Genesys Info Mart

schedule Section

5/9/2025

schedule Section

- aggregate-duration
- aggregate-schedule
- etl-end-time
- etl-frequency
- etl-start-time
- export-schedule
- maintain-start-time
- on-demand-migration
- run-aggregates
- run-export
- run-maintain
- run-scheduler
- run-update-stats
- timezone
- update-stats-schedule

This configuration section specifies the schedule that Genesys Info Mart Server uses to launch jobs. The Genesys Info Mart Server enables options to be modified while it is running. For those options that specify a time, the time format is HH:mm, where HH represents the number of hours (00–24), and mm represents the number of minutes (00–59).

aggregate-duration

Default Value: 5:00
Valid Values: 00:00-24:00
Changes Take Effect: Immediately
Dependencies: run-aggregates, aggregate-schedule

Specifies the amount of time, in 24-hour format, that Job_AggregateGIM will run after it is launched. When the run-aggregates option is set to TRUE, the scheduler will stop the aggregation job when this interval expires. The aggregation job is launched in accordance with a schedule defined by the aggregate-schedule option. After the aggregation job is launched, it runs continuously until the aggregation-duration interval expires.

aggregate-schedule

Default Value: 0 1
Valid Values: A valid CRON expression
Changes Take Effect: Immediately
Dependencies: run-aggregates

Specifies the daily schedule for Job_AggregateGIM to start. The job will start in accordance with this schedule when aggregation is being controlled by the scheduler (in other words, the run-aggregates option is set to true). Between them, the aggregate-schedule and aggregate-duration options define daily time intervals within which Job_AggregateGIM will run continuously.

The schedule is defined in the format of a CRON expression that represents a set. The expression

comprises two fields, which are separated by whitespace:

- The first field specifies minutes. Valid values are 0–59 and optional special characters (see below).

- The second field specifies hours. Valid values are 0–23 and allowed special characters.

The following special characters are allowed in the CRON expression:

- , (comma)—Separates items in a list. For example, specifying the first field (minutes) as 0,30,45 means the 0th, 30th, and 45th minutes of the hour.

- - (hyphen)—Defines a range. For example, specifying the first field (minutes) as 30-35 means every minute between the 30th and 35th minute of the hour, inclusive; this is the same as specifying 30,31,32,33,34,35.

- * (asterisk)—Indicates that the CRON expression will match for all values of the field. For example, specifying the second field (hours) as * means every hour in the day.

- / (forward slash)—Describes increments. For example, specifying the first field (minutes) as 0/10 means the 0th minute of the hour and every 10 minutes thereafter.

Examples

The following values for **aggregate-schedule** illustrate sample schedules:

- 0 1 means that the aggregation job will be launched once a day at 01:00.

- 30 0,3/2 means that the aggregation job will be launched every day at 00:30, 03:30, and every 2 hours after that for the rest of the day.
  This schedule assumes that the value of **aggregate-duration** is 02:00 or less. The scheduler will not launch a new instance of Job_AggregateGIM while an existing instance is running. For aggregation to run on the specified schedule, the value of **aggregate-duration** must not exceed the intervals between scheduled start times.

- 30 * means that the aggregation job will be launched every hour during the day on the half-hour (00:30, 01:30, 02:30, and so on), assuming that the value of **aggregate-duration** is 01:00 or less.

Genesys recommends against configuring a schedule that has the aggregation job running in a series of short bursts—for example, **aggregate-schedule**=30 * and **aggregate-duration**=00:15. When the time specified by **aggregate-duration** expires, the scheduler immediately stops the aggregation job, even if it is in the middle of processing a batch of data.

If you want Job_AggregateGIM to run continuously for 24 hours a day, without any breaks for maintenance activities (which is not recommended), set **aggregate-schedule**=0 0 and **aggregate-duration**=24:00.

etl-end-time

Default Value: 22:00
Valid Values: 00:00-23:59
Changes Take Effect: Immediately
Dependencies: run-scheduler, etl-start-time

Specifies the time of day, in 24-hour format, when the last ETL cycle can start running. If the value

that you specify is before the ETL start time, the end time is for the next day (past midnight).

If etl-start-time=etl-end-time, the ETL cycle will run continuously.

Ensure that you configure **etl-start-time** and **etl-end-time** so that there is sufficient time for the last ETL cycle of the day to complete and for Job_MaintainGIM to run (see the **maintain-start-time** option), before the start of the first ETL cycle of the next day.

etl-frequency

Default Value: 1
Valid Values: 0-1440
Changes Take Effect: On the next ETL cycle
Dependencies: None

Specifies the number of minutes that pass between the start times of each ETL cycle. If the amount of time that it takes to complete a cycle is shorter than the specified value, the next cycle is delayed until the time elapses. If the amount of time that it takes to complete a cycle is longer than the specified value, the next cycle is started immediately.

