



This PDF is generated from authoritative online content, and is provided for convenience only. This PDF cannot be used for legal purposes. For authoritative understanding of what is and is not supported, always use the online content. To copy code samples, always use the online content.

Reporting and Analytics Aggregates

agg Section

12/17/2025

agg Section

- `agg-level-<level>-delay`
- `cancel-threshold`
- `deadlock-threshold`
- `default-tz-offsets`
- `level-of-log`
- `realtime-offset`
- `sub-hour-interval`
- `warning-threshold`
- `writer-schedule`
- `zone-offset`

Use the aggregate section of a Genesys Info Mart application to define the general behavior of the aggregation process. The values of options in this section impact all aggregation hierarchies. This section must be named **[agg]**.

agg-level-<level>-delay

Default Value: 0

Valid Values: Any positive integer

Changes Take Effect: After restart

Specifies the minimum delay (seconds) between aggregation runs, on a level-by-level basis. This option is available beginning with RAA release 8.5.001.45, and applies to materialized levels only (day and higher).

The delay is applied based on the aggregation level, as follows:

`agg-level-<level>-delay=<seconds>`

where `<level>` is one of: day, week, month, quarter, year and `<seconds>` is the number of seconds (or a calculation, such as `60*60*12`) to delay aggregation for that level. The *minimum* delay permitted for each level is as follows:

- day: 3600 seconds (1 hour)
- week: 3600*3 seconds (3 hours)
- month and quarter: 3600*12 seconds (12 hours)
- year: 86400 seconds (1 day)

Note that delaying aggregation at any given level also delays aggregation at higher levels, because each level is based on the levels below it. So, delaying aggregation of day-level [materialized] data, also delays aggregation of week, month, quarter, and year levels.

cancel-threshold

Default Value: 0

Valid Values: A value in seconds.

Changes Take Effect: After restart of the aggregation process

Introduced: 100.0.004.0100

Causes RAA to automatically exit the process using System.exit if aggregation has not stopped as expected within the number of seconds configured for this option. Set this value to a positive integer (seconds).

If this option is set to 0 (No System.Exit) and aggregation fails to stop correctly, aggregation can hang, preventing it from starting when called on subsequent occasions.

deadlock-threshold

Default Value: 1800 (30 min)

Valid Values: Any positive integer

Changes Take Effect: After restart

Specifies the amount time, in seconds, within which each aggregation writer thread must return the results of its aggregation of a batch of data. If a writer thread does not respond within this time frame, RAA assumes either that the process is deadlocked or that the database is too slow and cannot process aggregation in a timely fashion. When the deadlock-threshold time period has elapsed, RAA cancels all database queries and closes all sessions. To resume processing, aggregation must be restarted.

Genesys recommends that you do not set this option's value to less than 900 seconds.

default-tz-offsets

Default Value: 0,0

Valid Values: a,b

where: a = the number of seconds of the winter offset and b = the number of seconds of the summer offset.

Changes Take Effect: After restart

Specifies the winter and summer Universal Coordinated Time (UTC) offset, in seconds, from the time zone of the DATE_TIME table for environments:

- Whose offsets are in increments other than one hour—that is, whose offset is not evenly divisible by 3600.
- That configure more than one time zone.

For example, an offset of six and a half hours (UTC+06:30) with recognition of daylight saving time in the summer of one hour would be configured as follows:

default-tz-offsets=23400,27000

level-of-log

Default Value: .:INFO

Valid Values: [category]: [<value>][, category: [<value>]...]

where category is either “.” (for the root logging category) or “Agg”, and value corresponds to the desired level of log information: SEVERE, WARNING, INFO, CONFIG, FINE, FINER, FINEST, ALL, OFF.

Changes Take Effect: After restart

Specifies the detail level of log messages that the Genesys Info Mart Server generates for aggregation-related activity, by category. Specify “.” for the root logging category; otherwise, specify “Agg”.

The lower the value level, the greater the detail that the Genesys Info Mart Server logs. When you specify no value at all, Genesys Info Mart Server uses the default value, .:INFO. Valid levels of log detail are:

- SEVERE—Genesys Info Mart Server logs only severe messages from the corresponding category.
- WARNING—Genesys Info Mart Server logs severe and warning messages from the corresponding category.
- INFO—Genesys Info Mart Server logs severe, warning, and informational messages from the corresponding category.
- CONFIG—Genesys Info Mart Server logs severe, warning, informational, and configuration messages from the corresponding category.
- FINE—Same as CONFIG plus an even finer detail of messages from the corresponding category.
- FINER—Same as FINE plus an even finer detail of messages from the corresponding category.
- FINEST—Same as FINER plus an even finer detail of messages.
- ALL—Genesys Info Mart Server logs all messages from the corresponding category.
- OFF—Genesys Info Mart Server logs no messages from the corresponding category.

The lower the value, the greater the detail that the Genesys Info Mart Server logs. When you specify no value at all, Genesys Info Mart Server uses the default value, .:INFO.

realtime-offset

Default Value: 900 (15 minutes)

Valid Values: 0-7200 (2 hours)

Changes Take Effect: Either:

- In autonomous mode, upon restart of the aggregation process.
- In integrated mode, immediately upon every 5-minute reevaluation.

Specifies the number of seconds that the upper boundary of Zone 1 is offset from aggregation. Zone 1 contains the most recent aggregation notification requests. Use this option in conjunction with the writer-schedule and zone-offset configuration options to fine-tune aggregation dispatching.

