Genesys Info Mart

Genesys Configuration Options Current

12/29/2021

Table of Contents

Genesys Info Mart Options Reference

Welcome to the Options Reference for Genesys Info Mart. This document provides full information about all the configuration options that are set on the Genesys Info Mart application object and in Genesys Info Mart–related configuration sections on other objects, such as DNs.

Important

The valid values shown in the option descriptions reflect the acceptable values that control the described behavior. In some cases, Genesys Info Mart will apparently accept other values without necessarily generating an error when Genesys Info Mart checks the configuration or during runtime. To prevent unexpected behavior, always observe the valid values reported in this document.

Genesys Info Mart Application

Options for this component are contained in the following configuration sections:

- date-time
- elasticsearch-<data-source-id>
- error-policy
- gim-etl
- gim-etl-media-chat
- gim-etl-media-email
- gim-etl-populate
- gim-export
- gim-transformation
- kafka-<cluster-name>
- log
- log4j
- schedule

Tip

In the summary table(s) below, type in the Search box to quickly find options, configuration sections, or other values, and/or click a column name to sort the table. Click an option name to link to a full description of the option. Be aware that the default and valid values are the values in effect with the latest release of the software and may have changed since the release you have; refer to the full description of the option to see information for earlier releases.

Power users: Download a CSV file containing default and valid values and descriptions.

The following options are configured at the application level (in other words, on the application object).

Section	Option	Default	Changes Take Effect
date-time	date-time-max-days-ahead	366	At the next run of Job_MaintainGIM
date-time	date-time-min-days-ahead	183	At the next run of Job_MaintainGIM
date-time	date-time-start-year	2012	At the next run of Job_MaintainGIM
date-time	date-time-table-name	DATE_TIME	At the next run of Job_MaintainGIM
date-time	date-time-tz	GMT	At the next run of Job_MaintainGIM
date-time	first-day-of-week	1	At the next run of Job_MaintainGIM
date-time	fiscal-year-start	No default value	At the next run of Job_MaintainGIM
date-time	fiscal-year-week-pattern	none	At the next run of Job_MaintainGIM
date-time	min-days-in-first-week	1	At the next run of Job_MaintainGIM
date-time	simple-week-numbering	true	At the next run of Job_MaintainGIM
elasticsearch-<data-source-id>	client	off	On the next ETL cycle
elasticsearch-<data-source-id>	g:index-interval	No default value	On the next ETL cycle
elasticsearch-<data-source-id>	g:tenant-prefix	No default value	On the next ETL cycle
elasticsearch-<data-source-id>	rest:max-content-length	100 MB	At the next run of Job_TransformGIM
elasticsearch-<data-source-id>	sources:extra	None	On the next ETL cycle
error-policy	error-policy-call-mergecall-missing	resume	On the next ETL cycle
error-policy	error-policy-campaign-group-missing	exception	On the next ETL cycle
error-policy	error-policy-cfg-check-backup-data-source	error	At the next configuration check
error-policy	error-policy-ipurpose-numberformat	resume	On the next ETL cycle
error-policy	error-policy-irf-exception	log_db_resume	On the next ETL cycle
error-policy	error-policy-irf-exception-resumable	Exception	On the next ETL cycle
Section	**Option**	**Default**	**Changes Take Effect**

Section	Option	Default	Changes Take Effect
error-policy	error-policy-islink-dangling	resume	On the next ETL cycle
error-policy	error-policy-islink-multiple-sources	resume	On the next ETL cycle
error-policy	error-policy-islink-multiple-targets	resume	On the next ETL cycle
error-policy	error-policy-islink-multiple-vertices	resume	On the next ETL cycle
error-policy	error-policy-islink-source-party-missing	resume	On the next ETL cycle
error-policy	error-policy-party-created-duplicated	resume	On the next ETL cycle
error-policy	error-policy-party-created-missing	resume	On the next ETL cycle
error-policy	error-policy-party-parent-missing	resume	On the next ETL cycle
gim-etl	aggregation-engine-class-name	none	On restart
gim-etl	days-to-keep-active-facts	30	At the next run of Job_TransformGIM or Job_MaintainGIM
gim-etl	days-to-keep-cfg-facts	400 or days-to-keep-gim-facts	At the next run of Job_MaintainGIM
gim-etl	days-to-keep-deleted-annex	2	At the next run of Job_MaintainGIM
gim-etl	days-to-keep-discards-and-job-history	600	At the next run of Job_MaintainGIM
gim-etl	days-to-keep-gdpr-history	15	At the next run of Job_MaintainGIM
gim-etl	days-to-keep-gidb-facts	14	At the next run of Job_MaintainGIM
gim-etl	days-to-keep-gim-facts	400	At the next run of Job_MaintainGIM
gim-etl	delayed-data-threshold	900	On the next ETL cycle
gim-etl	etl-start-date	Initialization date minus 30 days	On database initialization
gim-etl	extract-data-cfg-facts-chunk-size	90000	On the next ETL cycle for the DAP with role=ICON_CFG
gim-etl	extract-data-chunk-size	900	On the next ETL cycle
gim-etl	extract-data-max-conn	128	On the next ETL cycle
gim-etl	extract-data-stuck-threshold	28860	On the next ETL cycle
Section	**Option**	**Default**	**Changes Take Effect**

