

# Adding PostgreSQL Data Warehouses

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



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

## 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 Database:** Ensure that you know the database to use for aggregate tables to be built in the data warehouse.

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.

## Add A PostgreSQL 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 **PostgreSQL**.
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 Aggregate **Database**. AtScale reads and writes aggregate data to this database.
6. Enter the name of the **Aggregate 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 PostgreSQL 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 PostgreSQL connections are configured, you can assign which types of queries are permitted per connection (see Query Roles below).

1. Expand the drop-down adjacent to the listed data warehouse and click **Create a PostgreSQL Connection** (or **Add PostgreSQL Connection** if you have other existing connections).
2. For **Name**, enter the name of your PostgreSQL connection.
3. For **Host**, enter the host address where your PostgreSQL server resides.
4. For **Port**, enter the port number to route transmitted data. The default port number is 5432.
5. For **Database**, the value input during the Data Warehouse setup will be used because the same database must be applied to all connections.
6. Enter the **Username** and **Password** associated with the PostgreSQL connection.

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

7. (Optional) Specify any extra JDBC flags in the **Extra JDBC Flags** field only if you are certain that they are required.
8. Test the connection.
9. Click **Save** to complete the setup.

## Query Roles

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

Expand the drop-down for the PostgreSQL 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.