



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.

Genesys Info Mart Private Edition Guide

[GIM configuration options](#)

6/11/2026

Contents

- 1 level
- 2 console-pattern-layout
- 3 appender.ConsoleLogger.Threshold
- 4 days-to-keep-active-facts
- 5 days-to-keep-deleted-annex
- 6 says-to-keep-discards-and-job-history
- 7 days-to-keep-gidb-facts
- 8 days-to-keep-gim-facts
- 9 How Purge Options Work
 - 9.1 Purge examples
- 10 etl-start-date
- 11 max-chunks-per-job
- 12 max-time-deviation
- 13 memory-threshold
- 14 partitioning-ahead-range
 - 14.1 Example
- 15 partitioning-interval-size-gidb
- 16 partitioning-interval-size-gidb-mm
- 17 partitioning-interval-size-gidb-ocs
- 18 partitioning-interval-size-gim
- 19 purge-thread-pool-size
- 20 purge-transaction-size
- 21 date-time-max-days-ahead
- 22 date-time-min-days-ahead
- 23 date-time-start-year
- 24 date-time-tz
- 25 first-day-of-week
- 26 fiscal-year-start
- 27 fiscal-year-end
- 28 fiscal-year-week-pattern
- 29 min-days-in-first-week

-
- [30 simple-week-numbering](#)
 - [31 aggregate-duration](#)
 - [32 aggregate-schedule](#)
 - [33 Examples](#)
 - [34 etl-end-time](#)
 - [35 etl-frequency](#)
 - [36 etl-start-time](#)
 - [37 export-schedule](#)
 - [38 maintain-start-time](#)
 - [39 run-aggregates](#)
 - [40 run-export](#)
 - [41 run-maintain](#)
 - [42 run-scheduler](#)
 - [43 run-update-stats](#)
 - [44 on-demand-migration](#)
 - [45 timezone](#)
 - [46 update-stats-schedule](#)
 - [47 chunk-size-seconds](#)
 - [48 days-to-keep-output-files](#)
 - [49 max-retries](#)
 - [50 output-directory](#)
 - [51 output-files-encoding](#)
 - [52 retry-delay-seconds](#)
 - [53 start-date](#)
 - [54 thread-pool-size](#)
 - [55 use-export-views](#)

Options you can use to control GIM behavior.

Related documentation:

-
-
-
-

RSS:

- [For private edition](#)

GIM supports the configuration options described on this page. See [Configure GIM behavior](#) for information about how to change the option values in the Helm chart.

GIM behavior affects data in the Info Mart database dimensional model. For full information about the Info Mart database tables and columns referenced in the descriptions on this page, see the *Genesys Info Mart Historical Database Reference*.

- aggregate-duration
- aggregate-schedule
- appender.ConsoleLogger.Threshold
- chunk-size-seconds
- console-pattern-layout
- date-time-max-days-ahead
- date-time-min-days-ahead
- date-time-start-year
- date-time-tz
- days-to-keep-active-facts
- days-to-keep-deleted-annex
- days-to-keep-gidb-facts
- days-to-keep-gim-facts
- days-to-keep-output-files
- etl-end-time
- etl-frequency
- etl-start-date
- etl-start-time
- export-schedule
- first-day-of-week
- fiscal-year-end
- fiscal-year-start
- fiscal-year-week-pattern
- level
- maintain-start-time
- max-chunks-per-job
- max-retries
- max-time-deviation
- memory-threshold
- min-days-in-first-week
- on-demand-migration
- output-directory
- output-files-encoding
- partitioning-ahead-range
- partitioning-interval-size-gidb-mm
- partitioning-interval-size-gidb-ocs
- partitioning-interval-size-gidb
- partitioning-interval-size-gim
- purge-thread-pool-size
- purge-transaction-size
- retry-delay-seconds
- run-aggregates
- run-export
- run-maintain
- run-scheduler
- run-update-stats
- says-to-keep-discards-and-job-history
- simple-week-numbering
- start-date
- thread-pool-size
- timezone
- update-stats-schedule
- use-export-views

Important

When you specify an override value in the **values.yaml** file, you indicate the configuration level for that option value. The configuration level specifies the scope for which an option value is in effect. For example, for certain ***-threshold** options, you can specify different thresholds for the various media types, and you can further refine behavior by specifying different thresholds for individual queues. Only some options support configuration at different levels.

