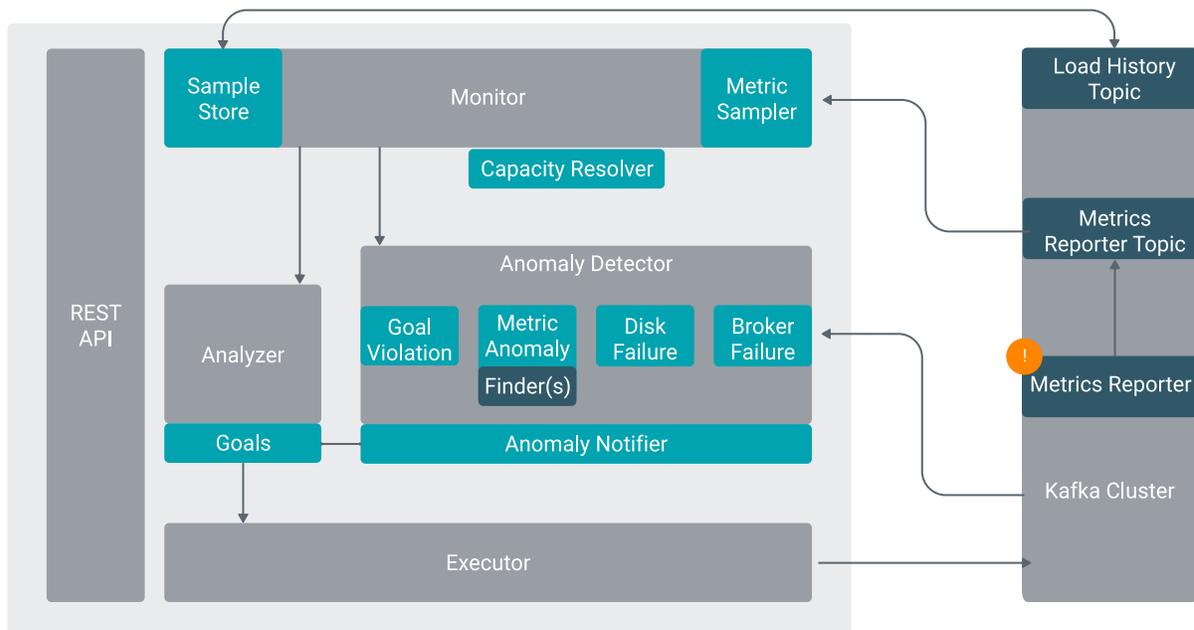
Load Monitor

Generates a cluster workload model based on standard Kafka metrics and resource metrics to utilize disk, CPU, bytes-in rate, and bytes-out rate. Feeds the cluster model into Anomaly Detector and Analyzer.

Analyzer

Generates optimization proposals based on optimization goals provided by the user, and cluster workload model from Load Monitor. Hard goals and soft goals can be set. Hard goals must be fulfilled, while soft goals can be left unfulfilled if hard goals are reached. The optimization fails if the hard goal is violated by optimization results.

Anomaly Detector

Responsible for detecting anomalies that can happen during the rebalancing process. The following anomaly detections are supported in Cruise Control:

* Broker failure
* Goal violations
* Disk failure
* Slow broker as Metric Anomaly
* Topic replication factor

The detected anomalies can be resolved by the self-healing feature of Cruise Control. For more information, see the *How Cruise Control self-healing works* documentation.

Executor

Carries out the optimization proposals and it can be safely interrupted when executing proposals. The executions are always resource-aware processes.

Related Information

How Cruise Control self-healing works

How Cruise Control rebalancing works

For the rebalancing process, Cruise Control creates a workload model based on the resources of the cluster, such as CPU, disk and network load, also known as capacity estimates. This workload model will be the foundation of the optimization proposal to rebalance the Kafka cluster, which can be further customized for your requirements using optimization goals.

During the rebalancing operation, the Kafka brokers are checked to determine if they met the requirements set by the selected goals. If the goals are fulfilled, the rebalancing process is not triggered. If a goal is not fulfilled, the rebalancing process is triggered, and partitions, replicas and topics are reassigned between the brokers until the requirements set by the goal are met. The rebalancing fails in case the hard goals are not met at the end of the reassignment process.

The following goals can be defined for the rebalancing process:

Hard goals

- List of goals that any optimization proposal must fulfill. If the selected hard goals are not met, the rebalancing process fails.

Default goals

- 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 if self-healing goals are not specified.
- The value of the Default goals must be a subset of the Supported goals, and a superset of the Hard and Self-healing goals.
- If there are no goals specified as query parameters, the Default goals will be used for the rebalancing process.

Supported goals

- List of supported goals to assist the optimized rebalancing process.

When the hard, default and supported goals are fulfilled, the rebalancing is successful. If there are any goal violations, self-healing can be used.