The ETL frequency must not be greater than the chunk size for data extraction, as specified by the **extract-data-chunk-size** option. Otherwise, Genesys Info Mart will not be able to keep pace with ICON. When it checks the deployment, Genesys Info Mart verifies the internal consistency between the ETL frequency and extraction chunk size.

By default, the value of **etl-frequency** is much smaller than the value of **extract-data-chunk-size**. Genesys recommends that you retain this relationship, to minimize data latency. For example, say that **extract-data-chunk-size**=900 (15 minutes), **etl-frequency**=1, and all data from the last chunk has been processed; when the next ETL cycle starts 1 minute later, there is only 1 minute's worth of new data, and this can be processed very quickly. Alternatively, if there is a backlog of data, and it takes less than 15 minutes to process a 15-minute chunk, the next ETL cycle starts almost immediately, to continue catching up.

etl-start-time

Default Value: 06:00
Valid Values: 00:00-23:59
Changes Take Effect: Immediately
Dependencies: run-scheduler

Specifies the time of day, in 24-hour format, when the first ETL cycle starts running.

export-schedule

Default Value: 20 0/8
Valid Values: A valid CRON expression
Changes Take Effect: Immediately
Dependencies: run-export
Introduced: 8.5.005

Defines the time intervals at which Job_ExportGIM will run. The job will start and then run periodically in accordance with this schedule when the **run-export** option is set to **true**. By default, the job runs at 00:20, 08:20, and 16:20 every day.

The default schedule, run in conjunction with the default **chunk-size-seconds** option in the **[gim-export]** section, is designed to keep daily disruptions or delays from carrying over to the next day.

Job_ExportGIM can run in conjunction with the ETL jobs, but not in conjunction with Job_MaintainGIM.

The schedule is defined in the format of a CRON expression that represents a set. The expression comprises two fields, which are separated by whitespace:

- The first field specifies minutes. Valid values are 0–59 and optional special characters (see below).
- The second field specifies hours. Valid values are 0–23 and allowed special characters.

The following special characters are allowed in the CRON expression:

- , (comma)—Separates items in a list.
- - (hyphen)—Defines a range.
- * (asterisk)—Indicates that the CRON expression will match for all values of the field.
- / (forward slash)—Describes increments.

maintain-start-time

Default Value: 03:00
Valid Values: 00:00-23:59
Changes Take Effect: Immediately
Dependencies: run-maintain

Specifies the time of day, in 24-hour format, when Job_MaintainGIM is started. This job is scheduled to start at this time when the run-maintain option is set to TRUE. The value that you specify must be outside the range that is specified by etl-start-time and etl-end-time.

> Tip
> If the time of day that is represented by the new value has already passed, the new value is applied to the following day.

on-demand-migration

Default Value: false
Valid Values: true, false
Changes Take Effect: When Genesys Info Mart next enters the migration state
Dependencies: None
Introduced: 8.5.007

Controls whether Genesys Info Mart will run Job_MigrateGIM automatically if the Info Mart database schema is not up to date following migration of the Info Mart server.

- true — Genesys Info Mart will launch Job_MigrateGIM automatically if the schema is not up to date.

- false — Genesys Info Mart will not launch Job_MigrateGIM automatically if the schema is not up to date and Genesys Info Mart enters the migration state.

> ### Important
> Genesys does not recommend enabling migration on demand unless policies and procedures are in place to ensure that essential pre-migration and post-migration steps are also performed without manual intervention — for example, frequent database backup and re-creation of read-only views following migration.

The value of false preserves Genesys Info Mart legacy behavior, which requires you to run Job_MigrateGIM manually before ETL functioning will resume if Genesys Info Mart has entered the migration state.

Enabling on-demand migration is suitable only if you always want to apply schema updates immediately, without review and without controlling the timing of the schema update or required system preparation.

run-aggregates

Default Value: false
Valid Values: true, false
Changes Take Effect: Immediately
Dependencies: None

Specifies whether the scheduler will manage the aggregation job, to run the aggregation engine inside the Genesys Info Mart process.

- When the value of this option is set to true, the scheduler will start Job_AggregateGIM at the scheduled time, as specified by the **aggregate-schedule** option; Job_AggregateGIM will then run continuously until the scheduler stops the job after the scheduled interval, as specified by the **aggregate-duration** option. The scheduler will not allow a second aggregation process to be launched while the job is running. The scheduler will also not allow any other aggregation process to be launched outside the intervals that are defined by the **aggregate-schedule** and **aggregate-duration** options.