All options are configured at the tenant level. Below the tenant level, the available configuration levels are:

- **standard** — Applies across media types.
- **media** — Applies to the specified media type, such as voice or email.
- **dn: media** — Applies to the specified Virtual Queue DN. An option value specified on the DN level overrides a value specified for that option at the media level.
- **script: media** or **script: standard** — Applies to the specified Script object representing an Interaction Queue or Interaction Workbin for digital interactions. An option value specified on the Script level overrides a value specified for that option at the media or standard levels.

level

Default value: info

Valid values: debug, info, warn, error, none, all, off

Configuration level: standard

Determines whether local logging is enabled, and specifies the minimum level of events to log.

The DEBUG, INFO and WARN values correspond to the Genesys Management Layer DEBUG, TRACE, and STANDARD logging values, respectively.

[>> Back to list](#)

console-pattern-layout

Default value: %d{ISO8601} %-5p %-12t %m%n

Valid values: Any valid pattern string

Configuration level: standard

For information about this option, refer to the Apache logging site: <http://logging.apache.org/index.html>

>> [Back to list](#)

appender.ConsoleLogger.Threshold

Default value: info

Valid values: trace, debug, info, warn, error, all, off

Configuration level: standard

For information about this option, refer to the Apache logging site: <http://logging.apache.org/index.html>

>> [Back to list](#)

days-to-keep-active-facts

Default value: 27

Valid values: Any positive integer

Configuration level: standard

Specifies the maximum number of days to retain active multimedia interactions in GIDB and the dimensional model, including certain Staging tables, after which the interactions become eligible for purging. Depends on the value of **days-to-keep-gim-facts**.

>> [Back to list](#)

days-to-keep-deleted-annex

Default value: 2

Valid values: Any positive integer

Configuration level: standard

Specifies the number of days to retain deleted records in the *_ANNEX dimension tables.

The default value of days-to-keep-deleted-annex is small, because there is likely little reason to retain deleted data for significant periods of time. The major reason that Genesys Info Mart provides

*_ANNEX data is to support GCXI visibility controls, and GCXI uses only active *_ANNEX data for this purpose.

For example, with the default value (2 days), if Job_MaintainGIM is running on September 4, 2022, the job purges *_ANNEX records that terminated (in other words, the configuration setting on the object's Annex tab was deleted) before September 2, 2022.

>> [Back to list](#)

says-to-keep-discards-and-job-history

Default value: 60

Valid values: Any positive integer

Configuration level: standard

Specifies the number of days to retain data in the discard tables and audit and history tables.

The discard tables are Staging tables that store operational data that the transformation job was unable to process—for example, voice interaction data with unresolved IS-Links, or Configuration details records with missing configuration objects. The audit and history tables are Control tables that store information about data lineage and about ETL processing activity. Information in the discard, audit, and history tables is useful for troubleshooting.

Records in the discard, audit, and history tables are purged based on the timestamp of the ETL processing event—ETL_TS for the discard (Staging) tables and CREATED_TS for the audit and history (Control) tables.

For example, if Job_MaintainGIM is running on August 23, 2022 (day 235 of the year) and days-to-keep-discards-and-job-history=600, Job_MaintainGIM purges all records in the discard, audit, and history tables that were written by instances of the ETL jobs that ran before January 1, 2021 (day 1 of the previous year).

>> [Back to list](#)

days-to-keep-gidb-facts

Default value: 27

Valid values: Any positive integer

Configuration level: standard

Specifies the number of days to retain fact data in GIDB. Facts that have a start time that is earlier than the retention period are eligible to be purged. Related Options: days-to-keep-active-facts

The retention period for multimedia facts in GIDB is determined solely by the value of **days-to-keep-gidb-facts**.

