Writing query results

This document describes how to write or save query results.

Temporary and permanent tables

BigQuery saves all query results to a table, which can be either permanent or temporary.

- BigQuery uses temporary tables to cache query results (/bigquery/docs/cached-results) that aren't written to a permanent table. The tables are created in a special dataset and named randomly. You can also create temporary tables for your own use. For further information, see Temporary tables (/bigquery/docs/reference/standard-sql/data-definition-language#temporary_tables).

  After a query finishes, the temporary table exists for up to 24 hours. To view table structure and data, go to the BigQuery console (https://console.cloud.google.com/bigquery), click **Query history**, and choose the query that created the temporary table. Then, in the **Destination table** row, click **Temporary table**.

  You cannot query or share temporary tables, and they are not visible using any of the standard list or other table manipulation methods. You are not charged for storing temporary tables.

- A permanent table can be a new or existing table in any dataset to which you have access. If you write query results to a new table, you are charged for storing (/bigquery/pricing#storage) the data. When you write query results to a permanent table, the tables you're querying must be in the same location as the dataset that contains the destination table.

Required permissions

At a minimum, to write query results to a table, you must be granted the following permissions:

- `bigquery.tables.create` permissions to create a new table
- `bigquery.tables.updateData` to write data to a new table, overwrite a table, or append data to a table

- `bigquery.jobs.create` to run a query job

Additional permissions such as `bigquery.tables.getData` may be required to access the data you're querying.

The following predefined IAM roles include both `bigquery.tables.create` and `bigquery.tables.updateData` permissions:

- `bigquery.dataEditor`
- `bigquery.dataOwner`
- `bigquery.admin`

The following predefined IAM roles include `bigquery.jobs.create` permissions:

- `bigquery.user`
- `bigquery.jobUser`
- `bigquery.admin`

In addition, if a user has `bigquery.datasets.create` permissions, when that user creates a dataset, they are granted `bigquery.dataOwner` access to it. `bigquery.dataOwner` access gives the user the ability to create and update tables in the dataset.

For more information on IAM roles and permissions in BigQuery, see Predefined roles and permissions (/bigquery/docs/access-control).

Writing query results to a permanent table

When you write query results to a permanent table, you can create a new table, append the results to an existing table, or overwrite an existing table. You can write query results to a permanent table by:

- Using the Cloud Console or the classic BigQuery web UI
- Using the command-line tool's `bq query` command
- Calling the `jobs.insert` (/bigquery/docs/reference/rest/v2/jobs/insert) API method and configuring a `query` (/bigquery/docs/reference/rest/v2/jobs#configuration.query) job

- Using the client libraries

Writing query results

Use the following procedure to write your query results to a permanent table. To help control costs, you can preview data (/bigquery/docs/best-practices-costs#preview-data) before running the query.

ConsoleDDL (#ddl)Classic UI (#classic-ui)bq (#bq)API (#api)Go (#go)Java (#java)Node.js (#node.js)Py

1. Open the BigQuery web UI in the Cloud Console.

   Go to the Cloud Console (https://console.cloud.google.com/bigquery)

2. In the navigation panel, in the **Resources** section, expand your project and select a dataset.
3. If the query editor is hidden, click **Show editor** at the top right of the window.
4. Enter a valid SQL query in the **Query editor** text area.
5. Click **More** below the editor, then select **Query settings**.

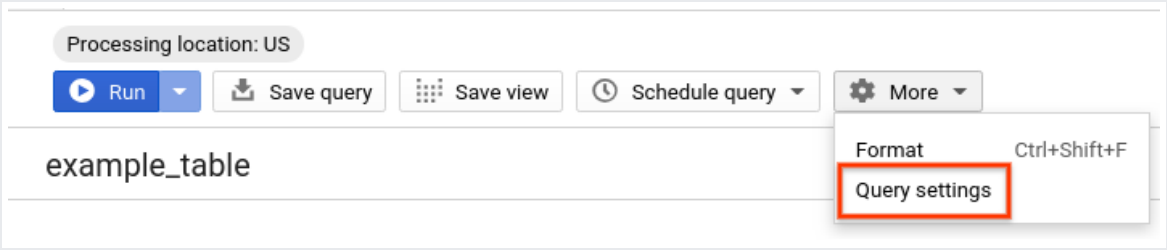

6. Check the box to **Set a destination table for query results**.

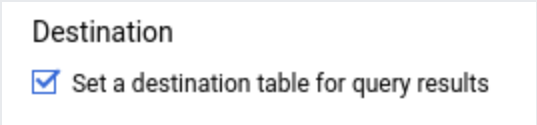

7. In the **Destination** section, select the appropriate **Project name** and **Dataset name** where the table will be created, and choose a **Table name**.
8. In the **Destination table write preference** section, choose one of the following:
   - **Write if empty** — Writes the query results to the table only if the table is empty.
   - **Append to table** — Appends the query results to an existing table.
   - **Overwrite table** — Overwrites an existing table with the same name using the query results.