Section	Option	Default	Changes Take Effect
gim-etl	extract-data-thread-pool-size	32	On the next ETL cycle
gim-etl	extract-last-second	false	On the next ETL cycle
gim-etl	link-msf-userdata-mm	false	On the next ETL cycle
gim-etl	link-msf-userdata-voice	false	On the next ETL cycle
gim-etl	max-call-duration	3600	On the next ETL cycle
gim-etl	max-camp-group-session-duration-in-hours	168	On the next ETL cycle
gim-etl	max-camp-group-state-duration-in-hours	168	On the next ETL cycle
gim-etl	max-chain-processing-duration-in-hours	8	On the next ETL cycle
gim-etl	max-chunks-per-job	10	On the next ETL cycle
gim-etl	max-msfs-per-irf	50	On the next ETL cycle
gim-etl	max-parties-per-call	100	On the next ETL cycle
gim-etl	max-session-duration-in-hours	24	On the next ETL cycle
gim-etl	max-state-duration	14400	On the next ETL cycle
gim-etl	max-thread-duration-after-inactive-in-days	31	On the next ETL cycle
gim-etl	max-time-deviation	30	On the next ETL cycle
gim-etl	memory-threshold	0	Immediately
gim-etl	merge-chunk-size	200000	On the next ETL cycle
gim-etl	merge-failed-is-link-timeout	0	On the next ETL cycle
gim-etl	partitioning-ahead-range	14	At the next run of Job_MaintainGIM
gim-etl	partitioning-interval-size-gidb	86400	At the next run of Job_MaintainGIM
gim-etl	partitioning-interval-size-gidb-mm	86400 or partitioning-interval-size-gidb	At the next run of Job_MaintainGIM
gim-etl	partitioning-interval-size-gidb-ocs	86400 or partitioning-interval-size-gidb	At the next run of Job_MaintainGIM
gim-etl	partitioning-interval-size-gim	86400	At the next run of Job_MaintainGIM
gim-etl	purge-thread-pool-size	32	At the next run of Job_MaintainGIM
gim-etl	purge-transaction-size	100000	At the next run of Job_MaintainGIM
gim-etl	q-answer-threshold-voice	60	On the next ETL cycle
Section	**Option**	**Default**	**Changes Take Effect**

Section	Option	Default	Changes Take Effect
gim-etl	q-short-abandoned-threshold-voice	10	On the next ETL cycle
gim-etl	short-abandoned-threshold	10	On the next ETL cycle
gim-etl	sm-resource-state-priority	ACW, NOT_READY, BUSY, READY	On the next ETL cycle
gim-etl	user-event-data-timeout	3600	On the next ETL cycle
gim-etl-media-chat	q-answer-threshold	60	At the next run of Job_TransformGIM
gim-etl-media-chat	q-short-abandoned-threshold	10	On the next ETL cycle
gim-etl-media-chat	short-abandoned-threshold	10	On the next ETL cycle
gim-etl-media-email	q-answer-threshold	60	At the next run of Job_TransformGIM
gim-etl-media-email	q-short-abandoned-threshold	10	On the next ETL cycle
gim-etl-media-email	short-abandoned-threshold	10	On the next ETL cycle
gim-etl-populate	populate-irf-asm-engage-duration	false	At the next run of Job_TransformGIM
gim-etl-populate	populate-media-neutral-sm-facts	false	On the next ETL cycle
gim-etl-populate	populate-mm-ixnqueue-facts	false	On the next ETL cycle
gim-etl-populate	populate-mm-workbin-facts	true	On the next ETL cycle
gim-etl-populate	populate-sip-im-facts	false	On the next ETL cycle
gim-etl-populate	populate-sm-busy-from-mm-ixns	false	On the next ETL cycle
gim-etl-populate	populate-sm-chat-resource-activity	true	On the next ETL cycle
gim-etl-populate	populate-sm-email-resource-activity	true	On the next ETL cycle
gim-etl-populate	populate-sm-resource-session-facts	true	On the next ETL cycle
gim-etl-populate	populate-sm-resource-state-facts	true	On the next ETL cycle
gim-etl-populate	populate-sm-resource-state-reason-facts	true	On the next ETL cycle
gim-etl-populate	populate-sm-voice-resource-activity	true	On the next ETL cycle
gim-etl-populate	populate-thread-facts	false	At the next run of
Section	**Option**	**Default**	**Changes Take Effect**