Genesys expects the value of **days-to-keep-gidb-facts** to be less than 30. Although many GIDB tables contain personally identifiable information (PII), Genesys Info Mart does not include most GIDB tables in General Data Protection Regulation (GDPR) processing.

[>> Back to list](#)

days-to-keep-gim-facts

Default value: 400

Valid values: Any positive integer

Configuration level: standard

Specifies the number of days to retain fact data in the dimensional model. Facts that have a start time that is earlier than the retention period are eligible to be purged. Job_MaintainGIM does not purge active fact data, dimension data, or aggregate tables (if aggregation is enabled). Genesys Info Mart enforces that the value of **days-to-keep-active-facts** *be less than the value of days-to-keep-gim-facts* **and** *days-to-keep-gidb-facts be less than the value of days-to-keep-gim-facts*.

Genesys Info Mart does not enforce a limit on the length of the retention period. However, Genesys strongly recommends against setting **days-to-keep-gim-facts**, and likewise **days-to-keep-active-facts**, to an excessively large value (for example, thousands of days). To avoid performance issues or job failures when Genesys Info Mart operates against an excessively large Info Mart database, be realistic about your requirements and available database resources, and apply best practices and recommendations from your RDBMS vendor. If you choose a very long retention period, you must carefully choose and tune the storage used for the Info Mart database; also carefully consider settings for partition size and other database- and performance-related configuration options, as well as parameters on the RDBMS side.

By definition, all voice interaction data in the dimensional model relates to completed interactions.

How Purge Options Work

To illustrate how various days-to-keep-* options combine to determine when data is purged, the following tables provide examples of different scenarios for GIDB and the dimensional model, respectively.

Purge examples

GIDB

Scenario	Start Time of Facts to Be Purged
<ul style="list-style-type: none">• days-to-keep-gidb-facts=14.• days-to-keep-active-facts=30.• Job_MaintainGIM is running on December 4, 2020, after the last ETL cycle for the day. In other words:	November 21, 2020

Scenario	Start Time of Facts to Be Purged
<ul style="list-style-type: none"> • days-to-keep-gidb-facts threshold is November 21, 2020 • days-to-keep-active-facts threshold is November 5, 2020 	

Dimensional Model

Scenario	Start Time of Facts to Be Purged
<ul style="list-style-type: none"> • days-to-keep-gim-facts=400. • days-to-keep-active-facts=30. • Job_MaintainGIM is running on December 4, 2020, after the last ETL cycle for the day. In other words: <ul style="list-style-type: none"> • days-to-keep-gim-facts threshold is November 1, 2019 • days-to-keep-active-facts threshold is November 5, 2020 	<p>November 1, 2019, or earlier</p> <p>The oldest timestamp for artificially terminated multimedia interactions is November 5, 2020.</p>

>> [Back to list](#)

etl-start-date

Default value:

Valid values: Any date after 1970 in the format yyyy-mm-dd hh:mm:ss

Configuration level: standard

Specifies the earliest date for which Genesys Info Mart considers IDB data for extraction in a new deployment or when the Info Mart database is re-initialized. IDB data that has timestamps earlier than the ETL start date is never extracted.

The option is used only when Job_InitializeGIM initializes the database. If the **etl-start-date option** is not specified, the earliest starting point for Genesys Info Mart processing is IDB data that has timestamps 30 days prior to the Info Mart database initialization.

The main purpose of the option is to pre-empt performance and maintenance issues when Genesys Info Mart is introduced into a deployment with much older existing IDB data. Specifying the ETL start date enables users to:

- Shorten the backlog of IDB data if they do not need to include all of the old data in their reporting.
- In deployments with a partitioned Info Mart database, specify the starting point for creating partitions, which is important for proper partitioning.

Genesys Info Mart ignores this option value for Configuration details; after the database has been initialized or re-initialized, Genesys Info Mart extracts all cfg data going back as far as 2010, regardless of the value of **etl-start-date**.

