

Adding Azure Synapse Warehouses

An Azure Synapse data warehouse contains the tables and views that you want to import as fact datasets and dimension datasets into AtScale cubes. It also contains aggregate-table instances that either you or the AtScale engine creates in a database 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.
- ▶ Ensure that you know the database to use for aggregate tables to be built in the data warehouse. You should create this database if it is not available; AtScale does not automatically create a database if one does not exist.
- ▶ Ensure that you know the schema to use for aggregate tables to be built in the data warehouse. AtScale creates tables, drops tables, and reads and writes aggregate data to this schema. You should create this schema if it is not available.
- ▶ To be able to load data into Synapse through AtScale you need to create `AtscaleAzureDataLakeStore` external data source, with credentials of the user that sets up the data warehouse in AtScale. For details, see [Use external tables with Synapse SQL](#).

Consider that AtScale supports ADLS Gen 2. For details, see [Azure Data Lake Storage Gen2](#).

Adding Azure Synapse Data Warehouse

1. Go to **Settings > Data Warehouses** and choose **Create Data Warehouse** to open the Add a Data Warehouse dialog box.
2. Select **Azure Synapse Analytics SQL** in **Type of Data Warehouse** .
3. Enter a unique name for the data warehouse. AtScale displays this name in the Design Center and uses it in log files.
4. 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.
6. Enter the name of the **Aggregate Schema** for AtScale to use when creating aggregate tables.
7. Choose whether to set the data warehouse as read-only:
 - ▶ Select the **Read-only source** 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.

8. If you need to use objects in a file system:
 1. Go to the **Configure Filesystem** section and enable the **Enable Filesystem** option.
 2. In the **Filesystem Type** field, select **Azure DLS Gen2 Filesystem**, then enter the following details:
 - ▲ URI
 - ▲ Storage account name
 - ▲ Container name
 - ▲ SAS token
 3. Optionally, test the file system connection.
9. Click **Save**.

Add An Azure Synapse Connection

After setting up your data warehouse, you must add one or more connections to it before you can run queries:

1. Expand the Azure Synapse data warehouse, and select **Create Connection**.
2. Enter the **Name** of your Azure Synapse connection.
3. Enter the **Host** location of the Azure Synapse server.
4. For **Port**, enter the port number to route transmitted data. The default port number is 443.
5. Enter the **Database**.
6. If necessary, enter the JDBC flags that are needed for the connection in the **Extra JDBC Flags** field.

To use Active Directory authentication, enter `authentication=ActiveDirectoryPassword`.

7. In the **Authorization** section, choose **Password** and enter the Username and Password.
8. Optionally, test the connection.
9. Click **Save**.