

# Aggregate Creation Settings

You can view and modify a lot of cube-level settings for aggregate creation.

## Joins

- ▶ `aggregate.create.joins.allowPreventIncremental.enabled` : Whether or not to consider joining to a dataset that is not safe for incremental update, if it would prevent this aggregate from otherwise being an incremental aggregate. The default value is True.
- ▶ `aggregates.create.joins.compression` : Specify the minimum compression ratio for any proposed join. This ratio is calculated as the cardinality of the join key in the fact table (or in the dimension table if that is not available) to the cardinality of the grouped dimension values; for example,  $\#(\text{Key Cardinality}) / \#(\text{Dim Table grouped by Dim Value})$ . Joins for which the compression ratio is below this minimum will not be used. The default value is 100.0.
- ▶ `aggregates.create.joins.enabled` : Set to True (default) to allow the AtScale engine to use joins when defining aggregates. This setting must be set to True for the other settings starting with `aggregates.create.joins.` to have an effect.
- ▶ `aggregates.create.joins.maximumDepth` : Specify the maximum number of dimensions that can be traversed in a join path. The default value is 3.
- ▶ `aggregates.create.joins.maximumKeyCardinality` : Specify the maximum cardinality that the AtScale engine will allow in join keys when the engine is determining whether to use a join in the definition of an aggregate. Higher cardinalities will cause the engine not to use a join. The default value is 10000000.
- ▶ `aggregates.create.joins.nonInner.enabled` : Whether to allow non-inner joins in system aggregates. Warning: it is unsafe to enable this option if a Cube has any role-played dimensions. The default value is False.
- ▶ `aggregates.create.joins.prime.compression` : Specify the minimum compression ratio for a proposed join from a prime query part (where aggregates cannot be stored anywhere except in preferred storage). The default value is 0.99.
- ▶ `aggregates.create.joins.smallJoins.enabled` : Whether to allow small joins to be added to system aggregates. The default value is False.
- ▶ `aggregates.create.joins.smallJoins.maximumCompressionRatio` : The maximum compression ratio for valid small joins. The default value is 10.0.
- ▶ `aggregates.create.joins.smallJoins.maximumKeyCardinality` : The maximum key cardinality to allow a small join. The default value is 100.

## Higher Levels

- ▶ `aggregates.create.includeHigherLevels.enabled` : Set this to True (default) to enable the addition of higher levels to system aggregates (without causing additional joins). For example, if an aggregate has the Day level, we will automatically add Month and Year.
- ▶ `aggregates.create.includeHigherLevels.maxHierarchies` : Specify the maximum number of hierarchies the dimension from system aggregates should be part of, in order to enable higher level expansion. Also depends on

`aggregates.create.includeHigherLevels.enabled` is enabled. The default value is 2.

## Widening

- ▶ `aggregates.create.widening.enabled` : Set to True (default) to allow the engine to define new aggregates as wider versions of existing aggregates. Wider aggregates contain more measures than their predecessors.
- ▶ `aggregates.create.widening.measure.limit` : Specify the maximum number of measures that can be added when widening. This setting requires `aggregates.create.widening.enabled` to be set to True. The default value is 20.

## Partitioning

- ▶ `aggregates.create.partition.hintedAggregate.enabled` : Whether to partition hinted aggregates from Query Data Sets using the Cube's partition key list. The default value is True. For this setting to have an effect, the `tables.create.partitions.enabled` setting must be set to True.
- ▶ `aggregates.create.partition.userDefinedAggregate.enabled` : Set to True to enable the AtScale engine to partition user-defined aggregates. The default value is True. For this setting to have an effect, the `tables.create.partitions.enabled` setting must be set to True.
- ▶ `aggregates.create.partition.systemDefinedAggregate.enabled` : Set to True (default) to enable the AtScale engine to partition system-defined aggregates. For this setting to have an effect, the `tables.create.partitions.enabled` setting must be set to True.
- ▶ `aggregates.create.partition.systemDefinedAggregate.threshold` : The estimated number of rows in the proposed aggregate table divided by the number of partitions, determines the estimated number of rows per partition. The default value is 50000.0.

For more information about partitioning settings, see [About Hinted Aggregate Tables](#), [Settings that Control Partitioning of System-Defined Aggregate Tables](#), and [Engine Settings for User-Defined Aggregates Only](#).

## Dimensionally Modified Aggregates

- ▶ `query.planning.dma.enabled` : Whether to use dimensionally modified aggregates during query planning. The default value is True.
- ▶ `aggregates.create.dimensionModifications.complexityLimit` : Estimated row-count limit for dimensionally modified aggregates. Used to prevent the execution of excessively complex queries that may be too expensive to run or may create excessively large tables. The default value is 1500000.
- ▶ `aggregates.create.dimensionModifications.enabled` : When set to True (default), aggregates containing dimensional modifications can be created.
- ▶ `aggregates.dimensionModifications.retentionLimit` : The number of active instances of dimensionally modified aggregates retained per cube. The default value is 30.

For more information, see [Dimensionally Modified Aggregates](#).

## Other Settings

- ▶ `aggregate.create.securityDimensions.enabled` : When set to "True", aggregates containing attributes from security dimensions can be created. The default value is False.
- ▶ `aggregates.create.aggressiveDimensionalCopyPromotion.enabled` : Enables aggressive promotion of dimensional hierarchy copies into preferred storage. The default value is True.
- ▶ `aggregates.create.allowExactDistinctCountMeasures.enabled` : When set to True, exact distinct count measures may be included in system-defined aggregates. The default value is False.
- ▶ `aggregates.create.compression.threshold` : Specify the compression factor that aggregates proposed by the engine must meet or exceed. This factor is a measure of the quality of a proposed aggregate. It is calculated as the number of rows in the fact table divided by the estimated number of rows in a proposed aggregate. The default value is 3.0.
- ▶ `aggregates.create.demandDefined.enabled` : Enables the creation of demand-defined system aggregates. The default value is True.
- ▶ `aggregates.demandDefined.disabled.deactivate` : Whether to deactivate aggregate active instances if demand-defined aggregates are disabled. This setting will take effect only if `aggregates.create.demandDefined.enabled` is set to false. The default value is True.
- ▶ `aggregates.create.narrowing.dimensional.enabled` : Allow aggregates narrowing when building dimension-only aggregates. The default value is True.
- ▶ `aggregates.create.threshold.enabled` : Set to True (default) to turn on the `aggregates.create.compression.threshold` setting.
- ▶ `aggregates.create.useIncidentalData.enabled` : Set to False to disable the creation of dimensional aggregates from not strictly related dimensions using incidental data. The default value is True.
- ▶ `aggregates.create.withoutCompressionEstimate.enabled` : Allow new aggregates to be created without estimated compression ratios (i.e. when statistics are not available). The default value is False.

## More Information

### Cube Level Configuration Settings