[>> Back to list](#)

max-chunks-per-job

Default value: 10

Valid values: 1-100

Configuration level: standard

Specifies the number of extracted data chunks that the transformation job processes in one ETL cycle. As long as it seems practical from the performance perspective, increase the option value to transform a larger amount of data in a single cycle.

In deployments using the Data Export feature, the option also controls the maximum number of export chunks that can be created by the export job during each job run. In situations where the time range to be exported exceeds the export chunk size (for example, if there is a backlog because data is being re-exported), a larger number of export chunks reduces the time required to process a large export backlog.

[>> Back to list](#)

max-time-deviation

Default value: 30

Valid values: 1-120

Configuration level: standard

Specifies the maximum time deviation, in seconds, to consider for time synchronization inaccuracy between host-system clocks. This option also specifies the maximum acceptable delay the transformation job accommodates in scenarios in which ICON creates delayed records in the GM_F_USERDATA table in IDB.

If the maximum time deviation or delay is exceeded, Genesys Info Mart results become unreliable.

Genesys recommends the following relationship between the value of this option and the Advanced Disconnect Detection Protocol (ADDP) timeout, which is the Local Timeout parameter configured for ADDP connections to data sources on the Connections tab of the ICON Application: $[(ADDP\ Local\ Timeout) * 2] + (actual\ maximum\ difference\ in\ time\ synchronization\ between\ hosts)$

In HA deployments, Genesys Info Mart uses max-time-deviation for reliable analysis of the “no data” information in IDB; (NoData – max-time-deviation) is considered to be reliable for all data sources for a particular ICON provider.

[>> Back to list](#)

memory-threshold

Default value: 0

Valid values: 0-99

Configuration level: standard

Specifies the percentage of available memory that must be exceeded before Genesys Info Mart logs a message (55-20101) that indicates that the memory threshold has been exceeded. If the value of this option is 0, the feature is disabled.

[>> Back to list](#)

partitioning-ahead-range

Default value: 31

Valid values: Any positive integer

Configuration level: standard

Specifies, in terms of number of days, how far ahead Job_InitializeGIM (in the first instance) and Job_MaintainGIM (on an ongoing basis) create partitions for GIDB, Control, and Info Mart fact tables that are partitioned. These jobs add partitions for the number of days ahead of the time that the job is running. (Job_InitializeGIM also adds partitions from etl-start-date up to the time that the job is running.) The number of partitions that Job_MaintainGIM actually creates during each run depends on the partition sizes and the job frequency.

Example

If `partitioning-interval-size-gidb=86400` (one day), `partitioning-interval-size-gim=604800` (one week), and `partitioning-ahead-range=14`, Job_MaintainGIM creates as many other partitions as necessary to provide partitions for GIDB, Control, and Info Mart fact tables up to 14 days ahead. If Job_MaintainGIM runs daily, then:

- For GIDB tables, each run of Job_MaintainGIM creates one new partition of size 1 day, for the fourteenth day. (Previous runs have created partitions for the other days.)
- For Info Mart fact tables and Control tables, the maintenance job creates one new partition of size 7 days at the start of each week. (A previous run will have created a partition for the other week.)

If the value of **partitioning-ahead-range** is not a multiple of **partitioning-interval-size-gim**, the maintenance job creates a new partition only when the last day of the partitioning ahead range falls in a week for which a partition has not been created. For example, if `partitioning-interval-size-gim=604800` (7 days) but `partitioning-ahead-range=10`, two new partitions are created on the first run of Job_MaintainGIM, and the next partition is created on the fifth run.

To guarantee that partitions are always available for use by the ETL, ensure that run-scheduler is set to true (the default value).

This option applies only in deployments that use partitioning.

[>> Back to list](#)

partitioning-interval-size-gidb

Default value: 604800

Valid values: Any positive integer

Configuration level: standard

Specifies the size of partitions, in seconds, for GIDB tables that are partitioned. Job_MaintainGIM creates partitions of the specified size in the Info Mart database in preparation for future ETL cycles. The default size of GIDB table partitions is 24 hours (86400 seconds).

