

Settings Specific For Data Warehouses

The aggregate settings described here should be applied only when you use the corresponding data warehouse. To modify them:

1. Make sure you are logged in as a super user.
2. Go to **Settings > Engine**.
3. Locate the setting you need, and update it.

Redshift

Aggregates.create.usectas.redshift

This setting is specific for the Redshift data warehouse. When enabled (default), it allows using a faster method for creating aggregates for very large data sets. In case of performance issues you can try disabling or enabling this setting.

Snowflake

Aggregates.create.usectas.snowflake

This setting is specific for the Snowflake data warehouse. When enabled (default), it allows using a faster method for creating aggregates for very large data sets. In case of performance issues you can try disabling or enabling this setting.

InterSystems IRIS

Aggregates.create.iris.nojourn.enabled

This setting is specific for the InterSystems IRIS data warehouse, versions 2022.2 or newer. It allows you to set whether to use the %NOJOURN parameter when creating aggregate tables. The default value is 'false'.

DB2

Db2.Value.compression.enabled

Specify whether value compression for DB2 storage should be used. The default value is false.

Db2.Row.compression.enabled

Specify whether row compression for DB2 storage should be used. The default value is false.

Db2.Row.compression.type

Specify which DB2 compression method to use. This setting is applicable only if you set the value of `db2.row.compression.enabled` to true. The value should be ADAPTIVE or STATIC (default).