

Adding DB2 Data Warehouses

A Db2 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.



To configure a Db2 data warehouse, you must have a special AtScale license. Ask your sales representative about obtaining a Db2-enabled license.

Db2 Limitations

The following features are not supported with Db2 platforms:

- ▶ Quantile / NTile Calculations

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.
- ▶ You must manually install the AtScale UDAFs. For manual UDAF installation steps, see [Customer-Managed UDAF Installation on a Db2 Cluster](#).

Add The Data Warehouse

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 **Db2**.
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 **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 **Aggregate Schema**. AtScale reads and writes aggregate data to this database. The AtScale service account user must have ALL privileges for this database. BI tool user accounts should not have the select permission for this database.
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.

8. Select the Custom Function Installation Mode:

- ▶ **Customer Managed** - If selected, the Db2 administrator installs the custom functions in a designated schema. See [Customer-Managed UDAF Installation on a Db2 Cluster](#). Enter the name of the AtScale UDAF schema to use.
- ▶ **None** - Available for certain management edge-cases. Do not use this settings when servicing end-user queries because system performance will be degraded.

9. Click **Save** to complete the setup.

Add A Db2 Connection

After setting up your data warehouse, you must add a connection to the data warehouse before you can run queries. You can establish multiple connections with different configurations to the same data warehouse. If multiple Db2 connections are configured, you can assign which types of queries are permitted per connection (see [Query Roles](#)).

1. Expand the drop-down adjacent to the listed data warehouse and click **Create Connection** (or **Add Db2 Connection** if you have other existing connections).
2. Enter the **Name** of your Db2 connection.
3. Enter the **Host** location of the Db2 server.
4. For **Port**, enter the port number to route transmitted data. The default port number is 1025.
5. Enter the **Username** and **Password** for connecting to the database.

If you have [external secret manager](#) enabled and configured, enter the corresponding credentials.

6. (Optional) Specify any extra JDBC flags in the **Extra JDBC Flags** field only if you are certain that they are required.
 - ▶ If using Kerberos authentication, you must add the JDBC flag `LOGMECH=KRB`.
 - ▶ If using LDAP authentication, you must add the JDBC flag `LOGMECH=LDAP`.
7. Test the connection.
8. Click **Save** to complete the setup.

Query Roles

If multiple Db2 connections are configured, you can toggle which types of queries (Large Interactive, Small Interactive or System) are permitted per Db2 connection. A query role can be assigned to only one connection at a time.

Expand the drop-down for the Db2 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.