

Adding InterSystems IRIS Data Warehouses

An InterSystems IRIS 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 InterSystems IRIS data warehouse, you must have a special AtScale license. Ask your sales representative about obtaining a InterSystems IRIS-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 An InterSystems IRIS 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 **InterSystems IRIS**.
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 namespace configured in InterSystems IRIS.
5. 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.
6. Enter the name of the **Aggregate Schema** for AtScale to use when creating aggregate tables. (AtScale will create the 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 An InterSystems IRIS 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 InterSystems IRIS 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 InterSystems IRIS Connection** (or **Add Terdata Connection** if you have other existing connections).
2. For **Name**, enter the name of your InterSystems IRIS connection.
3. For **Host**, enter the host address where your InterSystems IRIS server resides.
4. For **Port**, enter the port number to route transmitted data. The default port number is 1972.
5. Enter the name of the namespace configured in InterSystems IRIS.
6. Optionally, only if you are certain this is required, specify any extra JDBC flags in the **Extra JDBC Flags** field. For example, see the 'Connection logging' section below.
7. In the **Authorization** section, choose one of the following authorization methods:
 - ▲ For Password: enter the username and the password.
 - ▲ For Kerberos: enter the IRIS principal service name.
 - ▲ For TLS: enter the username, password, key store location, and key store password.
8. Test the connection.
9. Choose **Save** to complete the setup.

Query Roles

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

Expand the drop-down for the InterSystems IRIS 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.

Connection Logging

If you need to keep a log for the connection:

1. Start editing the connection.
2. Locate the **Extra JDBC Flags** field in the dialog.
3. Enter there the name of the log file; for example: `Iris.log`

Consider the following:

- ▲ It is allowed to add a slash sign (/) before the log name: `/Iris.log`
- ▲ It is allowed to put the log name in quotes: `"Iris.log"`
- ▲ When a log file with that name already exists, AtScale would append new entries to it.

If you want to delete the existing log and start a new one, add a plus sign (+) before the name: `+Iris.log`

4. Test the connection, then save it.

To review the log file created this way:

1. Log in to the AtScale system.
2. Go to the following directory (check your file system for detailed names): `/opt/atscale/versions/VERSION-NAME/pkg/ENGINE-VERSION/`
3. Open the file with a text editor.