9. (Optional) For **Processing location**, click **Auto-select** and choose your location (/bigquery/docs/locations).
10. Click **Run query**. This creates a query job that writes the query results to the table you specified.

Alternatively, if you forget to specify a destination table before running your query, you can copy the cached results table to a permanent table by clicking the **Save Results** (#save-query-results) button below the editor.

Writing large query results

Normally, queries have a maximum response size (/bigquery/quotas#query_jobs). If you plan to run a query that might return larger results, you can:

- In standard SQL, specify a destination table for the query results.
- In legacy SQL, specify a destination table and set the `allowLargeResults` option.

When you specify a destination table for large query results, you are charged for storing (/bigquery/pricing#storage) the data.

Limitations

In legacy SQL, writing large results is subject to these limitations:

- You must specify a destination table.
- You cannot specify a top-level `ORDER BY`, `TOP` or `LIMIT` clause. Doing so negates the benefit of using `allowLargeResults`, because the query output can no longer be computed in parallel.
- Window functions (/bigquery/query-reference#windowfunctions) can return large query results only if used in conjunction with a `PARTITION BY` clause.

Writing large results using legacy SQL

To write large result sets using legacy SQL:

1. Open the BigQuery web UI in the Cloud Console.
   Go to the Cloud Console (https://console.cloud.google.com/bigquery)
2. Click **Compose new query**.
3. Enter a valid SQL query in the **Query editor** text area. Use the `#legacySQL` prefix or be sure you have **Use Legacy SQL** checked in the query settings.
4. Click **More** then select **Query settings**.

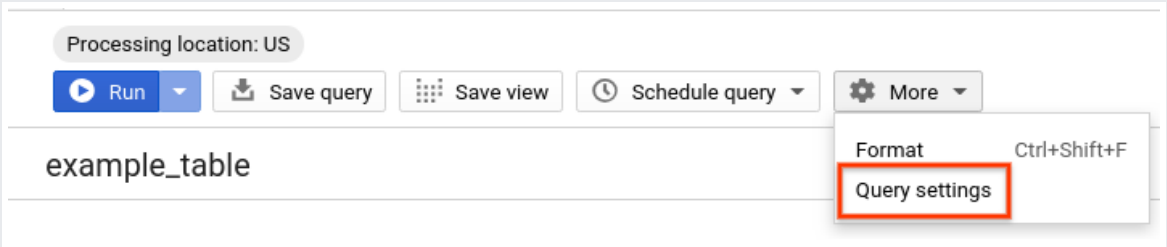

5. For **Destination**, check **Set a destination table for query results**.

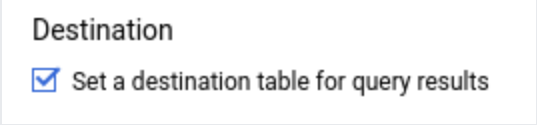

6. For **Project name**, choose the project where the destination table will be created.
7. For **Dataset name**, choose the dataset that will store the table.
8. In the **Table name** field, enter a table name.
9. If you are writing a large results set to an existing table, you can use the **Destination table write preference** options to control the write disposition of the destination table:
   - **Write if empty:** Writes the query results to the table only if the table is empty.
   - **Append to table:** Appends the query results to an existing table.
   - **Overwrite table:** Overwrites an existing table with the same name using the query results.

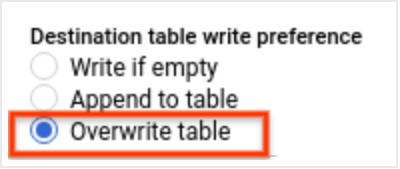

10. For **Results Size**, check **Allow Large Results (no size limit)**.

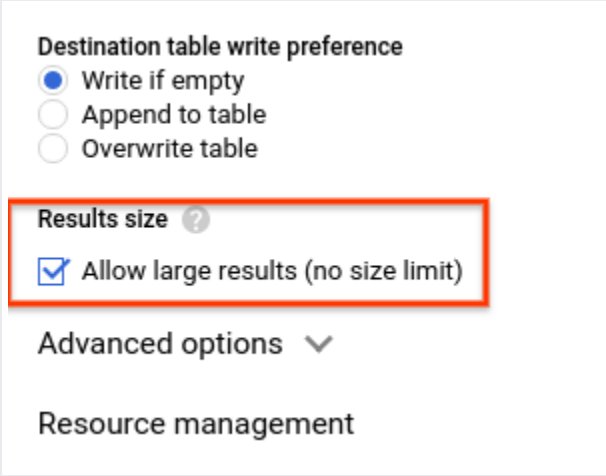

11. (Optional) For **Processing Location**, click **Auto-select** and choose your data's location (/bigquery/docs/locations).

