

Other Settings For System-Defined Aggregates Only

A number of features of system-defined aggregates are enabled or affected by particular settings. To access these settings:

1. Make sure you are a super user.
2. Choose **Settings** on top, then choose **Engine** or **Aggregates** on left.
3. If a setting is not displayed on the **Engine** page, choose Show Custom Settings in the Overview section, and configure it there.

Aggregates.systemgenerated.activeinstance.retentionlimit

The system default value for the maximum number of system-defined aggregates retained per cube. An Org-specific value may be set on **Settings > Aggregates**. Cube-specific values can be set using the Cube Settings user interface. Setting this value too high will cause long aggregate batch build times and may impact data warehouse workloads.

Aggregates.create.allowexactdistinctcountmeasures.enabled

Used for creating System-Defined Aggregates that contain Distinct Count measures. It is disabled by default. To enable it, set it's value to True.

See also `aggregates.systemgenerated.withdistinctcountmeasure.retentionpercentage` and `aggregates.systemGenerated.activeInstance.retentionLimit`.

Aggregates.systemgenerated.withdistinctcountmeasure.retentionpercentage

To prevent Distinct Count aggregates from dominating the Aggregate retention limit, AtScale defines a separate pool of Distinct Count aggregates as a percentage of the overall cube-scoped System-Defined aggregate retention limit. The percentage is controlled by this setting, at either the Engine or Cube level. It accepts values between 0 and 100, with a default of 40. The size of the Distinct Count aggregate pool is calculated as a percentage of the effective Cube value of the "retentionLimit" setting (see above).

For example if the Cube's effective retention limit is 100, and

`aggregates.systemGenerated.withDistinctCountMeasure.retentionPercentage` is 40, then the system will allow the creation of up to 40 System-Defined aggregates that use Distinct Count measures.

Aggregates.systemgenerated.activeinstance.extraallowance

The system default value for the maximum number of additional system-defined aggregates temporarily permitted per cube when the retention limit is reached. An Org-specific value may be set on **Settings > Aggregates**. Cube-specific values can be set using the Cube Settings user interface. Setting this value too high will cause long aggregate batch build times and may impact data warehouse workloads.

Aggregates.completeness.degeneratedimensionswithmeasures.enabled

Set to **True** to allow the AtScale engine to direct dimension-only queries to aggregate instances that contain complete degenerate dimensions.

Aggregates.dimensional.build

Set to **True** to allow the engine to create aggregates that contain dimensional attributes only. Such aggregates can be useful in Tableau for queries against fact datasets that contain degenerate dimensions.

Aggregates.create.includehigherlevels.enabled

Set to **True** to enable the aggregation process to evaluate higher levels of a hierarchy when building an aggregate. If a higher level meets the needed criteria, it will be built in. For example: if this setting is enabled and a hierarchy contains Year and YearMonth, when an aggregate build against YearMonth * myMeasure is requested, the build process evaluates Year as a possible candidate too.

Aggregates.prediction.checkForNonAdditiveMeasuresAndConstraints.enabled

AtScale checks for aggregate candidates to queries selecting distinct count measures and with a non-equals WHERE constraint. When this option is set to True (default), such aggregate candidates are rejected, because they attempt to re-aggregate the data.