Zone 1 represents a sliding window of time that is bound by two timestamps:

- The moment in time (t) that the aggregation process is running minus the number of seconds specified by this option defines the end of the sliding window.
- The moment in time represented by t minus the number of seconds specified by this option minus the number of seconds specified by the **zone-offset** option marks the beginning of the sliding window.

Once the aggregation process starts, every five minutes thereafter it redefines Zone 1 and the Zone 1 upper boundary. RAA will not aggregate data above this boundary.

You specify a real-time offset to eliminate the overlap that might occur between the Genesys Info Mart ETL transformation process (writing data to Info Mart) and the RAA aggregation process (reading Info Mart data). With this offset, you can prevent RAA from aggregating data that is currently being transformed—data that is likely to be significantly changed by Genesys Info Mart ETL. Moreover, on some RDBMSs—Microsoft SQL Server, in particular—database locks could result if you specify too short a value for this option. To minimize this possibility, specify a large enough value to instruct RAA to avoid processing data that Genesys Info Mart ETL is likely to be writing.

sub-hour-interval

Default Value: 30min

Valid Values: 15min, 30min

Changes Take Effect: After restart

Specifies the lowest time level of aggregation, in minutes, for the AG2_*_SUBHR tables.

You must choose a value for this option before the aggregation engine writes data to the subhour aggregation tables and avoid changing it afterwards. Otherwise, aggregation results will be difficult to interpret. If you do want to change the value of this option and data has been written to the subhour tables, you must first stop aggregation and purge all data from the SUBHR tables before resetting the value of this option. In addition, if it is necessary to have data for the period of time that data was purged, you must rerun aggregation for that period.

warning-threshold

Default Value: 300 seconds (5 minutes)

Valid Values: Any positive integer

Changes Take Effect: After restart

Introduced: 8.5.005.02

Specifies the amount of time, in seconds, within which aggregation is expected to complete. If it has not completed within the specified period of time, the plan of the SQL query of aggregation is written to the log with the log level WARNING.

writer-schedule

Default Value: `default=flex(3:1)`

(Three writers that are dedicated to Z1 and one writer that is dedicated to Z2.)

Valid Values: `default=p(a:b) [, hour(HH-HH)=p(c:d) ] [, hour(HH-HH)=p(e:f) ]` (no spaces) where:

- The `default` keyword indicates that the writer assignments for each zone define the schedule for hours that you do not explicitly configure using the `hour` keyword. Where:
 - `p` represents the degree of pliability: `flex` (for a flexible schedule) or `strict`. A flexible schedule enables RAA to borrow writer threads from the other zone when there are insufficient idle threads dedicated to the current zone to handle aggregation requests. Conversely, RAA will never borrow threads when the degree of pliability is `strict`.
 - `default=strict(3:5)` means that the default schedule mandates that 3 writers always be dedicated to Z1, and 5 always to Z2. The schedule indicated by the `hour` keyword supersedes the default schedule.
- The `hour` keyword indicates that the immediate schedule defines the writer assignments for the indicated span of whole hours using a 24-hour clock. For example:
 - `hour(8-19)` defines the immediate schedule from 8:00 am to 6:59 pm.
 - `hour(20-7)` defines the immediate schedule from 8:00 pm to 6:59 am.

This parameter also accepts the argument `hour(#-#)=purge`, which enables and schedules purging of aggregate data. For more information about purging, see **RAA Aggregation Runtime Parameters** in the *Reporting and Analytics Aggregates Deployment Guide*. There are no resets at midnight, and you can configure any number of hour constructs. RAA uses the schedule of the first encountered.

- `a`, `c`, and `e` specify the number of writers for Zone 1.
- `b`, `d`, and `f` specify the number of writers for Zone 2. The maximum number of writer Z1-Z2 pairings must not exceed 10. `default=strict(10:0)` is valid, whereas `hour(0,6)=flex(2,9)` is not; $(2+9>10)$.

Changes Take Effect: Either:

- In autonomous mode, upon the next start of the aggregation process.
- In integrated mode, immediately upon every 5-minute reevaluation.

writer-schedule controls the schedule for the number of writers that RAA dedicates to the aggregation of notifications received in Zone 1 (Z1) and Zone 2 (Z2).

- Z1 consists of the more recent notifications about pending aggregation requests of the most recent data and is bound by the timestamps implied by the values of the **realtime-offset** and **zone-offset** configuration options.
- Z2 consists of notifications about older data and is bound only by the timestamp implied by the value of the **zone-offset** configuration option. (Refer to the descriptions of these options to learn how RAA determines these timestamps.)

For more information about aggregation dispatching, see **How Do I Configure Genesys Info Mart for Aggregation?** in the *Reporting and Analytics Aggregates Deployment Guide*.

zone-offset

Default Value: 115200 (32 hours)

Valid Values: Integers between 8100 (>2 hours) and 8000000000 (>25 years) inclusive. Use of the largest values is designed to effectively eliminate Zone 2.

Changes Take Effect: After restart of the aggregation process

Specifies the length of Zone 1 (housing the most recent aggregation notification requests) in seconds. This option also indirectly defines the boundary between Zone 2 and Zone 1.

The zone offset represents a sliding window of time bound by two timestamps:

- The end of the sliding window is defined by the moment in time (t) that the aggregation process starts minus the number of seconds specified by the **realtime-offset** configuration option.
- The moment of time represented by t minus the real-time offset minus the number of seconds specified by this option marks the beginning of the sliding window.

The beginning of Zone 1 also marks the end of Zone 2. Refer to the *Reporting and Analytics Aggregates User's Guide* for an illustration of zones 1 and 2 delineated by the zone and real-time offsets. Use this option in conjunction with the **realtime-offset** and **writer-schedule** configuration options to fine-tune aggregation dispatching.