In PostgreSQL deployments, Genesys recommends setting the size of GIDB table partitions to one week (604800 seconds).

This option applies only in deployments that use partitioning.

[>> Back to list](#)

partitioning-interval-size-gidb-mm

Default value: 604800

Valid values: Any positive integer

Configuration level: standard

Specifies the size of partitions, in seconds, for partitioned GIDB tables that store multimedia interaction data. When this option is set, Job_MaintainGIM creates partitions of the specified size in the Info Mart database in preparation for future ETL cycles. If the option is not specified, the value of partitioning-interval-size-gidb, which has a default value of 24 hours (86400 seconds), is used.

Genesys recommends increasing the size of GIDB partitions for multimedia interactions, which typically live longer than voice interactions but generate a smaller volume of data.

In PostgreSQL deployments, Genesys recommends setting the size of GIDB table partitions to one week (604800 seconds).

Transformation performance is optimal when partitions are large enough that data that is being actively used is in the fewest number of partitions, and the partitions are small enough that their indexes can fit into the cache.

[>> Back to list](#)

partitioning-interval-size-gidb-ocs

Default value: 604800

Valid values: Any positive integer

Configuration level: standard

Specifies the size of partitions, in seconds, for partitioned GIDB tables that store Outbound Contact-related data. When this option is set, Job_MaintainGIM creates partitions of the specified size in the Info Mart database in preparation for future ETL cycles. If the option is not specified, the value of partitioning-interval-size-gidb, which has a default value of 24 hours (86400 seconds), is used.

In PostgreSQL deployments, Genesys recommends setting the size of GIDB table partitions to one week (604800 seconds).

Transformation performance is optimal when partitions are large enough that data that is being actively used is in the fewest number of partitions, but small enough that their indexes can fit into the cache.

[>> Back to list](#)

partitioning-interval-size-gim

Default value: 2592000

Valid values: Any positive integer

Configuration level: standard

Specifies the size of partitions, in seconds, for Info Mart fact and Control tables that are partitioned. Job_MaintainGIM creates partitions of the specified size in the Info Mart database in preparation for future ETL cycles.

In PostgreSQL deployments, the recommended size of partitions for dimensional-model data depends on your plans for data retention in the Info Mart database. For PostgreSQL, Genesys recommends setting the size of fact table partitions to:

- One month (2592000 seconds) if data retention is under three years (**days-to-keep-gim-facts** is less than 1095)
- Two or three months (5184000 or 7776000 seconds) if data retention is more than three years (**days-to-keep-gim-facts** is greater than 1095).

This option applies only in deployments that use partitioning.

[>> Back to list](#)

purge-thread-pool-size

Default value: 32

Valid values: Any positive integer

Configuration level: standard

Specifies the maximum number of concurrent purging transactions. The optimum value for this option depends on the characteristics and capacity of your deployment. Consider increasing the value of this option if you think that there is scope to improve performance of the purge operation.

[>> Back to list](#)

purge-transaction-size

Default value: 100000

Valid values: Any positive integer

Configuration level: standard

Specifies the number of deleted records per table that are committed in a single transaction.

For example:

- If there are 150,000 records in a particular table that are eligible for purging and **purge-transaction-size=100000**, Job_MaintainGIM deletes and commits 100,000 records in one transaction and 50,000 records in a separate transaction.
- If there are 90,000 records in one table, 10,000 records in another table, and **purge-transaction-size=100000**, Job_MaintainGIM deletes and commits 90,000 records from the first table in one transaction and 10,000 records from the second table in a separate transaction.

[>> Back to list](#)

date-time-max-days-ahead

Default value: 400

Valid values: Any positive integer

Configuration level: standard

Specifies, in number of days, how far ahead the calendar dimension table is populated. The default value specifies that the calendar dimension is populated up to a year in advance (365 days + 1 day)

for leap years). Genesys does not recommend that you populate the calendar tables more than a year in advance, in case there are changes to DST or other international time standards that might invalidate the prepopulated data.

