Add A Normal Dimension

This section explains how to model a logical dimension from a normalized dimension dataset.

> **Note:** Unhandled `NULL` values in key columns will result in incomplete aggregate tables and unexpected query results. See the Custom Empty Member feature for more details.

About This Task

Normal dimensions are based on a normalized dataset (a single table or view). You model a normal dimension by choosing a dataset and specifying the dataset columns that represent the key level of the dimension. For normal dimensions, the key level is unique for each record or row in the dimension dataset.

Once the dimension is created, you can then edit it to model additional hierarchies, levels, and secondary attributes.

Procedure

To add a normal dimension:

1. Open the main cube canvas.
2. Open the Dimensions panel.
3. Click the plus icon on the Dimensions panel, or drag a dataset from the Library panel to the Dimensions panel.

   In case you do not see a particular dataset in the Library panel, you can drag it from the Data Sources panel to the Library panel.

4. Complete the requested information in the **Create a Dimension** dialog.

Field	Description
Name	The user-friendly label name for the dimension. This label name will also be given to the initial hierarchy and key level that gets created for this dimension.
Dataset	The dataset that this dimension is based on. A dataset represents a table or view that has been imported into your AtScale project.
Key Columns	The dataset column that uniquely identifies a row or record in this dimension. If the dataset has a compound key, click **Add Key Column** to specify other columns that comprise the dataset's key. The choices you make here will be used to create the key level of this dimension, so the columns you choose here should have a corresponding foreign key in the fact dataset (or other dimension datasets) that you plan to join to.
Unique	The AtScale engine uses this as input when joining rows from this dimension level to other datasets in the cube model. For a normal dimension (where the Key Column values are unique for each row) **select this option**. For a multi-valued dimension (where the Key Column values can have duplicates) **leave this option unselected**.
Value Column	The column whose values are used when a user selects this dimension in a visualization or report. For example, an ID column may be the dataset key, but you want users to see the Name column values instead when used in a report or visualization.

5. Click Save.

The new dimension appears in the Dimensions panel and on the canvas.

6. Double-click the dimension to open it in the dimension editor canvas.

> **Performance Best Practice:** MDX Tools like Excel will specify filters using Attribute Key bindings. However, SQL and DAX-based tools like Tableau will specify filters using Attribute Name (Value) bindings. Data architects should be aware of their user community's tool usage so they may adopt a partitioning strategy that maximizes their query performance. If a cube services a mix of BI Tools, then it is recommended that Data Architects either partition by both name and key columns or use the same physical column for the AtScale Key and Name attribute bindings.

What To Do Next

After the dimension is created, open it in the dimension editor canvas and complete the following modeling tasks:

- Edit the initial hierarchy and make sure the dimension properties are correct.
- Edit the initial key level and make sure the level properties are correct.
- Model additional levels, secondary attributes, and hierarchies as needed.
- Create a one-to-many relationship between the dimension's key level and the corresponding foreign key columns in the fact dataset (or other dimension dataset if it is a secondary dimension).