Section	Option	Default	Changes Take Effect
			Job_TransformGIM
gim-etl-populate	populate-workbin-as-hold	false	On the next ETL cycle
gim-export	chunk-size-seconds	86400	On the next run of Job_ExportGIM
gim-export	days-to-keep-output-files	14	On the next run of Job_ExportGIM
gim-export	encrypt-certificate	No default value	
gim-export	max-retries	3	On the next run of Job_ExportGIM
gim-export	output-directory	output	On the next run of Job_ExportGIM
gim-export	output-files-encoding	utf8	On the next run of Job_ExportGIM
gim-export	retry-delay-seconds	30	On the next run of Job_ExportGIM
gim-export	start-date	No default value	On the next run of Job_ExportGIM
gim-export	thread-pool-size	10	On the next run of Job_ExportGIM
gim-export	use-export-views	false	On the next run of Job_ExportGIM
gim-transformation	adjust-vq-time-by-strategy-time	false	At the next run of Job_TransformGIM
gim-transformation	canceled-queues	iWD_Canceled	At the next run of Job_TransformGIM
gim-transformation	cb-virtual-queue-pattern	.*	At the next run of Job_TransformGIM
gim-transformation	chunk-size	7200	At the next run of Job_TransformGIM
gim-transformation	completed-queues	iWD_Completed	At the next run of Job_TransformGIM
gim-transformation	default-ivr-to-self-service	false	At the next run of Job_TransformGIM
gim-transformation	expand-mediation-time-for-gapless	true	On the next ETL cycle
gim-transformation	fix-missing-party-links	false	At the next run of Job_TransformGIM
gim-transformation	ignored-reason-codes	INTERACTION_WORKSPACE	At the next run of Job_TransformGIM
gim-transformation	introduced-transfer-threshold	0	At the next run of Job_TransformGIM
gim-transformation	irf-io-parallelism	4	At the next run of
Section	**Option**	**Default**	**Changes Take Effect**

Section	Option	Default	Changes Take Effect
			Job_TransformGIM
gim-transformation	ixn-data-limit	10000	At the next run of Job_TransformGIM
gim-transformation	kafka-idle-timeout	10 (seconds)	On the next ETL cycle
gim-transformation	link-vrp-vq-msf-to-irf	false	At the next run of Job_TransformGIM
gim-transformation	msf-target-route-thru-queue	false	At the next run of Job_TransformGIM
gim-transformation	ocs-allowed-lateness	PT0S	On the next ETL cycle
gim-transformation	ocs-caf-aggregates-calls	true	At the next run of Job_TransformGIM
gim-transformation	ocs-chain-history-limit	5000	At the next run of Job_TransformGIM
gim-transformation	pipeline-timeout-in-hours	1	At the next run of Job_TransformGIM
gim-transformation	routing-target-regular-dn-fold-external	true	At the next run of Job_TransformGIM
gim-transformation	show-non-queue-mediation-mm	false	On the next ETL cycle
gim-transformation	stop-ixn-queues	No default value	At the next run of Job_TransformGIM
gim-transformation	ud-io-parallelism	5	At the next run of Job_TransformGIM
kafka-<cluster-name>	bootstrap.servers	No default value	On the next ETL cycle
kafka-<cluster-name>	g:topic:<topic-name>	No default value	On the next ETL cycle
log	standard	network	Immediately
log	verbose	standard	Immediately
log4j	console-pattern-layout	%d{ISO8601} %-5p %-12t %m%n	Immediately
log4j	file-pattern-layout	%d{ISO8601} %-5p %-12t %m%n	Immediately
log4j	log-file-name	gim_etl.log	Immediately
log4j	log4j.appender.ConsoleLogger.Threshold	info	Immediately
log4j	logging-level	info	Immediately
log4j	max-backup-index	10	Immediately
log4j	max-log-file-size	50MB	Immediately
schedule	aggregate-duration	5:00	Immediately
schedule	aggregate-schedule	0 1	Immediately
schedule	etl-end-time	22:00	Immediately
schedule	etl-frequency	1	On the next ETL cycle
schedule	etl-start-time	06:00	Immediately
schedule	export-schedule	20 0/8	Immediately
Section	**Option**	**Default**	**Changes Take Effect**

