

Adding Snowflake Data Warehouses

A Snowflake data warehouse is a cloud-based data warehouse that contains the tables and views that you want to access as cube facts and dimensions. It also contains aggregate-table instances that either you or the AtScale engine creates in a schema that you specify.

Before You Begin

- ▶ Ensure that your user ID in the Design Center is assigned the Super User role or is assigned the Manage Data Warehouses role permission.
- ▶ **Aggregate Schema:** Ensure that you know the schema to use for aggregate tables to be built in the data warehouse.
AtScale reads and writes aggregate data to this schema. The AtScale service account user must have ALL privileges for this schema. BI tool user accounts should not have the select permission for this schema.

Procedure

1. Go to **Settings > Data Warehouses** and click **Create Data Warehouse**.
2. In the Add a Data Warehouse dialog box, under **Type of Data Warehouse**, select **Snowflake**.
3. Enter a unique name for the data warehouse. AtScale displays this name in the Design Center and uses it in log files.
4. (Optional) Enter the name of the **External Connection ID**. This input will be inherited by the **Name** value, but can be overridden by enabling the **Override generated value** slider.
5. Enter the name of the **Database** for AtScale to use when creating aggregate tables. (You must create this database; AtScale does not automatically create a database if one does not exist.)
6. Enter the name of the schema for AtScale to use when creating aggregate tables. (You must create a schema; AtScale does not create a schema if one does not exist.)
7. Specify the Data Warehouse as a **Read-only source**.
 - ▶ Select this option if AtScale will be configured to access this database with credentials that do not have permission to create new tables, or if you need to prevent AtScale from placing any aggregate tables in this data warehouse.
 - ▶ If **Read-only source** is enabled upon editing an existing data warehouse with aggregates present, all existing aggregates stored in this data warehouse will be deactivated.

Add A Snowflake Connection

1. Enter a unique name for the connection.
2. Enter the name of the Snowflake Data Warehouse.
3. For **Host**, enter the full name of your Snowflake account.

4. For **Database**, enter the name of the default database to use after connecting. This value was entered on the Data Warehouse creation above, and should default to the database set there.
5. For **Username** and **Password**, enter the user name and password that AtScale should use when connecting.
If you have [external secret manager](#) enabled and configured, enter the corresponding credentials.
6. If using **Key Pair Authentication**, select Key Pair, Encrypted Key Pair, Key File or Encrypted Key File and then enter the **Credential File Location**. Then enter private key, or file containing the private key.
7. Enter the name of the Snowflake Data Warehouse.
8. (Optional) Specify any extra JDBC flags only if you are certain that they are required.
9. Test the connection.
10. Click **Save** to complete the setup.

Query Roles

If multiple Snowflake connections are configured, you can toggle which types of queries (Large Interactive, Small Interactive or System) are permitted per Snowflake connection. A query role can be assigned to only one connection at a time.

Expand the drop-down for the Snowflake data warehouse to choose the allowed query roles for each connection:

- ▲ **Large Interactive Queries:** Use this engine for large user queries. These are usually BI queries involving full scans of the raw fact table or ones that involve high-cardinality dimensions.
- ▲ **Small Interactive Queries:** Use this engine for small user queries. These are usually BI queries that can be optimized by AtScale aggregates or ones that involve low-cardinality dimensions.
- ▲ **System Queries:** Use this engine for AtScale system processing, such as building aggregate tables or querying system statistics. The System query role user must not be subject to resource limitation rules that would prevent it from building aggregate tables.

Aggregates

By default AtScale creates aggregate tables with quoted lowercase table names on Snowflake. To configure AtScale to create aggregate table names in uppercase on Snowflake do the following:

1. As an AtScale System Administrator, log in to Design Center and navigate to the Engine Settings page by clicking: SETTINGS > ENGINE.
2. Search page for `aggregate.snowflake.table.names.uppercase`
3. Set value to True
4. Click Save.
5. Restart the Engine for the setting to take effect.