

Using Snowflake

Contents:

- *Limitations*
 - *Uses of Snowflake*
 - *Before You Begin Using Snowflake*
 - *Secure Access*
 - *Storing Data in Snowflake*
 - *Reading from Snowflake*
 - *Writing to Snowflake*
-

This section describes how you interact through Trifacta® Self-Managed Enterprise Edition with your Snowflake data warehouse.

- Snowflake is a scalable data warehouse solution. For more information, see <https://www.snowflake.com>.

Limitations

- To ingest data from a Snowflake table, one of the following must be enabled:
 - A named stage must be created for the table. For more information, see the Snowflake documentation.
 - Snowflake must be permitted to create a temporary stage, which requires:
 - Write permissions on the table's database, and
 - A schema named PUBLIC must exist and be accessible.
- No schema validation is performed as part of writing results to Snowflake.
- Credentials and permissions are not validated when you are modifying the destination for a publishing job.
- For Snowflake, no validation is performed to determine if the target is a view and is therefore not a supported target.

Uses of Snowflake

The Trifacta platform can use Snowflake for the following tasks:

1. Create datasets by reading from Snowflake tables.
2. Write to Snowflake tables with your job results.

Before You Begin Using Snowflake

- **Enable S3 Sources:** Snowflake integration requires the following:
 - Installation of the product on a customer-managed AWS infrastructure.
 - S3 is set to the base storage layer.
 - For more information, see *Enable Snowflake Connections*.
- **Read Access:** Your Snowflake administrator must configure read permissions. Your administrator should provide a database for upload to your Snowflake data warehouse.
 - **Read-only Access:** If you are creating a read-only connection to Snowflake, you must provide a database for staging. The accessing user must have write permission to the specified database. See *Create Snowflake Connections*.
- **Write Access:** You can write and publish jobs results to Snowflake.

Secure Access

SSL is the default connection method.

Storing Data in Snowflake

Your Snowflake administrator should provide database access for storing datasets. Users should know where shared data is located and where personal data can be saved without interfering with or confusing other users.

NOTE: Trifacta Self-Managed Enterprise Edition does not modify source data in Snowflake. Datasets sourced from Snowflake are read without modification from their source locations.

Reading from Snowflake

You can create a Trifacta dataset from a table stored in Snowflake.

NOTE: The Snowflake cluster must be in the same region as the default S3 bucket.

For more information, see *Snowflake Browser*.

Writing to Snowflake

You can write back data to Snowflake using one of the following methods:

- Job results can be written directly to Snowflake as part of the normal job execution. Create a new publishing action to write to Snowflake. See *Run Job Page*.
- For more information on how data is converted to Snowflake, see *Snowflake Data Type Conversions*.

Data Validation issues:

- No validation is performed for the connection and any required permissions during job execution. So, you can be permitted to launch your job even if you do not have sufficient connectivity or permissions to access the data. The corresponding publish job fails at runtime.
- Prior to publication, no validation is performed on whether a target is a table or a view, so the job that was launched fails at runtime.