

Operations 1

Data Analytics Studio Operations

Date of Publish: 2018-11-16



<http://docs.hortonworks.com>

Contents

DAS overview.....	3
Manage queries.....	3
Search for queries.....	4
Refine Search using Filters.....	4
Save the search.....	4
Compose queries.....	4
Compare queries.....	5
Manage databases.....	5
Using the Database Explorer.....	5
Searching for tables.....	6
Managing tables.....	6
Creating tables.....	6
Uploading tables.....	7
Editing tables.....	7
Deleting tables.....	8
Managing columns.....	8
Managing partitions.....	8
Viewing storage information.....	8
Viewing detailed information.....	8
Viewing table and column statistics.....	9
Previewing tables using Data Preview.....	9
Manage reports.....	9
View Read and Write Reports.....	9
View Join Report.....	9
Running a query on a different Hive instance.....	10

DAS overview

Data Analytics Studio (DAS) helps you to perform operations on Hive tables and provides recommendations to optimize the performance of your queries.

DAS has two variants: DAS and DAS-Lite. DAS is the one that has all the features. The free version, called as DAS-Lite, comes with fewer features. The following features that are available in DAS are not available in DAS-Lite:

- Query optimization recommendation
- Table optimization recommendation
- Read and write report
- Join report

To upgrade from DAS-Lite to DAS, see *Data Analytics Studio Installation*.



Note: DAS and DAS-Lite have the same release versions, as well as the same [Support Matrix](#) for a version number. However, the binaries for DAS and DAS-Lite are different for a given version. Make sure that you download the appropriate binary before starting the installation.

Administrators can perform operations such as managing the databases.

You can perform the following operations using DAS:

- Search queries: You can search for queries executed on Hive tables in a database. You can further refine your search results using filters. Data Analytics Studio provides recommendations to optimize the performance of your queries on Hive tables. You can view the recommendations and edit your queries.



Note: The query optimization recommendations are not provided in the DAS-Lite version.

- Compose and execute queries: You can compose queries using the intuitive query composer. It has context based auto-complete to help editing a query faster. You can also view the visual explain of the query after executing it. You can save queries to view them later and edit them. You can edit the existing, saved queries and then save them as new queries. When you try to edit a query, you can use the query composer to easily create and execute your queries.
- Compare queries: You can compare two queries to know how each query is performing in terms of speed and cost effectiveness. DAS compares various aspects of the two queries, based on which you can identify what changed between the execution of those two queries, and you can also debug performance-related issues between different runs of the same query.
- Manage databases: Using the Database Explorer, you can manage existing databases by creating new tables, editing existing tables, and deleting tables. You can also create new database and add tables to it. You can manage existing tables by editing them to modify existing columns or add new columns. You can create new tables in DAS or upload existing tables available in CSV, JSON, and XML formats. You can edit columns in tables and also view suggestions for partitions and implement these recommendations.
- View reports: You can view which columns and tables are used for joins and make changes to the data layout to optimize the performance of the query with different search criteria.



Note: The view reports feature is not available in the DAS-Lite version.

Manage queries

Using the Data Analytics Studio, you can manage search for queries, compose them, and refine your search results based on various filters.

Data Analytics Studio uses a simple SQL query language, called HiveQL, that enables users familiar with SQL to query the data.

Search for queries

You can search for queries and see the list of queries that have been searched. You can refine your search results on the basis of parameters such as status of the query, queue to which the query belongs, the user of the query, tables read for the query, and tables written for the query, and execution modes.

Procedure

1. Enter your search query.
2. Click the time range drop down to select the from and to dates. You can also click to select one of the quick ranges provided in the list.
3. In the Refine tab, click the plus icon to further refine your search result. All options to refine your search appear. Select the required parameters and click **Apply**.
4. In the Actions column, click the pencil icon to open and edit the query in Composer view.
5. In the Actions column, click the i icon to view the query details.

Refine Search using Filters

You can further refine your query search results using various filters.

Procedure

1. Click + in the Refine tab.
2. Select the filters from the list by clicking each filter.
3. Click **Apply**.

Save the search

You can save your search for future use.

Click the save icon to save your search and the results.

Compose queries

You can write and edit queries using the Composer.

Procedure

1. Select a database from the left panel.

The **Tables** column displays all the tables in the database.

2. Search for the required tables, if needed.

Click on a table name to view the columns within the table. Use the **Filter** field to enter text to further refine your search for the required column.

3. Enter the query in the worksheet tab.

Each worksheet is identified by a unique name and you can add worksheets using the plus icon. The worksheet tab provides auto completion features. As you start entering the query, the worksheet suggests SQL commands, keywords, and table columns according to the query.

4. Perform the desired operation on the query:
 - Click **Execute** to run the query.
 - Click **Save As** to save the worksheet with a different name.
 - Click **Visual Explain** to view the query details in the form of a block diagram.
 - Select **Show Results** if you want to view the results of the query.
 - Select **Download Results** to download the query results.
 - Click **Saved** tab to view saved queries and edit them.

