

Cloudera Runtime 7.2.7

Troubleshooting Data Analytics Studio

Date published: 2020-05-22

Date modified: 2021-02-04

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

Problem area: Queries page.....	4
Queries are not appearing on the Queries page.....	4
Query column is empty but you can see the DAG ID and Application ID.....	4
Cannot see the DAG ID and the Application ID.....	5
Cannot view queries of other users.....	5
Problem area: Compose page.....	5
Cannot see databases, or the query editor is missing.....	5
Unable to view new databases and tables, or unable to see changes to the existing databases or tables.....	6
Troubleshooting replication failure in the DAS Event Processor.....	6
Problem area: Reports page.....	7
How DAS helps to debug Hive on Tez queries.....	7

Problem area: Queries page

Issues related to queries can arise mostly because of misconfigured Hive or Tez hooks. You may also experience trouble viewing your queries if you are a non-admin user. Non-admin users can view only their queries.

Queries are not appearing on the Queries page

If queries are not appearing on the **Queries** page, then try the following:

- Check your role. Only admin users can view queries and data of other users. Non-admin users can view only the queries that they have run.

To enlist yourself as the admin user, go to Cloudera Manager Clusters DAS Configuration , and add your user name in the Admin Users field.

For example, hive,admin1,admin2. To make all users administrators, specify an asterik (*).

- Make sure that the DAS services are running as the hive service user.
- Go to Cloudera Manager Clusters HIVE_ON_TEZ and verify whether the logs are added to the path specified in the Base Directory for Hive Proto Hook field.
- Go to Cloudera Manager Clusters TEZ and verify whether the logs are added to the path specified in the Tez history events directory field.

Use the following command to get the path to the logs for Hive:

```
hdfs dfs -ls /warehouse/tablespace/external/hive/sys.db/query_data
```

Use the following commands to get the paths to the logs for Tez:

```
hdfs dfs -ls /warehouse/tablespace/external/hive/sys.db/dag_data
```

```
hdfs dfs -ls /warehouse/tablespace/external/hive/sys.db/dag_meta
```

```
hdfs dfs -ls /warehouse/tablespace/external/hive/sys.db/app_data
```

If you have not run any queries, then the log directories can be empty.

- Verify whether you have the required access and permission to directories specified in the Base Directory for Hive Proto Hook and Tez history events directory field by checking errors from EventProcessorPipeline in the /var/log/das/event-processor.log file.
- If you have a large number of users running Hive queries simultaneously using DAS, the connections from Hive can get exhausted. This sends the query to a queue in Hive which takes a long time to start running. This can result in the following:
 - If a query is fired from DAS, it does not return any log or indicate that the query has been queued.
 - The query does not appear on the UI, because no event is received by DAS.

Evaluate your cluster usage. You may need to add more machines to the clusters to resolve this issue.

Query column is empty but you can see the DAG ID and Application ID

If the query column is empty but you still see the DAG ID and the Application ID, then check the following:

- Check whether Hive has access to write into the directory specified in the Base Directory for Hive Proto Hook field.

- Check whether Hive can write into the directory specified in the Base Directory for Hive Proto Hook field.
- Check whether Hive has access to read from the directory specified in the Base Directory for Hive Proto Hook field.

Cannot see the DAG ID and the Application ID

If you cannot see the the DAG ID and the Application ID for a query, even when the query column is not empty, check the following:

- Check whether the Tez job has access to write into the directory specified in the Tez history events directory field.
- Check whether the data is being written into the following subdirectories of the directory specified in the Tez history events directory directory configuration: app_data, dag_data, and dag_meta.
- Check whether das has access to read from the following subdirectories of the directory specified in the Tez history events directory directory configuration: app_data, dag_data, and dag_meta.

Cannot view queries of other users

Check your role. Only admin users can view queries and data of other users.

To enlist yourself as the admin user, go to Cloudera Manager Clusters DAS Configuration , and add your user name in the Admin Users field.

For example, hive,admin1,admin2. To make all users administrators, specify an asterik (*).

Problem area: Compose page

Misconfigured repl in Hive can affect functionality on the **Compose** page. Issues can also occur if HIVE_ON_TEZ is not running.

Cannot see databases, or the query editor is missing

If you cannot see the databases or the query editor on the **Compose page**, then try the following:

- Verify whether the HS2 service is running by going to Cloudera Manager Clusters HIVE_ON_TEZ Status .
- Verify whether the Hive configuration parameter in the Metastore Transactional Listener List field is set to org.apache.hive.hcatalog.listener.DbNotificationListener by going to Cloudera Manager Clusters HIVE_ON_TEZ Configuration .
- Ideally, the DAS Event Processor runs the repl dump command as a hive service user. The hive service user can run the repl dump command if the user has repl admin privileges. Verify whether DAS can read the repl root directory permissions, especially if DAS and hive are running as different service users.

