

Cube Level Configuration Settings

Users can specify cube-scoped values for a subset of Aggregate and Query-related settings.

If an eligible setting is not defined for a specific cube, AtScale uses the value defined at the Engine or Organization level. If the value changes, all cubes that do not define a cube-level override will use the higher-level value. Changes to the cube-scoped settings are effective immediately and do not require an engine restart or project re-publish.

Attention!

If your project was published using the [rename publish](#) feature, you cannot configure cube-scoped settings associated with your project.

Before you begin

To view and edit cube settings, a user must be assigned the [Override Engine and Organization Settings](#) permission.

Viewing and Modifying Cube-Level Settings

To view cube settings:

1. Go to **Settings > Organization Settings > Cube Settings**.
2. Choose which cubes to see:

- **Overridden Values:** display only cubes where some of the settings are overridden.
- **Effective Values:** display all cubes.

3. Click the **Filter by cube** text box, and select the cubes you need.
4. Click the **Filter by setting** text box, and search or select the settings you want to be displayed; see [Filtering Settings](#) below for details.
5. Review the cube settings and their values:

- You can sort them by column.
- There is an information icon adjacent to the name of each setting. Hovering over it shows the setting's default value, and description of this setting.

To change the settings:

1. Follow the procedure above to display the settings you need.
2. Modify the settings by entering a value, or turning them on or off.
3. Choose Save.
4. In the *Save Cube Setting* dialog, review your changes and choose Continue to confirm.

To reset a setting to its default value:

1. Follow the procedure above to display the settings you need.
2. Locate a value that has been modified - it is marked with a colored asterisk.
3. Hover over the area to the right of the modified value, and when the Reset button appears choose it.
4. Save and confirm your changes.

Note

Some settings can have a significant impact on down-stream systems. For example, increasing the number of aggregates for a cube by overriding `aggregates.systemGenerated.activeInstance.retentionLimit` can increase the workload on a data warehouse. Before changing such a parameter it is recommended that you view the Effective Value for the setting across all cubes that share the same data warehouse. The user interface is designed to display these values in a column, to assess the cumulative impact of the proposed change.

Cube-level settings list

For details, see the following documents:

- [Aggregate Creation Settings](#)
- [Speculative Aggregates Settings](#)
- [Aggregate Incremental Builds Settings](#)
- [Other Aggregate Settings](#)
- [Query Settings](#)

Other cube-level settings you can view and modify are:

- `virtualization.subquery.result.max_rows`: The maximum number of rows that can be returned for a virtualized query. The default value is 10000000.
- `xmla.discover.properties.daxdialect`: The dialect to be used for DAX queries; ignored when connection-level settings are applied. For details, see [Using SuperDAX MD](#).

More information

[Engine Level Configuration Settings](#)