Note: Ensure that you populate the calendar far enough ahead to meet the requirements of your reporting intervals.

[>> Back to list](#)

date-time-min-days-ahead

Default value: 183

Valid values: Any positive integer

Configuration level: standard

Specifies, in number of days that remain in the prepopulated calendar, when the calendar table is updated with the next batch of days ahead. The default value specifies that the maintenance job updates this calendar approximately 6 months before it expires.

[>> Back to list](#)

date-time-start-year

Default value: 2020

Valid values: 1970-2038

Configuration level: standard

Specifies the year that the calendar starts. When you are setting this option, ensure that you choose a start year that provides sufficient buffer to prevent inconsistencies or unexpected missing dimensions around the start of the calendar. Genesys recommends that you set the value so that the calendar starts at least one year prior to any date that might be encountered in the data. Be aware that Genesys Info Mart uses GMT for internal time references, and this affects exactly when the calendar starts.

For example, if the other **[date-time]** options that affect the start date are set so that the calendar starts at 00:00 AM on January 1, 2022, and the **date-time-tz** option is set to Eastern European Time (GMT + 2), the calendar table is populated with dimensions starting at 02:00 AM on January 1, 2022.

[>> Back to list](#)

date-time-tz

Default value: GMT

Valid values: Any valid Java time zone

Configuration level: standard

Specifies the time zone for the calendar. You can use any valid time zone that is supported by the version of the Java Runtime Environment (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.

Sample public resources:

- http://www.java2s.com/Tutorial/Java/0120__Development/GettingallthetimezonesIDs.htm
- <http://en.wikipedia.org/wiki/Zone.tab>

Important

Particularly in deployments that use GCXI or RAA, ensure that the time zone is set appropriately for your deployment before you initialize Genesys Info Mart or before aggregation starts.

>> [Back to list](#)

first-day-of-week

Default value: 1

Valid values: 1-7 (Sunday-Saturday)

Configuration level: standard

Specifies the day of the week that is considered to be the start of the week. For example, 1 (Sunday) is usually the first day of the week in the United States; for countries that use the ISO 8601 standard, 2 (Monday) is the first day of the week.

>> [Back to list](#)

fiscal-year-start

Default value:

Valid values: Any valid combination of month and day, in M-d format

Configuration level: standard

Specifies the month and day that the fiscal year starts. For example, 1-1 means January 1; 10-1 means October 1. Set either this value, or **fiscal-year-end**, but not both.

If `simple-week-numbering=true`, every fiscal year starts on the fixed date that is specified by this option. If **simple-week-numbering=false**, the fiscal year starts on the first day of the week that contains the date that is specified by this option; however, the actual start date depends on the value of the **first-day-of-week** option. Genesys Info Mart adjusts automatically for the floating fiscal year. For example, if **simple-week-numbering=false**, **fiscal-year-start=3-1**, and **first-day-of-week=1**, then:

- Fiscal year 2012 starts on February 26.
- Fiscal year 2013 starts on February 24.
- Fiscal year 2014 starts on February 23.

This option requires that **fiscal-year-week-pattern** be set to a valid pattern.

[>> Back to list](#)

fiscal-year-end

Default value:

Valid values: Any valid combination of month and day, in M-d format.

Configuration level: standard

Specifies the month and day that the fiscal year ends. For example, 1-31 means January 31; 11-30 means November 30. Set either this value, or **fiscal-year-start**, but not both.

For example, to have calendar date 2022-11-01 in fiscal year 2023, set no value for **fiscal-year-start**, and set **fiscal-year-end=10-31**. This option requires that **fiscal-year-week-pattern** be set to a valid pattern.