There may be exceptions in the /var/log/das/event-processor.log file with repl dump failures because of read errors.



Note: If the DAS service cannot read the repl root directory permissions, then it can cause a number of repl directories to be created, but not deleted.

Unable to view new databases and tables, or unable to see changes to the existing databases or tables

If you cannot see the new databases or tables, or the changes that you made to the existing tables, then try the following:

- Verify whether the HIVE_ON_TEZ service is running by going to Cloudera Manager Clusters HIVE_ON_TEZ Status .
- This can also happen if the Hive metadata replication fails. To troubleshoot replication-related failures, see [Troubleshooting replication failure in the DAS Event Processor](#) on page 6.

Troubleshooting replication failure in the DAS Event Processor

DAS uses replication as a way to copy database and table metadata information from Hive to the DAS PostgreSQL database. If the replication fails, then you may not be able to see database or table information in DAS.

The Hive metadata replication process occurs in the following two phases:

- Bootstrap dump
- Incremental dump

When the DAS Event Processor is started for the first time, the entire Hive database and the table metadata is copied into DAS. This is known as the bootstrap dump. After this phase, only the differences are copied into DAS at one-minute intervals from the time the last successful dump was run. This is known as an incremental dump.

If the bootstrap dump never succeeded, then you may not see any database or table information in DAS. If the bootstrap dump fails, then the information regarding the failure is captured in the most recent Event Processor log.

If an incremental dump fails, then you may not see any new changes to the databases and tables in DAS. The incremental dump relies on events stored in the Hive metastore because these events take up a lot of space and are only used for replicating data. The events are removed from Hive metastore daily, which can affect DAS.

Fixing incremental dump failures

Required permission/role: You must be an admin user to complete this task.

If you see the message “Notification events are missing in the meta store”, then reset the Postgres database using the following command:

```
curl -H 'X-Requested-By: das' -H 'Cookie: JSESSIONID=<session id cookie>' http(s)://<hostname>:<port>/api/replicationDump/reset
```

where:

- session id cookie is the cookie value which you have to get for an admin user on the DAS UI, from the browser
- hostname is the DAS Webapp hostname
- port is the DAS Webapp port



Note: The error message "Notification events are missing in the meta store" is a symptom and not the cause of an issue, at most times. You can usually fix this by resetting the database. However, to determine the actual cause of the failure, you need to look at the first repl dump failure. Older Event Processor logs may contain information about the actual cause, but they are purged from the system at a set frequency. Therefore, if you hit this issue, we recommend that you first copy all the Event Processor logs that are available; else it might be too late to diagnose the actual issue.

If the first exception that you hit is an SQLException, then it is a Hive-side failure. Save the HiveServer and the Hive Metastore logs for the time when the exception occurred.

File a bug with Cloudera Support along with the above-mentioned logs.

Problem area: Reports page

Only the admin users can view the **Reports** page. If you are a non-admin user, then you lack the privileges to view the **Reports** page.

To enlist yourself as the admin user, go to Cloudera Manager Clusters DAS Configuration , and add your user name in the Admin Users field.

For example, hive,admin1,admin2. To make all users administrators, specify an asterik (*).

How DAS helps to debug Hive on Tez queries

Data Analytics Studio (DAS) provides a comprehensive framework to analyze and debug Hive and Hive on Tez queries. Using the query recommendations, Explain Plan, Directed Acyclic Graph (DAG) information, task attempts, and DAS logs, you can debug why a certain query failed to run or is taking longer than expected.

Understanding how Hive on Tez queries are executed

When you run a complex query from DAS, the query is executed on the Hive engine. In the first stage, a Directed Acyclic Graph (DAG) is created. It constitutes the time taken to compile, parse, and build the DAG for the next phase of the query execution. In the next stage of query execution, the DAG which is generated in Hive is submitted to the Tez engine for execution.

Tez is the default execution engine for Hive. When you run a complex query, Tez creates vertices (mappers and reducers) to provide results. The task and event details are written into the Tez logs. Tez writes the task and event details into a history file and Hue reads the history files to generate DAGs. Tez schedules tasks in such a way that the long-running tasks are run first and the shorter ones are scheduled to run later.

In the post-execution stage, the HDFS files are moved or renamed. The DAG Runtime shows the time taken by the Tez engine to execute the DAG.

What do you need to analyse and debug a query?

The crux of the query execution lies in the time taken by mappers and reducers and all tasks within them. When you are debugging a query to optimize its performance, you look into the vertex that took the longest time to run.

