

Modeling Relationships

This section explains how to create relationships within your model to get the desired matching and filtering behavior when the model is queried. A relationship links a logical dimension to a physical dataset. Relationships are not modeled between two datasets directly.

When you model a relationship, you provide information that the AtScale engine can use to join the underlying tables at query time.

Before you can add relationships to your model, you must have your dimensions defined. For more information, see [Add Dimensions to Models](#).

- ▶ [Types of Relationships in an AtScale Model](#) This section explains the types of relationships you can create in a model. A relationship creates a link between a physical dataset and a logical dimension, and adds an instance of the dimension to the model.
- ▶ [Model a Relationship](#) You don't model relationships directly between two datasets - you create them between a dataset and a logical dimension. A dimension is not considered part of a model until it has a relationship to the fact data (either directly or indirectly).
- ▶ [About Keys and Foreign Keys](#) Two datasets can be joined together on columns that they have in common. The column or columns in the dimension dataset are modeled as the key level of a dimension hierarchy. This key level can then be linked to the corresponding columns in the fact dataset (known as the foreign key).
- ▶ [About Join Behavior](#) When a query involves data from two different datasets, AtScale uses the relationships defined in the model to plan how to join the underlying tables in the connected data warehouse. This section explains the join behavior considered by the AtScale query engine.
- ▶ [Configure Join Behavior on a Dimension](#) You can configure a dimension hierarchy to specify the default matching and filtering behavior used by the AtScale query engine.
- ▶ [About Factless Queries](#) This section explains factless, dimension-only queries and use cases associated with them.