[>> Back to list](#)

fiscal-year-week-pattern

Default value: none

Valid values: none, 544, 454, 445

Configuration level: standard

Specifies the pattern for the number of weeks in each month of a fiscal quarter. For example, 544 means 5 weeks in the first month, 4 weeks in the second month, and 4 weeks in the third month of

each quarter. A value of none means that the calendar is not a fiscal one.

[>> Back to list](#)

min-days-in-first-week

Default value: 1

Valid values: 1-6

Configuration level: standard

Specifies the minimum number of days from the new year that must be in the first week of the year, if simple week numbering is not used and there are no partial weeks in the calendar year. The ISO 8601 standard does not use simple week numbering.

The ISO 8601 definition of the first week in the year is the week that has the first Thursday in it. To conform to the ISO 8601 standard, set `simple-week-numbering=false`, `first-day-of-week=2`, and `min-days-in-first-week=4`.

For example, if `simple-week-numbering=false`, `first-day-of-week=2`, and January 1 of the new year is on a Friday, there are 3 days from the new year in the week that starts on Monday, December 28. Therefore:

- If the value of this option is set to 1, the calendar counts the first week of the new year as starting on Monday, December 28.
- If the value of this option is set to 4, the week that starts on Monday, December 28, is assigned to the previous year, and the calendar counts the first week of the new year as starting on Monday, January 4.

[>> Back to list](#)

simple-week-numbering

Default value: true

Valid values:

Configuration level: standard

Specifies whether the calendar year and the week-numbering year coincide.

For simple week numbering, Week 1 always begins on the first day of the calendar year (for Gregorian calendars, January 1; for fiscal calendars, the day that is specified in the `fiscal-year-start` option). As a result, the first and last weeks of the year might be partial weeks, because the first week does not necessarily start with the day that is specified by the `first-day-of-week` option. To comply with ISO 8601 week numbering, set the value of this option to false.

[>> Back to list](#)

aggregate-duration

Default value: 23:00

Valid values: 00:00-24:00

Configuration level: standard

Specifies the amount of time, in 24-hour format, that Job_AggregateGIM runs after it is launched. When the run-aggregates option is set to TRUE, the scheduler stops 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.

[>> Back to list](#)

aggregate-schedule

Default value: 30 0

Valid values: A valid CRON expression

Configuration level: standard

Specifies the daily schedule for Job_AggregateGIM to start. The job starts 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 runs 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 matches 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 launches once a day at 01:00.
- `30 0,3/2` means that the aggregation job launches 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 does 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 launches 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`.

[>> Back to list](#)

etl-end-time

Default value: 23:30

Valid values: 00:00-23:59

Configuration level: standard

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 runs 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.

[>> Back to list](#)

etl-frequency

Default value: 1

Valid values: 0-1440

Configuration level: standard

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 can not 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.

[>> Back to list](#)

etl-start-time

Default value: 00:00

Valid values: 00:00-23:59

Configuration level: standard

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

[>> Back to list](#)

export-schedule

Default value: 0/30

Valid values: A valid CRON expression

Configuration level: standard

Defines the time intervals at which Job_ExportGIM runs. The job starts and runs 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.

[>> Back to list](#)

maintain-start-time

Default value: 23:40

Valid values: 00:00-23:59

Configuration level: standard

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.

[>> Back to list](#)

run-aggregates

Default value: true

Valid values: true, false

Configuration level: standard

Specifies whether the scheduler manages 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 starts `Job_AggregateGIM` at the scheduled time, as specified by the `aggregate-schedule` option; `Job_AggregateGIM` then runs continuously until the scheduler stops the job after the scheduled interval, as specified by the `aggregate-duration` option. The scheduler does not allow a second aggregation process to be launched while the job is running. The scheduler also does 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. This means that, if you change the value of **run-aggregates** to `false` while the aggregation job is running, the scheduler never stops the job.

For example, if **run-aggregates**=`true`, **aggregate-schedule**=`01`, and **aggregate-duration**=`05:00`, the aggregation job runs continuously between 01:00 AM and 06:00 AM daily. The scheduler does 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 identifies that the job is not supposed to be running at that time and stops 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`.

