

# 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 model facts and dimensions. It also contains aggregate-table instances that either you or the AtScale engine creates in a schema that you specify.

## Prerequisites

To add data warehouses, you must be logged in to AtScale as an admin user.

Additionally, you must know the schema to use for building aggregate tables in the data warehouse. AtScale reads and writes aggregate data to this schema, so the AtScale service account user must have ALL privileges for it. BI tool user accounts should not have the select permission.

## Procedure

1. In Design Center, open the main menu and select **Settings**. The **Settings** page opens.
2. In the **Settings** panel, click **Data Warehouses**.
3. In the **Add Data Warehouse** section, click the **Snowflake** icon. The **Add Data Warehouse** panel opens.
4. Enter a unique **Name** for the data warehouse. AtScale displays this name in the Design Center and uses it in log files.
5. (Optional) Enter the name of the **External connection ID**. This field inherits the **Name** value but can be overridden by enabling the **Override generated value?** option.
6. Enter the name of the **Aggregate Database**, which AtScale uses when creating aggregate tables. You must create this; AtScale does not create one automatically.
7. Enter the name of the **Aggregate Schema** for AtScale to use when creating aggregate tables. You must create this schema; AtScale does not create one automatically.
8. (Optional) Specify whether the data warehouse is 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.
9. (Optional) Enable impersonation for the data warehouse by selecting the **Enable to impersonate the client when connecting to data warehouses** checkbox. When enabled, you can also enable the following options:

- ▲ **Always use Canary Queries:** Determines whether canary queries are required in the event of aggregate misses.
- ▲ **Allow Partial Aggregate Usage:** Enables mixed aggregate and raw data queries.

10. Click **Apply**.

## Add A Snowflake Connection

Once you have added your Snowflake data warehouse, you need to add a connection.

To add a Snowflake connection:

1. On the **Data Warehouses** page, click the expand icon for your data warehouse.
2. In the **Add Connection** field, click the plus icon. The **Add Snowflake Connection** panel opens.
3. Enter the following fields:
  - ▲ **Name:** A unique name for the connection.
  - ▲ **Snowflake data warehouse name:** The full name of the data warehouse.
  - ▲ **Host:** The full name of your Snowflake account.
  - ▲ **Database:** The name of the default database to use after connecting. This defaults to the value set when adding the data warehouse and cannot be changed.
  - ▲ **Username:** The username AtScale uses when connecting to the data warehouse.
4. In the **Authentication Method** field, select the authentication method to use when connecting to the data warehouse, then complete the credential fields as needed.
5. (Optional) Specify any extra JDBC flags only if you are certain that they are required.
6. Click **Test connection** to test the connection.
7. Click **Apply**.

## Aggregates

By default, AtScale creates aggregate tables on Snowflake with quoted lowercase table names. To configure AtScale to create aggregate table names in uppercase, set the `aggregate.snowflake.table.names.uppercase` engine setting to `True`.

Restart the AtScale engine for your changes to take effect. You can do this through your cluster.

For more information on changing engine settings, see [Changing Engine Settings](#).

## Next Steps

After configuring your data warehouses and connections, you can configure which types of queries are permitted for each connection. For more information, see [Configuring Query Mapping](#).