Dimensionally Modified Aggregates

Some of the queries processed by AtScale request Calculated Measures as well as Calculation Groups. The run-time performance of such queries can be significantly improved by enabling and configuring Dimensionally Modified Aggregates (DMA). This would allow you to:

- Reduce outbound query complexity.
- Improve aggregate hit-rate for dimensionally modified queries like "ParallelPeriod", "Lag", and "Moving Windows".

After Upgrade

Configuring DMA after upgrading from AtScale versions that do not support them is easy. When you enter a model that contains a Time dimension, a message would be displayed in the Warning Tray in the Design Center, asking you if you wish to enable the default DMA configurations for your Time Hierarchies:

- If you choose "Yes", the system would add the default DMA configuration for each Level in a Time Hierarchy.
- If you choose "No", all DMA settings for each Level in a Time Hierarchy would be turned off.

Consider that you would be able to choose Yes or No only when your user account does have the necessary permissions (see below).

Configuration

DMA are enabled by default. They can be configured as follows:

- On cube level, you can use the `query.planning.dma.enabled`, `aggregates.create.dimensionalModifications.complexityLimit`, `aggregates.create.dimensionalModifications.enabled`, and `aggregates.dimensionalModifications.retentionLimit` settings. For details, see Aggregate Creation Settings.
- You can also modify these settings on engine level (no need to restart), as described in Changing Engine Settings.

You should also enable or disable DMA creation for a specific combination of model attribute and calculation type:

- This can be done by users who have the Manage Aggregates permission. For details, see Creating and Editing Roles.
- The calculation types are available in the dialogs for dimension levels and secondary dimensional attributes. For details, see Edit a Level and Edit a Secondary Dimensional Attribute.