

Configuration

Date published: 2019-12-17

Date modified: 2019-12-17



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

| | |
|---|----------|
| How to configure Apache Flink?..... | 4 |
| Setting max parallelism..... | 4 |
| Configuring Flink application resources..... | 4 |
| Configuring RocksDB state backend..... | 5 |

How to configure Apache Flink?

Cloudera Streaming Analytics includes Flink with configuration that works out of the box. It is not mandatory to configure Flink to production, but you can use the available configurations to optimize the application behavior in production. Cloudera Manager includes all the necessary configurations for Flink that can also be accessed from the `flink-conf.yaml` file.

Setting max parallelism

The max parallelism is the most essential part of resource configuration for Flink applications as it defines the maximum jobs that are executed at the same time in parallel instances. However, you can optimize max parallelism in case your production goals differ from the default settings.

In a Flink application, the different tasks are split into several parallel instances for execution. The number of parallel instances for a task is called parallelism. Parallelism can be defined at the operator, client, execution environment and system level. Cloudera recommends setting parallelism to a lower value at the initially and increasing it over time, if the job cannot keep up with the input rate.

To configure the max parallelism, `setMaxParallelism` is called as it controls the number of key-groups created by the state backends. A key-group is a partition of an operator state. The number of key-groups determines how data is going to be distributed among the parallel operators. If the key-groups are not distributed evenly, the data distribution is also uneven.

Consider the following aspects when setting the max parallelism:

- The number should be large enough to accommodate expected future load increases as this setting cannot be changed without starting from an empty state.
- If `P` is the selected parallelism for the job, the max parallelism should be divisible by `P` to get even state distribution.
- Please note that larger max parallelism settings have greater cost on the state backend side, for large scale production jobs benchmarking the size of the state based on the maximum parallelism is useful before changing this parameter.

Based on these criteria, Cloudera recommends setting the max parallelism to factorials or other numbers with a large number of divisors (120, 180, 240, 360, 720, 840, 1260), which will make parallelism tuning easier.

Table 1: Reference values

| Stateless | In-memory state | RocksDB state |
|-------------------------------|------------------------------|-----------------------------|
| 1 million record / sec / core | 100 000 records / sec / core | 10 000 records / sec / core |

Configuring Flink application resources

Generally, Flink automatically identifies the resources for an application. You can configure the Taskmanager, buffer timeout and high availability for a Flink application to maintain resources in a more efficient way.

YARN automatically kills application containers that use more memory than their allocated limit. To avoid Flink TaskManagers getting killed by YARN use the following calculation to set the size: $TM\ total\ size = TM\ Heap + Heap\ -cutoff$. The default heap cutoff is 25% and it is also configurable.

Furthermore, use a YARN queue with preemption disabled to avoid long running jobs being affected when the cluster reaches its capacity limit.

Flink uses network buffers to transfer data from one operator to another. These buffers are filled up with data during the specified time for the timeout. In case of high data rates, the set time is usually never reached. For cases when the data rate is high the throughput can be further increased with setting the buffer timeout to an intentionally higher value due to the characteristics of the TCP channel. However this in turn increases the latency of the pipeline.

Flink on YARN jobs are configured to tolerate a maximum number of failed containers before they terminate. Configure the YARN maximum failed containers setting in proportion to the total parallelism and the expected lifetime of the job.

High Availability is enabled by default in CSA. This eliminates the Job Manager as a single point of failure. You can also tune the application resilience by setting the YARN maximum application attempts, which determines how many times the application will retry in case of failures.

Table 2: Reference values

| Configuration | Parameter | Recommended value |
|----------------------------|---------------------------------|-----------------------------------|
| TM container memory | -ytm / taskmanager.heap.size | <i>TM Heap + Heap-cutoff</i> |
| Max parallelism | env.setMaxParallelism(num) | <i>120,720,1260,5040</i> |
| Container heap cutoff | containerized.heap-cutoff-ratio | <i>0.25-0.75</i> |
| Buffer timeout | env.setBufferTimeout(millis) | <i>1-100</i> |
| YARN queue | -yqu | <i>A queue with no preemption</i> |
| YARN max failed containers | yarn.maximum-failed-containers | <i>3*num_containers</i> |
| YARN max AM failures | yarn.application-attempts | <i>3-5</i> |

Configuring RocksDB state backend

You can use RocksDB as a state backend when your Flink streaming application requires a larger state that doesn't fit easily in memory. The RocksDB state backend uses a combination of fast in-memory cache and optimized disk based lookups to manage state.

You can configure the state backend for your streaming application by using the state.backend parameter directly or in Cloudera Manager under the Configuration tab:



State Backend

FLINK-1 (Service-Wide) [Undo](#)

state.backend

 state_backend

FILESYSTEM

ROCKSDB

You can adjust how much memory RocksDB should use as a cache to increase lookup performance by setting the memory managed fraction of the TaskManagers in Cloudera Manager under the Configuration tab:



TaskManager Managed Memory Fraction

FLINK-1 (Service-Wide)

taskmanager.memory.managed.fraction

 taskmanager_managed_memory_fraction

The default fraction value is 0.4, but with larger cache requirements you need to increase this value together with the total memory size.

Related Information

[Stateful Tutorial: Configure the application for production](#)