

Setting Properties To Allow Incremental Rebuilds Of Aggregates

If you want to allow the AtScale engine the option of performing incremental rebuilds of the aggregates for cubes in a project, you can configure this in the properties of the fact dataset and dimensions.

Before You Begin

Ensure that you understand how incremental builds work by reading these topics:

- ▲ [About Incremental Rebuilds](#)
- ▲ [Aggregates for Fact Datasets that Use Joins](#)

About This Task

Allowing incremental rebuilds of aggregates does not cause rebuilds to be incremental. After an aggregate is defined, the AtScale engine determines the type of rebuild it will use for the instances of the aggregate. The first instance of an aggregate is built with a full build. Rebuilds occur to refresh aggregates with new or changed data. After the AtScale engine determines which type of rebuild to use for an aggregate, the engine always uses that type of rebuild for it unless the cube is edited and the project is published again. After such a publication, the AtScale engine again determines which type of rebuild to use for a defined aggregate. The choice that the engine makes can differ from the choice it made for the previous version of the cube.

Procedure: Fact Dataset

1. Open the Cube Designer canvas of the cube.
2. Locate the fact dataset on the canvas, and preview the data. Identify a **Long, Integer, Decimal (38,0)** (Snowflake only), **Timestamp**, or **DateTime** type column that can be used as the incremental indicator.

For example, a UNIX timestamp (seconds from epoch time) whose values are sequential numbers (e.g. 1446292800). If you do not have a column that is numeric and contains sequential values, you may have to create a calculated column to transform string or datetime values into the right type.

3. Open the **Library** panel, locate the fact dataset, open its menu, and select **Edit**.
4. In the **Edit Project-Scoped Dataset** dialog, go to the 'Aggregate Preferences' section and turn on the **Allow incremental builds** option.
5. Choose the dataset column to use as the **Incremental Indicator**.
6. Specify the **Grace Period**.

This is the number of previous increments since the last build to re-process, which allows a buffer window for late arriving data. If a user chooses a Timestamp or DateTime column, then the Grace Period will be in seconds. If the user chooses a column that is not Timestamp or DateTime, the grace period will restore to the unit of the column type selected.

For example, if your indicator column values represented milliseconds (e.g 1446292800000 milliseconds from epoch time), then the grace period should be in milliseconds instead (1 day = 86400000 milliseconds).

If your indicator column values represent a **Timestamp** or **DateTime**, to configure a 1 day grace Period you would enter a value of 86400. To configure a 5 day grace Period, they would enter a value of 432000.



When the incremental indicator column type is as a Timestamp, the Grace Period field is always interpreted as Seconds.

7. Save your changes.

Procedure: Dimension

Follow these steps for each dimension that you want to allow to be included in aggregates that are rebuilt incrementally:



This way you can edit the incremental rebuild settings for the whole project. You cannot edit those settings for a cube (and override project-level settings).

1. Open the Cube Designer canvas of the cube.
2. Open the **Library** panel, locate the dimension, open its menu, and select **Edit**.
3. In the **Edit Project-Scoped Dataset** dialog, go to the 'Aggregate Preferences' section and turn on the **Allow incremental builds** option.



You must be sure that the joined dimensions rarely, if ever, change outside of the grace period. When the AtScale engine performs an incremental rebuild, it does not search through existing rows in the aggregate to replace old values or include newly appended values. As more changes are made to a dimension outside of the grace period, the aggregate instance becomes less accurate. If changes are made outside of the grace period, you should trigger a full rebuild of the aggregate table. You can trigger a full rebuild from the Aggregate Configuration section of the Design Center. See [Performing Full Rebuilds of Incremental Aggregates](#).

However, if changes are made to joined dimensions within the grace period, there is no need for a full rebuild of the aggregate.

For full details on how to allow incremental rebuilds of aggregates that use joins, see [Aggregates for Fact Datasets that Use Joins](#).

4. Configure the **Incremental Indicator** and **Grace Period** options as described above.
5. Save your changes.

Example

Remember that the AtScale engine is able to rebuild incrementally only aggregates that join to dimensions for which you follow these steps!

For example, suppose that your cube has three dimensions: Customer, Date, Order, and Product. Further suppose that these dimensions have the following settings for **Safe when joined to incremental**:

Dimension	Setting of 'Safe when joined to incremental'
Customer	True
Date	True
Order	False
Product	False

After you publish the cube and queries start to run against it, the AtScale engine defines the following aggregates and creates instances of them:

Name of Aggregate	Included Dimensions	Possible Types of Rebuild
Agg A	Customer	Full, Incremental
Agg B	Order	Full
Agg C	Customer, Date	Full, Incremental
Agg D	Order, Product	Full
Agg E	Customer, Order	Full
Agg F	Customer, Date, Order, Product	Full

The only possible type of rebuild is a full rebuild if the following two conditions are true:

- ▲ A proposed aggregate definition contains a join to at least one dimension that is not declared safe for incremental rebuilds.
- ▲ The value of the engine setting `AGGREGATE.CREATE.JOINS.ALLOWPREVENTINCREMENTAL.ENABLED` is **True**.

In this case, a full rebuild is used whether the aggregate is defined by the AtScale engine or a data modeler has defined it.

If `AGGREGATE.CREATE.JOINS.ALLOWPREVENTINCREMENTAL.ENABLED` is set to False, then the aggregate table will use incremental rebuilds, but will not join to the dimension that is not declared safe for incremental rebuilds.

Results

The settings take effect only after the cube is published.