Self-healing goals

- List of goals to be used for self-healing relevant anomalies. If there are no self-healing goals defined, the Default goals are used for self-healing.
- If the rebalancing process is triggered by self-healing and the Self-healing goals list is not empty, then the Self-healing goals will be used to create the optimization proposal instead of the Default goals.
- The value of the Self-healing goals must be a subset of the Supported goals and a superset of the Hard goals.

Anomaly detection goals

- List of goals that the Anomaly detector should detect if violated. It must be a subset of the self-healing goals and thus also of Default goals.
- The value of the Anomaly detection goals must be a subset of the Self-healing goals. Otherwise, the self-healing process is not able to resolve the goal violation anomaly.

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 goal 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 fulfillment of any optimization or self-healing process.

How Cruise Control retrieves metrics

Cruise Control creates metric samples using the retrieved raw metrics from Kafka. The metric samples are used to set up the cluster workload model for the Load Monitor.

In Load Monitor, the Metric Fetcher Manager is responsible for coordinating all the sampling tasks: the Metric Sampling Task, the Bootstrap Task, and the Linear Model Training Task.

Each sampling task is carried out by a configured number of Metric Fetcher threads. Each Metric Fetcher thread uses a pluggable Metric Sampler to fetch samples. Each Metric Fetcher is assigned with a few partitions in the cluster to get the samples. The metric samples are organized by the Metric Sample Aggregator that puts each metric sample into a workload snapshot according to the timestamp of a metric sample.

The cluster workload model is the primary output of the Load Monitor. The cluster workload model reflects the current replica assignment of the cluster and provides interfaces to move partitions or replicas. These interfaces are used by the Analyzer to generate optimization solutions.

The Sample Store stores the metric and training samples for future use.

With the metric sampler, you can deploy Cruise Control to various environments and work with the existing metric system.

Cruise Control Metrics Reporter produces raw metrics directly to a Kafka topic by CruiseControlMetricsReporter. Then these metrics are fetched by Cruise Control and metric samples are created and stored back to Kafka. The samples are used as the building blocks of cluster models.

How Cruise Control self-healing works

The Anomaly detector is responsible for the self-healing feature of Cruise Control. When self-healing is enabled in Cruise Control, the detected anomalies can trigger the attempt to automatically fix certain types of failure such as broker failure, disk failure, goal violations and other anomalies.

Broker failures

> The anomaly is detected when a non-empty broker crashes or when a broker is removed from a cluster. This results in offline replicas and under-replicated partitions.

> When the broker failure is detected, Cruise Control receives an internal notification. If self-healing for this anomaly type is enabled, Cruise Control will trigger an operation to move all the offline replicas to other healthy brokers in the cluster. Because brokers can be removed in general during cluster management, the anomaly detector provides a configurable period of time before the notifier is triggered and the self-healing process starts.

> If a broker disappears from a cluster at a timestamp specified as T, the detector will start a countdown. If the broker did not rejoin the cluster within the broker.failure.alert.threshold.ms since the specified T, the broker is defined as dead and an alert is triggered. Within the defined time of broker.failure.self.healing.threshold.ms, the self-healing process is activated and the failed broker is decommissioned.

Goal violations

> The anomaly is detected when an optimization goal is violated.

> When an optimization goal is violated, Cruise Control receives an internal notification. If self-healing for this anomaly type is enabled, Cruise Control will proactively attempt to address the goal violation by automatically analyzing the workload, and executing optimization proposals.

You need to configure the anomaly.detection.goals property if you enable self-healing for goal violations to choose which type of goal violations should be considered. By default, the following goals are set for anomaly.detection.goals:

- com.linkedin.kafka.cruisecontrol.analyzer.goals.RackAwareGoal
- com.linkedin.kafka.cruisecontrol.analyzer.goals.ReplicaCapacityGoal
- com.linkedin.kafka.cruisecontrol.analyzer.goals.DiskCapacityGoal

Disk failure

The anomaly is detected when one of the non-empty disks dies.

> **Note:** This is the only anomaly that is related to Kafka broker running on a JBOD disk.

When a disk fails, Cruise Control will send out a notification. If self-healing for this anomaly type is enabled, Cruise Control will trigger an operation to move all the offline replicas by replicating alive replicas to other healthy brokers in the cluster.

Slow broker

The anomaly is detected when based on the configured thresholds, a broker is identified as a slow broker.

Within the Metric anomaly, you can set the threshold to identify slow brokers. In case a broker is identified as slow, the broker will be decommissioned or demoted based on the configuration you set for the remove.slow.broker. The remove.slow.broker configuration can be set to true, this means the slow broker will be removed. When setting the remove.slow.broker configuration to false, the slow broker will be demoted.

For more information about configuring the slow broker finder, see the official Cruise Control documentation.

Topic replication factor

The anomaly is detected when a topic partition does not have the desired replication factor.

The topic partition will be reconfigured to match the desired replication factor.