

Enable Snowflake Connections

Contents:

- *Limitations*
 - *Pre-requisites*
 - *General*
 - *AWS permissions*
 - *OAuth2 requirements*
 - *Enable*
 - *Create Stage*
 - *Create Snowflake Connection*
 - *Testing*
-

This section describes how to integrate the Trifacta® platform with Snowflake databases.

- Snowflake provides a cloud-database datawarehouse designed for big data processing and analytics. For more information, see <https://www.snowflake.com>.
- For more information on supported versions, see *Connection Types*.

Limitations

NOTE: This integration is supported only for deployments of Trifacta Self-Managed Enterprise Edition in customer-managed AWS infrastructures. These deployments must use S3 as the base storage layer. For more information, see *Supported Deployment Scenarios for AWS*.

- SSO connections are not supported.

Pre-requisites

General

- If you do not provide a stage database, then the Trifacta platform must create one for you under the `PUBLIC` schema in the default database.
 - In this default database, you must include a schema named `PUBLIC`.
 - For more information, please see the Snowflake documentation.
- The user-created stage must point to the same S3 bucket that is the default S3 bucket in use by Trifacta Self-Managed Enterprise Edition.

AWS permissions

Your Snowflake stage and other databases must be permitted to use S3 resources for your users.

NOTE: If users in your deployment are using IAM roles in user mode for AWS access, then the Snowflake stage must have permissions to write to the user's S3 bucket.

For more information, see *Required AWS Account Permissions*.

OAuth2 requirements

If you are integrating with Snowflake using OAuth2, additional configuration is required:

- OAuth2 must be enabled in the product. For more information, see *Enable OAuth 2.0 Authentication*.
- You must create a client app and client to manage authentication between the Trifacta application and your Snowflake deployment. For more information, see *OAuth 2.0 for Snowflake*.

Enable

When relational connections are enabled, this connection type is automatically available. For more information, see *Enable Relational Connections*.

Create Stage

In Snowflake terminology, a **stage** is a database object that points to an external location on S3. This stage must contain access credentials.

- If a stage is used, it should be in the default bucket used on S3 for storage.

NOTE: For read-only connections to Snowflake, you must specify a Database for Stage. The connecting user must have write access to this database.

Tip: You can specify a separate database to use for your stage.

- If a stage is not specified, a temporary stage is created using the current user's AWS credentials. Please verify the Pre-requisites again.

NOTE: Without a defined stage, you must have write permissions to the database from which you import. This database is used to create the temporary stage.

For more information on stages, see <https://docs.snowflake.net/manuals/sql-reference/sql/create-stage.html>.

In Trifacta Self-Managed Enterprise Edition, the stage location is specified as part of creating the Snowflake connection.

Create Snowflake Connection

For more information, see *Create Snowflake Connections*.

Testing

Steps:

1. After you create your connection, load a small dataset based on a table in the connected Snowflake database.

NOTE: For Snowflake connections, you must have write access to the database from which you are importing.

See *Import Data Page*.

2. Perform a few simple transformations to the data. Run the job. See *Transformer Page*.
3. Verify the results.

For more information, see *Verify Operations*.