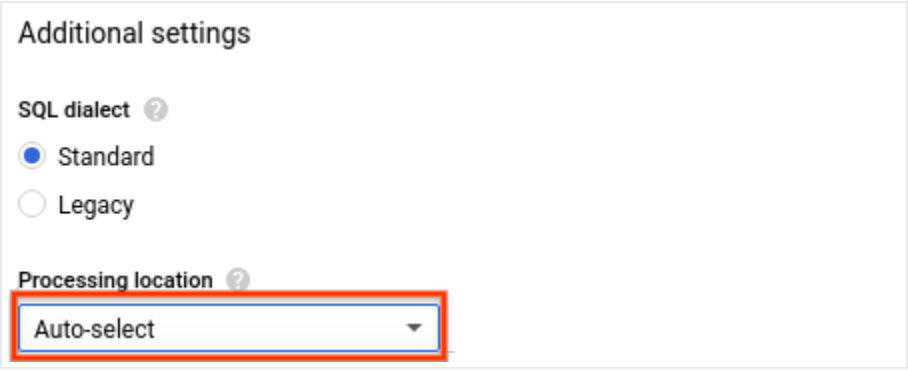

12. Click **Save** to update the query settings.
13. Click **Run**. This creates a query job that writes the large results set to the table you specified.

Downloading and saving query results from a UI

After you run a SQL query by using either classic BigQuery web UI or the Cloud Console, you can save the results to another location. You can download query results to a local file or to Sheets by using either Cloud Console or the classic BigQuery web UI, or you can download the results to Drive by using the Cloud Console. Saving results to a local file, Sheets, or Drive is not supported by the command-line tool or the API.

Limitations

Downloading and saving query results are subject to the following limitations:

- To download query results using the classic BigQuery web UI, the results set must contain fewer than 16,000 rows, and it must be 10 MB or less. If your results are larger than 10 MB or 16,000 rows you can save them to a table instead (#writing_query_results).
- You can download query results locally only in CSV or newline-delimited JSON format.
- You cannot download query results containing nested and repeated data in CSV format.
- You cannot save query results containing nested and repeated data to Sheets.
- When you save query results to Sheets using the classic BigQuery web UI, the results set must contain fewer than 16,000 rows, and it must be 10 MB or less. If your results are larger than 10 MB or 16,000 rows you can save them to a table instead (#writing_query_results).
- To save query results to Drive using the Cloud Console, the results set must be 1 GB or less. If your results are larger than 1 GB, you can save them to a table instead.
- You can save query results to Drive only in CSV or newline-delimited JSON format.

Downloading query results to a local file

Downloading query results to a local file is not supported by the command-line tool or the API.

To download query results as a CSV or newline-delimited JSON file by using the web UI:

Console Classic UI (#classic-ui)

1. Open the BigQuery web UI in the Cloud Console.
   Go to the Cloud Console (https://console.cloud.google.com/bigquery)
2. Click **Compose new query**.
3. Enter a valid SQL query in the **Query editor** text area.
4. (Optional) To change the processing location, click **More** and select **Query settings**. Under **Processing location**, click **Auto-select** and choose your data's location (/bigquery/docs/locations).
5. Click **Run**.
6. When the results are returned, click the **Save results** and select the format/location where you want to save the results.

The file is downloaded to your browser's default download location.

Saving query results to Drive

ature is covered by the Pre-GA Offerings Terms (/terms/service-terms#1) of the Google Cloud Platform Terms
e. Pre-GA features may have limited support, and changes to pre-GA features may not be compatible with oth
sions. For more information, see the launch stage descriptions (/products#product-launch-stages).

Saving query results to Drive is not supported by the command-line tool, the API, or the classic BigQuery web UI.

To save query results to Drive by using the Cloud Console:

Console

1. Open the BigQuery web UI in the Cloud Console.

   Go to the Cloud Console (https://console.cloud.google.com/bigquery)

2. Enter a valid SQL query in the **Query editor** text area.
3. Click **Run**.
4. When the results are returned, click **Save Results**.