- When the value of this option is set to false, the scheduler does not manage the aggregation job at all, leaving aggregation in whatever state it is in when the option value is set. Note that this means that, if you change the value of **run-aggregates** to false while the aggregation job is running, the scheduler will never stop the job.

For example, if **run-aggregates**=true, **aggregate-schedule**=0 1, and **aggregate-duration**=05:00, the aggregation job will run continuously between 01:00 AM and 06:00 AM daily. The scheduler will not allow you to launch a second instance of Job_AggregateGIM manually from the management GUI (Genesys Info Mart Manager or the Genesys Info Mart Administration Console) within that time period. Furthermore, if you try to launch an instance of Job_AggregateGIM manually from the management GUI outside that time period (for example, at 08:00 AM), the scheduler will

identify that the job is not supposed to be running at that time and will stop it. If you want to run Job_AggregateGIM manually from the management GUI outside the scheduled times, you must first set **run-aggregates** to `false`.

For more information about starting Job_AggregateGIM and managing the aggregation job from the management GUI, see the *Genesys Info Mart Operations Guide*.

run-export

Default Value: false
Valid Values: true, false
Changes Take Effect: Immediately
Dependencies: None
Introduced: 8.5.005

Specifies whether Job_ExportGIM will run. When the value of this option is set to **true**, the scheduler will start and run the job at the time and intervals specified by the **export-schedule** option.

run-maintain

Default Value: true
Valid Values: true, false
Changes Take Effect: Immediately
Dependencies: None

Specifies whether to run Job_MaintainGIM at the scheduled time, as specified by the maintain-start-time option.

run-scheduler

Default Value: false
Valid Values: true, false
Changes Take Effect: Immediately
Dependencies: None

Specifies whether to stop or start the scheduler. If the value of this option was set to true, so that the scheduler is currently scheduling jobs, and you change the value of this option to FALSE, the scheduler pauses, with no effect on any jobs that might already be running. If you then reset the value to TRUE, the scheduler resumes at the point at which it stopped.

run-update-stats

Default Value: false
Valid Values: true, false
Changes Take Effect: Immediately
Dependencies: None

Specifies whether Job_UpdateStats will run in PostgreSQL deployments, at the time and intervals

specified by the update-stats-schedule option.

timezone

Default Value: GMT
Valid Values: Any valid Java time zone
Changes Take Effect: Immediately
Dependencies: None

Specifies the time zone in which the schedule is defined. Internally, Genesys Info Mart maintains the schedule in UTC time. For convenience, you can use this option to specify a local time zone that makes it easier for you to plan and manage the schedule. You can use any valid time zone that is supported by the version of the JRE that runs the Genesys Info Mart Server.

For more information about supported time zones, see the documentation about calendar time zones on the Java developer website or other public resources. For sample reference sites, see the description of the date-time-tz option.

update-stats-schedule

Default Value: 0/10 *
Valid Values: A valid CRON expression
Changes Take Effect: Immediately
Dependencies: run-update-stats

Defines the time intervals at which Job_UpdateStats will run. The job will start and then run periodically in accordance with this schedule. By default, the job runs every 10 minutes throughout the day. Job_UpdateStats can run in conjunction with the ETL jobs, but not in conjunction with Job_MaintainGIM.

The schedule is defined in the format of a CRON expression that represents a set. The expression comprises two fields, which are separated by whitespace:

- The first field specifies minutes. Valid values are 0–59 and optional special characters (see below).

- The second field specifies hours. Valid values are 0–23 and allowed special characters.

The following special characters are allowed in the CRON expression:

- , (comma)—Separates items in a list. For example, specifying the first field (minutes) as 0,30,45 means the 0th, 30th, and 45th minutes of the hour.

- - (hyphen)—Defines a range. For example, specifying the first field (minutes) as 30-35 means every minute between the 30th and 35th minute of the hour, inclusive; this is the same as specifying 30,31,32,33,34,35.

- * (asterisk)—Indicates that the CRON expression will match for all values of the field. For example, specifying the second field (hours) as * means every hour in the day.

- / (forward slash)—Describes increments. For example, specifying the first field (minutes) as 0/10 means the 0th minute of the hour and every 10 minutes thereafter.

The schedule that you configure for Job_UpdateStats does not need to specifically allow for a

maintenance window: A running instance of Job_UpdateStats does not prevent Job_MaintainGIM from starting, and once Job_MaintainGIM has started as part of the schedule, the scheduler suspends the schedule for Job_UpdateStats until Job_MaintainGIM finishes.

For values that illustrate sample schedules, see the examples for the aggregate-schedule option.