Section	Option	Default	Changes Take Effect
schedule	maintain-start-time	03:00	Immediately
schedule	on-demand-migration	false	When Genesys Info Mart next enters the migration state
schedule	run-aggregates	false	Immediately
schedule	run-export	false	Immediately
schedule	run-maintain	true	Immediately
schedule	run-scheduler	false	Immediately
schedule	run-update-stats	false	Immediately
schedule	timezone	GMT	Immediately
schedule	update-stats-schedule	0/10 *	Immediately
Section	**Option**	**Default**	**Changes Take Effect**

Other Configuration Objects

Database Access Point (DAP)

The following options are configured at the DAP level (in other words, on the DAP object).

Section	Option	Default	Changes Take Effect
gim-etl	agg-jdbc-url	No default value	On restart of the Genesys Info Mart Server.
gim-etl	cp-reuse-count	-1	On restart of the Genesys Info Mart Server
gim-etl	default-schema	No default value	For an extraction DAP, at the next run of the extraction job for the particular data domain; for the Info Mart DAP, on restart of the Genesys Info Mart Server.
gim-etl	geo-location	""	For an extraction DAP, at the next run of the extraction job for the particular data domain; for the Info Mart DAP, on restart of the Genesys Info Mart Server.
Section	**Option**	**Default**	**Changes Take Effect**

Section	Option	Default	Changes Take Effect
gim-etl	jdbc-host	No default value	For an extraction DAP, at the next run of the extraction job for the particular data domain; for the Info Mart DAP, on restart of the Genesys Info Mart Server.
gim-etl	jdbc-port	1433 (for Microsoft SQL Server), 1521 (for Oracle), or 5432 (for PostgreSQL)	For an extraction DAP, at the next run of the extraction job for the particular data domain; for the Info Mart DAP, on restart of the Genesys Info Mart Server.
gim-etl	jdbc-sid	No default value	For an extraction DAP, at the next run of the extraction job for the particular data domain; for the Info Mart DAP, on restart of the Genesys Info Mart Server.
gim-etl	jdbc-url	No default value	For an extraction DAP, at the next run of the extraction job for the particular data domain; for the Info Mart DAP, on restart of the Genesys Info Mart Server.
gim-etl	role	No default value	For an extraction DAP, the next time that extraction jobs are launched; for the Info Mart DAP, on restart of the Genesys Info Mart Server.
Section	**Option**	**Default**	**Changes Take Effect**

Switch

The following options are configured at the Switch level (in other words, on the Switch object).

Section	Option	Default	Changes Take Effect
gim-etl	factor-dnd-into-sm-resource-states	false (for voice-handling and SIP switches); true (for multimedia-handling switches)	On the next ETL cycle
gim-etl	network-switch	false	At the next run of
Section	**Option**	**Default**	**Changes Take Effect**

Section	Option	Default	Changes Take Effect
			Job_ExtractICON
gim-etl	q-answer-threshold-voice	60	On the next ETL cycle
gim-etl	q-short-abandoned-threshold-voice	10	On the next ETL cycle
gim-etl-media-chat	q-answer-threshold	60	At the next run of Job_TransformGIM
gim-etl-media-chat	q-short-abandoned-threshold	10	On the next ETL cycle
gim-etl-media-email	q-answer-threshold	60	At the next run of Job_TransformGIM
gim-etl-media-email	q-short-abandoned-threshold	10	On the next ETL cycle
Section	**Option**	**Default**	**Changes Take Effect**

Media Type Business Attribute

The following options are configured on the Media Type business attribute.

Section	Option	Default	Changes Take Effect
gim-etl-media	q-answer-threshold	60	At the next run of Job_TransformGIM
gim-etl-media	q-short-abandoned-threshold	10	On the next ETL cycle
gim-etl-media	short-abandoned-threshold	10	On the next ETL cycle
Section	**Option**	**Default**	**Changes Take Effect**

DN

The following options are configured at the DN level (in other words, on the DN object).

Section	Option	Default	Changes Take Effect
gim-etl	link-msf-userdata	false	On the next ETL cycle
gim-etl	q-answer-threshold-voice	60	On the next ETL cycle
gim-etl	q-short-abandoned-threshold-voice	10	On the next ETL cycle
gim-etl-media-chat	q-answer-threshold	60	At the next run of
Section	**Option**	**Default**	**Changes Take Effect**

Section	Option	Default	Changes Take Effect
			Job_TransformGIM
gim-etl-media-chat	q-short-abandoned-threshold	10	On the next ETL cycle
gim-etl-media-email	q-answer-threshold	60	At the next run of Job_TransformGIM
gim-etl-media-email	q-short-abandoned-threshold	10	On the next ETL cycle
Section	**Option**	**Default**	**Changes Take Effect**

Script

The following options are configured at the Script level (in other words, on the Script object).

