

Monitoring Apache Impala

Date published: 2020-11-30

Date modified: 2021-10-29

The Cloudera logo is displayed in a bold, orange, sans-serif font. The word "CLOUDERA" is written in all caps, with the letter 'E' in "CLouDERA" featuring a unique design where the top bar is a horizontal line that extends to the right and then turns down to form the top of the letter 'R'.

<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

- Impala lineage.....4**
- Debug Web UI for Query Timeline..... 4**
- Improvement in Catalog Observability..... 5**

Impala lineage

You can use the Atlas lineage graph to understand the source and impact of data and changes to data over time and across all your data.

Atlas collects metadata from Impala to represent the lineage among data assets. The Atlas lineage graph shows the input and output processes that the current entity participated in. Entities are included if they were inputs to processes that lead to the current entity or they are output from processes for which the current entity was an input. Impala processes follow this pattern.

Note that lineage is not updated between a table and views that the table is a part of when an Impala ALTER TABLE operation runs on the table.


Debug Web UI for Query Timeline

For a detailed report on how a query was executed and to understand the detailed performance characteristics of a query, you can use the built-in web server's UI and look at the Gantt chart.

The Gantt chart is an alternative to the profile command and is a graphical display in the WebUI that renders timing information and dependencies. This chart allows you to determine where queries are spending time instead of looking for timing information in the profile. Once you identify which fragments or operators are consuming the most time you may need to look at the summary profile or full profile to drill down further.

In subsequent releases, the timeline will be enhanced to display additional details so that looking at the profiles will not be necessary.

Here is the screenshot of the Gantt chart that shows the Query timeline of a TPC_DS query.



Improvement in Catalog Observability

This release brings significant enhancements to the Impala Catalog Web UI, addressing performance issues related to delays in processing Hive Metastore (HMS) events, which can lead to queries using outdated metadata

These improvements specifically target slow Data Definition Language (DDL) and Data Manipulation Language (DML) operations such as INSERT, LOAD DATA statements, and REFRESH operations. Additionally, it addresses operations stuck in various stages such as startup, among other scenarios. Prior to this release, troubleshooting these issues required analyzing logs, but now the process has been streamlined by enhancing the observability of the catalog across multiple dimensions, including the catalogd WebUI, DDL/DML profiling, and additional log reports on sluggish operations.

Key enhancements include:

- Catalogd WebUI: Utilize the catalogd WebUI to monitor catalog operations.
- DDL/DML Profiling: Monitor DDL/DML profile for detailed execution insights within catalogd.
- Log Reports: Analyze logs for a comprehensive review of slow operations, offering valuable insights into the root causes of performance issues.

These improvements empower users with a more user-friendly and effective means of diagnosing and resolving catalog-related performance challenges.