

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. In Design Center, open the model you want to add a dimension to and select the **Canvas** tab.
2. Click the plus icon at the top of the page and select **New dimension**. A new dimension is created and opens in the **Edit Dimension** panel.
3. Complete the following fields.

| Field | Description |
|-------------|---|
| Type | Determines whether this is a standard or time dimension. |
| Unique name | The unique name for the dimension. This must be unique across all of your repositories and subrepositories. |
| Label | The user-friendly label name for the dimension. |
| Description | Optional. A description of the dimension. |

4. (Optional) Edit the dimension's default hierarchy:

1. In the **Hierarchies** section, click **Create Hierarchy**. The **Create Hierarchy** panel opens.
2. Enter a **Display name**, **Unique name**, and **Description** for the hierarchy. The **Unique name** must be unique within the dimension.
3. In the **Data Handling** section, enable the following options as needed:
 - ▲ **Exclude values with no fact data:** Configures the [join behavior](#) for the hierarchy, which determines how empty values are handled in client BI tools:
 - ▲ When enabled, query results in BI tools only include members that join to the fact dataset (inner join behavior). Members with no matching entries in the fact dataset are still included if the client BI tool requests them.
 - ▲ When disabled, query results include all members of the dimension, even those that have no matching entries in the fact dataset (outer join behavior). This occurs unless the client BI tool specifically requests to have these values filtered out.
 - ▲ **Use default member literal expression:** When enabled, defines a member of the hierarchy to use as the default filter for MDX queries on the hierarchy. Enter an MDX expression in the text box. For more information, see [About Default Hierarchical Members](#).
4. (Optional) In the **Folder** field, define the folder in which the hierarchy appears in BI tools.
5. Click **Apply**.

5. (Optional) Edit the dimension's default level:

1. In the **Hierarchies** section, click **Create Level**. The **Create Level** panel opens.
2. Complete the following fields as needed:

| Field | Description |
|--------------|--|
| Display name | Enter the display name to use for the level within AtScale. |
| Unique name | Enter a unique name for the level. This value must be unique within the dimension. |
| Time Unit | For levels on time dimensions only. Select the time unit to use for the level. |
| Description | Enter a description for the level. |
| | Select the Dataset you want to base the |

| | |
|---|---|
| Dataset | level on. |
| Key Columns | <p>The dataset column that uniquely identifies a row or record in this dimension. If the dataset has a compound key, click the plus icon to add 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.</p> |
| Unique | <p>Select this checkbox if the key columns are unique for each row. The AtScale engine uses this as input when joining rows from this dimension level to other datasets in the model.</p> <p>Select this option if you are creating a normal dimension (Key Column values are unique for each row). For multi-valued dimensions (where the Key Column values can have duplicates), leave this option unselected.</p> |
| Value Column | <p>Select the column whose values will appear for this level in BI tools. For example, the key may be a product ID number, but you want users to see product names instead.</p> |
| Sort Column | <p>Select the column to sort query results on.</p> |
| Exclude from System-Generated Dimension-Only Aggregates | <p>When enabled, this level is excluded from system generated dimension-only aggregates. This is useful if the attribute contains a large number (millions) of distinct values that you don't want to aggregate.</p> |
| Exclude from System-Generated Fact-Based Aggregates | <p>When enabled, this level is excluded from system generated fact-based aggregates. This is useful if the attribute contains a large number (millions) of distinct values that you don't want to aggregate.</p> |

| | |
|--|--|
| <p>Contains Unique Names</p> | <p>Determines whether each member of this level attribute has a unique name. Do not enable this functionality if two members have different keys but the same name.</p> |
| <p>Custom Empty Member</p> | <p>Allows Fact data with NULL or invalid foreign key values to be isolated and independently aggregated from fact records with valid foreign key values. For more information, see Using Custom Empty Members for Levels and Attributes.</p> |
| <p>Dimensionally Modified Aggregates</p> | <p>Select any calculations you want to use to create dimensionally modified aggregates for the level. Note that Periods To Date and Dates Periods To Date should only be selected when working with a time dimension.</p> |
| <p>Visualization</p> | <p>Determines whether the level appears in BI tools.</p> |

3. Click **Apply** to add the level.

6. In the **Edit Dimension** panel, click **Apply** to create the dimension.



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 model 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.