Compare queries

You can compare two queries to know how each query is performing in terms of speed and cost effectiveness. DAS compares various aspects of the two queries, based on which you can identify what changed between the executions of those two queries, and you can debug performance-related issues between different runs of the same query.

About this task

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 Directed Acyclic Graphs (DAG) information and DAG counters.

To compare two queries:

Procedure

1. Sign in to DAS and click **Queries**.
2. Under **ACTIONS**, click the **Compare** button to select a query.
The selected query is displayed on the comparison toolbar.
3. Next, select the query that you want to compare by clicking the **Compare** button.
The selected query is displayed on the comparison tool bar next to the query that you initially selected.
4. Click **COMPARE**.
The comparison report is displayed.
5. To remove any query from the comparison toolbar, click **x**.
6. To change the order of the queries that are being compared, click the **Swap queries** button on the comparison toolbar.

Manage databases

You can select the database on which you want to perform operations.

You can also create and drop a database, and search for tables within a database.

Using the Database Explorer

Using the Database Explorer, you can select a database to manage it, create a new database, or drop a database.

Once you select a database, the tables list gets refreshed to list all the tables from the selected database.

You can perform the following actions using the Database Explorer.

1. Select a database.
2. Create a database.

3. Drop a database.
4. Search for a table within a database.
5. Create a table.

Searching for tables

After selecting a database, the tables tab is refreshed to list all the tables in the database.

To search for a table, start entering the name of the table in the Search box. The list gets refreshed as you enter your search string to narrow the list down to your desired table.

Click the Refresh icon to refresh the list of tables.

Managing tables

Using the Database Explorer UI, you can view information about table columns, partitions, storage, and metadata, as well as preview a few rows from the table. This enables you to ensure that your table contains exactly the information that you intend it to.

For each table, you can view details from the following tabs:

- Columns

You can view the details of each column of the table. You can also search for columns using the Search box.

- Partitions

You can view the details of columns that are partitions in the table. You can also search from the list of columns.

- Storage Information

You can view the storage information such as input format, output format, if the table is compressed, number of buckets, buck columns, and so on.

- Detailed Information

You can view details such as the name of the database, the owner of the table, the created time, the last accessed time, table type, and so on.

- Statistics

You can view the table statistics and column statistics.

- Data Preview

You can preview a few rows from the table.

Creating tables

You can create a new table or upload an existing table to add it to the selected database.

Procedure

1. On the Database Explorer tab, click the + icon in the Tables section.

The Create Table screen appears.

2. Enter the name of the table.
3. Click **Add New Column** to add columns to the table.
4. For each column, specify the following details:
 - a) Enter the name of the column.
 - b) Select the data type of the column from the drop down list.

- c) Click **Advanced** to add advanced details of each column.
5. Click **Advanced** tab of the Create Table page.
 - a) Select Transactional if you want the table to be a transactional table.
 - b) Click Add Location if you want to specify a location for the table.
 - c) Select a file format from the drop down list of Add File Format section.
 - d) Click Add Row Format to specify details for the rows.
6. Click **Properties** to add key-value properties of the table.
7. Click **Create**.

Uploading tables

You can upload existing tables in CSV, JSON, or XML format to create new tables in the database. You can use the uploaded tables as other tables, run queries, and generate reports on them like other tables.

Procedure

1. Click **Upload Table**.
2. Enter the file format details in the Select File Format section.
 - a) File type - Select the file type from the drop down list.
 - b) Field Delimiter - This appears if you select CSV as the file format. Enter a delimiter character such as comma.
 - c) Escape Character - This appears if you select CSV as the file format. Enter an escape character such as \.
 - d) Quote Character - This appears if you select CSV as the file format. Enter a quote character.
 - e) If the first row is header, click to select Is first row header checkbox.
 - f) If the file contains endlines, click to select Contains endlines checkbox.
3. Select the source of the table file.
 - a) **Upload from HDFS** - Select this option to upload files from HDFS.
 - b) **Upload from Local** - Select this option to upload files from the local node.The table details automatically appear as preview. Click Preview if the table details do not load automatically. Data Analytics Studio updates the name of the table, names of the columns, and the data types of each column.
4. Modify the names of the table and columns as needed.
5. Click the **Advanced** tab to specify any additional options:
 - a) Select **Transactional** if you want the table to be a transactional table.
 - b) Click **Add Location** if you want to specify a location for the table.
 - c) Select a file format from **Add File Format**.
 - d) Click **Add Row Format** to specify details for the rows.
6. Click **Properties** to add key-value properties of the table.
7. Click **Create**.

Editing tables

You can edit tables to add new columns, update existing columns, and make changes to the properties of the table.

