

Running queries on system tables (Preview)

Date published: 2025-04-23

Date modified: 2025-04-23

Legal Notice

© Cludera Inc. 2025. All rights reserved.

The documentation is and contains Cludera 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 Cludera software may be found within the documentation accompanying each component in a particular release.

Cludera 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 Cludera software product page for more information on Cludera software. For more information on Cludera support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Cludera reserves the right to change any products at any time, and without notice. Cludera assumes no responsibility nor liability arising from the use of products, except as expressly agreed to in writing by Cludera.

Cludera, Cludera Altus, HUE, Impala, Cludera Impala, and other Cludera 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.

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cludera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

Contents

Legal Notice	2
Contents	3
Running queries on system tables	4
Coordinator-only request pools	4
Configuring the only coordinators request pool	4
Limitation	5

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

Running queries on system tables

Learn how queries against system table queries can be configured to consider only the coordinator resources that are needed.

Queries against Impala system tables, such as `sys.impala_query_live`, experienced delays due to admission control constraints. These queries, requiring only coordinator resources, were blocked by other queries competing for executor resources.

To address this, Impala introduces an "only coordinators" request pool. This allows system table queries to run without waiting for executor resources, even when those resources are fully used. In Cloudera Data Warehouse, queries submitted to an only coordinators request pool continue to run even if no executors are running.

Technical Preview: This feature is Technical Preview except for workload aware autoscaling virtual warehouses where it is not supported.

Coordinator-only request pools

You can configure coordinator-only request pools in Impala by setting the `<onlyCoordinators>` option to `true` in the `fair-scheduler.xml` file. When enabled, the request pool runs queries only on coordinators. No executors are required, and all fragment instances are executed exclusively on the coordinators.

Warning: Any query can be submitted to a coordinator-only request pool. If a large or non-system table query is submitted, it runs entirely on the coordinators and may lead to memory or CPU exhaustion.

To avoid resource issues, Cloudera recommends limiting access to coordinator-only request pools to a small group of users who understand the risks of running complex queries in this configuration.

Configuring the only coordinators request pool

Learn how to configure a coordinator-only request pool to optimize system table queries.

About the task

Configuring a coordinator-only request pool allows you to run queries that do not require executors, reducing resource constraints and improving performance.

Prerequisites

- Ensure you have administrator access to modify Impala configurations.

Steps

1. Log in to the Cloudera web interface and navigate to the Cloudera Data Warehouse service.
2. In the Cloudera Data Warehouse service, click Virtual Warehouses in the left navigation panel.
3. Select the Impala Virtual Warehouse, click options for the warehouse you want to include the `onlyCoordinators` setting for a request pool.
4. Click Edit and navigate to Impala Coordinator under the Configurations tab.
5. Select the `fair-scheduler.xml` under Configuration files.
6. Add `<onlyCoordinators>true</onlyCoordinators>`. Example:

Unset

```
<queue name="coords">
  <maxResources>10000 mb,0 vcores</maxResources>
  <onlyCoordinators>true</onlyCoordinators>
</queue>
```

7. Click Apply Changes and restart Impala.

Limitation

This feature is not supported for virtual warehouses that use workload-aware autoscaling.

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.