In a command-line mode, you can use the EXPLAIN command to view the dependency tree of a Tez job. The destinations are listed on the left and the sources are listed on the right, as shown in the following image:

```

DAG -> Dependency tree
Vertex dependency in root stage
|
| Map 1 <- Map 13 (BROADCAST_EDGE), Map 14 (BROADCAST_EDGE) |
| Map 11 <- Map 13 (BROADCAST_EDGE), Map 14 (BROADCAST_EDGE) |
| Map 3 <- Map 13 (BROADCAST_EDGE), Map 14 (BROADCAST_EDGE) |
| Map 9 <- Map 13 (BROADCAST_EDGE), Map 14 (BROADCAST_EDGE) |
| Reducer 10 <- Map 9 (SIMPLE_EDGE) |
| Reducer 12 <- Map 11 (SIMPLE_EDGE) |
| Reducer 2 <- Map 1 (SIMPLE_EDGE) |
| Reducer 4 <- Map 3 (SIMPLE_EDGE) |
| Reducer 5 <- Reducer 2 (CUSTOM_SIMPLE_EDGE), Reducer 4 (CUSTOM_SIMPLE_EDGE) |
| Reducer 6 <- Reducer 10 (CUSTOM_SIMPLE_EDGE), Reducer 5 (CUSTOM_SIMPLE_EDGE) |
| Reducer 7 <- Reducer 12 (CUSTOM_SIMPLE_EDGE), Reducer 6 (CUSTOM_SIMPLE_EDGE) |
| Reducer 8 <- Reducer 7 (SIMPLE_EDGE) |
|

```

There is only one connection between the source and the destination. Common types of edges are:

- Broadcast: sends the same data to every node
- Simple edge: sorts the data and sends it to the reducer
- Custom simple edge: sends data without sorting it

Using this information to troubleshoot Hive on Tez queries is a complex task requiring technical expertise. On older platforms, you could use Tez UI. On CDP, Cloudera recommends you to use DAS to debug queries. DAS provides a web interface to compose and run queries, shows intelligent query recommendations to optimize performance, query comparison, and other visual tools such as DAG info, DAG flow, visual explain plan, query timeline, and so on to debug queries.

Analyzing queries using DAS

After you run a query, go to the **Query Details** page. If there are any query recommendations on the top of the page, implement those in your query.

You can get the high-level information about the query from the Query Details section which you can use later:


Field	Description
Query ID	This is the Hive query ID
User	The user who ran the query
Status	The status of the query, whether it passed (success), failed, or is in progress (running)
Start time	The time when the query execution started
End time	The time when the query execution completed
Duration	Total time taken to execute the query
Tables read	Comma-separated list of Hive tables that were read while executing the query
Tables written	Comma-separated list of Hive tables to which the data was written while executing the query
Application ID	The YARN application ID, which can be used to investigate a query issue related to a YARN application
DAG ID	The DAG ID of the query that is run, which is also present in the application log
Session ID	The Tez session used through HiveServer2 (HS2) Application Master that you are working on
Thread ID	The thread that is used to process the request on HS2. You can correlate a query in a thread at the time when it is running on HS2.
Queue	The queue that is used for the Tez job. Currently, the Capacity Scheduler is the default scheduler. If no queues are configured, the default queue is used, which means that 100% of the cluster capacity is used to run Tez jobs. You can use the <code>tez.queue.name</code> property in the <code>hive-site.xml</code> file to specify the queue name.

To get information about the cluster ID and the HS2 instance to which you are connected, click the user icon on the top right corner of the page and select About.

By default, queries from the past 30 days are displayed on the **Queries** page. You can narrow down the queries run by a user within a specific time range by clicking the LAST 30 DAYS button on the search bar.

Next in the investigation, use the details from the Visual Explain plan as shown in the following image:

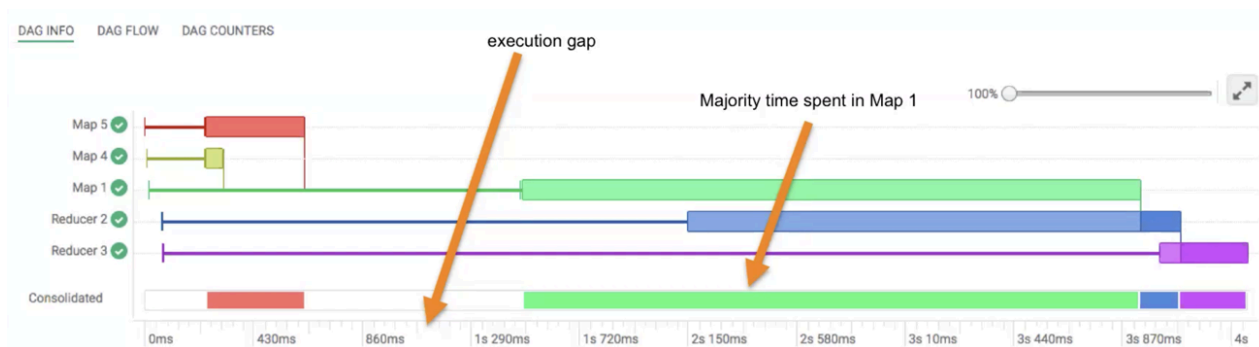
Visual Explain



The Visual Explain plan must be read from right to left. This is the DAG of how the query tasks are executed. You can get additional details that were captured while executing the query by clicking the boxes in the Visual Explain plan.

Next, check the Timeline section. It provides a break-up of the time taken to compile and parse the query, and build, submit, and run the DAG. This helps you to analyze query wait times. You can then go to the DAG INFO section to see which map or reducer has a dependency on other vertices and is consuming more time, and in turn resources.

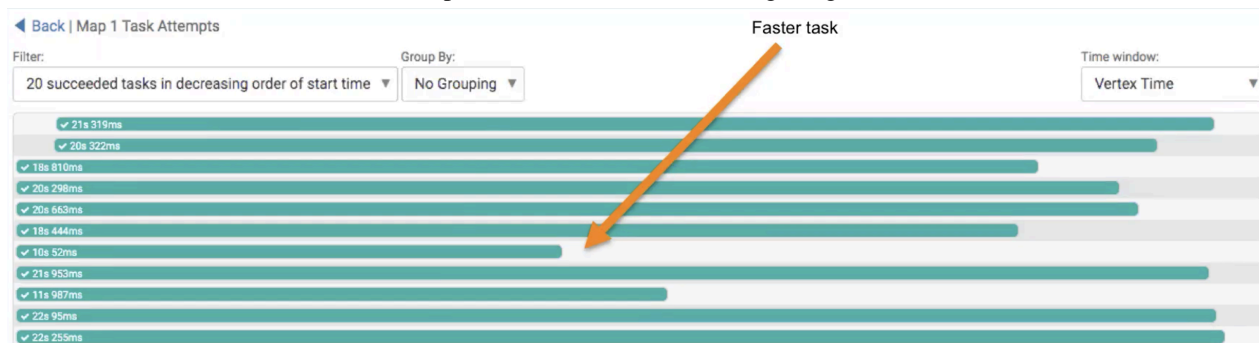
For example, in the following image, Map 1 consumes maximum time. Reducer 2 waits for an input from Map 1, and Reducer 3 depends on Reducer 2 for input:



The vertical lines connecting two vertices denote the dependency of one vertex on another. The lightly shaded part of the horizontal bar indicates that a vertex is waiting for an input from the other vertex. The darker part of the bar indicates the actual execution time for that task. The vertical lines indicate the time when the vertex initialized, the time when the vertex started, the time when the first task started, the time when the last task completed, and the time when the vertex finished its execution. When you hover your mouse over the vertical line, the bubble displays the stage of the vertex execution and provides a timestamp.

If you fix the white space (the execution gap) on the Consolidated bar, you can optimize this query to complete faster.

Each mapper and reducer has multiple sub-tasks known as task attempts. The Task Attempts section shows the number of times a certain task was attempted, as shown in the following image:



You can use this information to see what part of the query resulted in how many vertices and task attempts, such as the number of keys in the GROUP BY clause, or a JOIN condition.

To view the order in which the vertices are created, go to DAG FLOW. The DAG FLOW also shows which mapper or reducer is taking a long time or is not responding in case of long-running tasks or data skew.

To view the number of mapper and reducer counters for the job, the number of records accessed by the job, the number of tasks that were spun up, number of bytes read or written to HDFS or other storage platforms, and so on, go to the DAG COUNTERS tab. The numbers on the DAG COUNTERS may indicate any cartesian products or map side joins throwing OutOfMemory (OOM) errors, and so on.

To see the filtered view of Tez settings, go to the DAG CONFIGURATION tab.



Note: You can also compare two queries from the **Queries** page by clicking the Compare button under ACTIONS and selecting the second query that you want to compare. The query comparison report provides you a detailed side-by-side comparison of your queries, including recommendations for each query, metadata about the queries, visual explain for each query, query configuration, time taken at each stage of the query execution, and DAG information and DAG counters.

Accessing detailed logs of mapper and reducer tasks

DAS also allows you to download detailed logs of a task attempt. The log provides information about the container in which the task was running. To download the logs, click anywhere on the task attempt. On the **Details** pop-up, click Open log in new tab.