Procedure

1. On the Database Explorer tab, select the database.
2. From the list of tables, select the table that you need to edit.
3. From the Table view, click the Actions menu.
4. Select **Edit**.
5. Edit the details of existing columns.

6. Click **Add New Column** to add columns to the table.
7. For each column, specify the following details:
 - a) Enter the name of the column.
 - b) Select the data type of the column from the drop down list.
 - c) Click **Advanced** to add advanced details of each column.
8. Click **Advanced** tab of the Create Table page.
 - a) Select **Transactional** if you want the table to be a transactional table.
 - b) Click **Add Location** if you want to specify a location for the table.
 - c) Select a file format from the drop down list of Add File Format section.
 - d) Click **Add Row Format** to specify details for the rows.
9. Click **Properties** to add or edit key-value properties of the table.
10. Click **Edit** to save the changes you made to the table.

Deleting tables

You can delete a table using the Database Explorer.

Procedure

1. On the Database Explorer tab, select the database.
2. From the list of tables, select the table that you need to delete.
3. From the Table view, click the **Actions** menu.
4. Click **Delete**.
5. In the dialog box that appears, click **Confirm**.

Managing columns

Using the Columns tab, you can view the details of each column of the table.

You can also search for columns using the Search box.

To search for a column, enter the name of the column in the search box and click **Regex Search**.

Managing partitions

On the Partitions tab, you can view the details of columns that are partitions in the table.

You can also search from the list of columns.

To search for a column that is a partition, enter the name of the column in the search box and click **Regex Search**.

Viewing storage information

Using the Storage Information tab, you can view the storage information of the table such as input format, output format, if the table is compressed, number of buckets, number of columns, and more details.

To search for particular details, enter the name in the search box and click **Regex Search**.

Viewing detailed information

Using the Detailed Information tab, you can view details such as the name of the database, the owner of the table, the created time, the last accessed time, table type, and more details.

To search for particular details, enter the name in the search box and click **Regex Search**.

Viewing table and column statistics

On the Statistics tab, you can view table statistics and column statistics.

To view the statistics of each column, click **Show**. The following details are displayed:

- Min
- Max
- Number of Nulls
- Distinct Count
- Average Column Length
- Max Column Length
- Number of True
- Number of False

If a message appears that the statistics might be stale, select **Include Columns** and click **Recompute** at the top of the page.

Previewing tables using Data Preview

On the Data Preview tab, you can preview a few rows from the table.

Manage reports

As a database administrator, you can view which columns and tables are used for joins and make changes to the data layout to optimize the performance of the query.



Note: The view reports feature is not available in the DAS-Lite version.

View Read and Write Reports

You can see a list of tables belonging to the selected database along with the read count, write count, and other details for each table.

Once you click the **Read/Write Reports**, the Relations page appears.

On the right side of the page, you can also view these details in the form of an entity relationship diagram.

For each table, select one of the following counts from the drop down list of counts:

- Projection Count
- Aggregation Count
- Filter Count
- Join Count

You can refine your search on the basis of timeline by selecting the time when the read and write operations have been performed.

View Join Report

You can view the details of joins between all the tables in a database.

On the Reports page, click the **Join Report** icon to view the Join reports for each database. You can change the database from the drop down list to view the join diagram for that database.

You can further refine your join diagram using join algorithms:

- Join
- Hash Join
- Merge Join
- Lateral View Join
- Sorted Merge

The diagram will be updated on the basis of your selection.

Running a query on a different Hive instance

By default, ZooKeeper redirects your query to run on a Hive instance that is available at that time. However, if you have optimized a certain Hive instance, and you want to run your workload on that particular Hive instance, then you can switch to that instance by specifying the Hive Metastore host name and the port in the JDBC connection field on the DAS user interface.

About this task

The JDBC connection string is defined by the `hive.zookeeper.quorum` property under **Ambari > Services > Hive > CONFIGS > ADVANCED > Advanced hive-site**. It also resides in following DAS configuration file: `etc/das/conf/das-hive-site.conf`.



Note: The JDBC connection string that you specify through the web interface is a temporary configuration and is limited to that particular browser session.

To update the JDBC connection string:

Procedure

1. Sign in to the DAS portal.
2. Click the profile icon on the top right corner and then click **About**.
3. On the **Configurations** window, click **Edit** under the **JDBC connection** field.
4. Specify the Hive Metastore host name and port that you want to connect to in the following format:

```
jdbc:hive2//hive-metastore-host:port/
```



Note: The Hive instance that you want to switch to should be within the same cluster.

5. To save the changes, click **Update**.
6. To revert any changes that you have made to the original connection string, click **Reset**.
7. To go back to the **Configurations** window, click **Cancel**.
8. To exit, click **OK**.

Results

When you run your query after updating the JDBC connection string, the query is executed on that particular Hive instance instead of being redirected to any Hive instance by ZooKeeper.