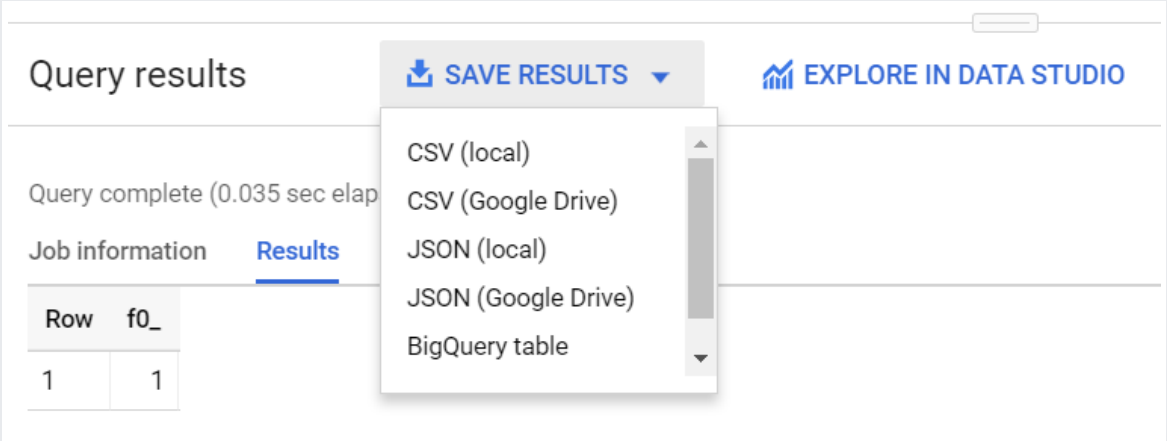

5. Select **CSV (Google Drive)** or **JSON (Google Drive)**. When you save results to Drive, you cannot choose the location. Results are always saved to the root "My Drive" location.

6. It may take a few minutes to save the results to Drive. When the results are saved, you receive a pop-up message that includes the filename — `bq-results-[TIMESTAMP]-[RANDOM_CHARACTERS].[CSV or JSON]`.

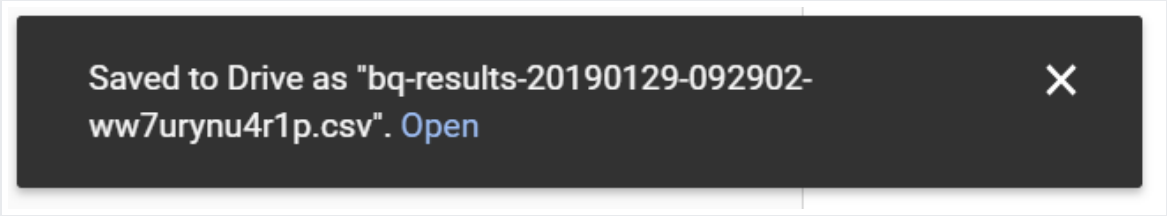

7. In the pop-up message, click **Open** to open the file, or navigate to Drive and click **My Drive**.

Saving query results to Sheets

Saving query results to Sheets is not supported by the command-line tool or the API.

To save query results to Sheets by using the web UI:

ConsoleClassic UI (#classic-ui)

1. Open the BigQuery web UI in the Cloud Console.
   Go to the BigQuery web UI (https://console.cloud.google.com/bigquery)
2. Click **Compose new query**.
3. Enter a valid SQL query in the **Query editor** text area.
4. (Optional) To change the processing location, click **More** and select **Query settings**. Under **Processing location**, click **Auto-select** and choose your data's location (/bigquery/docs/locations).
5. Click **Run**.
6. When the results are returned, click the **Save results** and select **Google Sheets**.
7. If necessary, follow the prompts to log into your Google Account and click **Allow** to give BigQuery permission to write the data to your Drive `MY Drive` folder.

   After following the prompts, you should receive an email with the subject "BigQuery Client Tools connected to your Google Account". The email contains information on the permissions you granted along with steps to remove the permissions.
8. When the results are saved, a message similar to the following appears below the query results in the console BigQuery web UI: `Saved to Sheets as "results-20190225-103531.`

`Open`. Click the link in the message to view your results in Sheets, or navigate to your `My Drive` folder and open the file manually.

When you save query results to Sheets, the filename begins with `results-[DATE]` where `[DATE]` is today's date in the format `YYYYMMDD`.

Note: Saving results to Sheets is not supported by the command-line tool or the API.

Last updated 2020-08-10 UTC.