[>> Back to list](#)

run-export

Default value: true

Valid values: true, false

Configuration level: standard

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

[>> Back to list](#)

run-maintain

Default value: true

Valid values: true, false

Configuration level: standard

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

[>> Back to list](#)

run-scheduler

Default value: true

Valid values: true, false

Configuration level: standard

Specifies whether to stop or start the scheduler. If the value of this option was set to TRUE so that the scheduler is 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.

[>> Back to list](#)

run-update-stats

Default value: true

Valid values: true, false

Configuration level: standard

Specifies whether Job_UpdateStats runs in PostgreSQL deployments, at the time and intervals specified by the **update-stats-schedule** option.

[>> Back to list](#)

on-demand-migration

Default value: true

Valid values: true, false

Configuration level: standard

Controls whether Genesys Info Mart runs 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 launches Job_MigrateGIM automatically if the schema is not up to date.
- false — Genesys Info Mart does 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 previous Genesys Info Mart behavior, which requires you to run Job_MigrateGIM manually before ETL functioning resumes 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.

[>> Back to list](#)

timezone

Default value: UTC

Valid values:

Configuration level: standard

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.

[>> Back to list](#)

update-stats-schedule

Default value: 0/10

Valid values: A valid CRON expression

Configuration level: standard

Defines the time intervals at which Job_UpdateStats runs. The job starts and then runs 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 white space:

- 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.

[>> Back to list](#)

chunk-size-seconds

Default value: 86400

Valid values: Any positive integer

Configuration level: standard

Specifies the size of the time interval, in seconds, for which data is exported in each job.

[>> Back to list](#)

days-to-keep-output-files

Default value: 30

Valid values: Any positive integer

Configuration level: standard

Specifies how many days to store exported files before deleting them.

[>> Back to list](#)

max-retries

Default value: 3

Valid values: Any non-negative integer

Configuration level: standard

Specifies the maximum numbers of retries the job does in the case of intermittent failures before failing the job.

[>> Back to list](#)

output-directory

Default value: gim-export

Valid values: Valid directory path (might not exist)

Configuration level: standard

Specifies the directory where exported files are stored.

[>> Back to list](#)

output-files-encoding

Default value: utf8

Valid values: Character encoding supported by Java

Configuration level: standard

Specifies the character encoding for exported files. For Java-supported encodings, see Supported Encodings.

[>> Back to list](#)

retry-delay-seconds

Default value: 30

Valid values: Any non-negative integer

Configuration level: standard

Specifies the amount of time, in seconds, that the job waits (in the case of intermittent failure) before attempting to run again.

[>> Back to list](#)

start-date

Default value:

Valid values: Any date after 1970 in the format yyyy-mm-dd hh:mm:ss

Configuration level: standard

Specifies the earliest date for which Job_ExportGIM exports data, and then continues with incremental exports. If no date is specified (the default), Job_ExportGIM starts exporting from the beginning of the Info Mart data. Requires that the value of **days-to-keep-discards-and-job-history** is greater than or equal to the value of **days-to-keep-gim-facts**.

- This option applies only when the directory in which Genesys Info Mart stores the exported files (the directory specified by the output-directory option) is empty.
- This option is not supported when audit log retention time is shorter than facts retention time.

[>> Back to list](#)

thread-pool-size

Default value: 10

Valid values: Any positive integer

Configuration level: standard

Specifies the maximum number of worker threads that are used to export data concurrently.

[>> Back to list](#)

use-export-views

Default value: true

Valid values: true, false

Configuration level: standard

Controls whether Genesys Info Mart exports data using export views, which represent a snapshot of the Info Mart schema at the time the export views were created.

- **false**—Genesys Info Mart exports data using the current schema. The export might include schema changes, such as new or renamed tables or columns, that occurred as part of a recent Info Mart migration.
- **true**—Genesys Info Mart exports data using export views. The export includes the same tables and columns as the previous export(s), regardless of any schema changes resulting from migrations that have occurred in the interim.

[>> Back to list](#)