Section	Option	Default	Changes Take Effect
gim-etl	link-msf-userdata	false	On the next ETL cycle
gim-etl-media-chat	q-answer-threshold	60	At the next run of Job_TransformGIM
gim-etl-media-chat	q-short-abandoned-threshold	10	On the next ETL cycle
gim-etl-media-email	q-answer-threshold	60	At the next run of Job_TransformGIM
gim-etl-media-email	q-short-abandoned-threshold	10	On the next ETL cycle
gim-etl-populate	populate-mm-ixnqueue-facts	false	On the next ETL cycle
gim-etl-populate	populate-mm-workbin-facts	true	On the next ETL cycle
Section	**Option**	**Default**	**Changes Take Effect**

Field

The following options are configured indirectly (through ICON) at the Field level (in other words, on the Field object).

Section	Option	Default	Changes Take Effect
gim-etl-mapping	column-name	No default value	On the next ETL cycle
gim-etl-mapping	table-name	No default value	On the next ETL cycle
Section	**Option**	**Default**	**Changes Take Effect**

date-time Section

- date-time-max-days-ahead
- date-time-min-days-ahead
- date-time-start-year
- date-time-table-name

- date-time-tz
- first-day-of-week
- fiscal-year-start
- fiscal-year-week-pattern

- min-days-in-first-week
- simple-week-numbering

Use this configuration section to specify options for populating the DATE_TIME table. To configure a custom calendar, create a similar section that has the same options; name the section by using the *date-time-* prefix.

Job_InitializeGIM populates data in all configured calendars when it initializes the Info Mart database.

Job_MaintainGIM subsequently maintains the calendars in accordance with options that are specified in the **[date-time]** and custom **[date-time-*]** configuration sections. The maintenance job automatically adjusts for special requirements such as daylight saving time (DST) and fiscal years that do not start on the same day every year (floating fiscal years).

> ## Important
>
> Consider the settings for the **date-time** options carefully before the calendar dimension tables are populated for the first time. You can subsequently change the values of the **date-time-min-days-ahead** and **date-time-max-days-ahead** options at any time. However, changing any of the other **date-time** options during runtime can introduce inconsistencies into the calendar data and affect reporting results adversely. For example, if you change the timezone option (**date-time-tz**) after Genesys Info Mart has been initialized, your reports might mix the results for different timezones within the same reporting interval. If you want to change calendar options during runtime, see Changing Calendar Dimensions in the *Genesys Info Mart Operations Guide,* which provides information about additional steps that are required to maintain reporting consistency.

date-time-max-days-ahead

Default Value: 366
Valid Values: Any positive integer
Changes Take Effect: At the next run of Job_MaintainGIM

Dependencies: None

Specifies, in number of days, how far ahead the calendar dimension table will be populated. The default value specifies that the calendar dimension will be 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.

date-time-min-days-ahead

Default Value: 183
Valid Values: Any positive integer
Changes Take Effect: At the next run of Job_MaintainGIM
Dependencies: None

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

date-time-start-year

Default Value: 2012
Valid Values: 1970-2038
Changes Take Effect: At the next run of Job_MaintainGIM
Dependencies: None
Modified: 8.1.1 (in release 8.1.0, the default value was 2010)

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 will start at 00:00 AM on January 1, 2012, and the **date-time-tz** option is set to Eastern European Time (GMT + 2), the calendar table will be populated with dimensions starting at 02:00 AM on January 1, 2012.

date-time-table-name

Default Value: DATE_TIME
Valid Values: Any string that is a valid table name for your RDBMS
Changes Take Effect: At the next run of Job_MaintainGIM
Dependencies: None

Specifies the name of the table in the Info Mart database schema. You must manually modify the script that creates the custom calendar table, to specify this value as the table name.

date-time-tz

Default Value: GMT
Valid Values: Any valid Java time zone
Changes Take Effect: At the next run of Job_MaintainGIM
Dependencies: None

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.

first-day-of-week

Default Value: 1
Valid Values: 1-7 (Sunday-Saturday)
Changes Take Effect: At the next run of Job_MaintainGIM
Dependencies: None

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.

fiscal-year-start

Default Value: No default value
Valid Values: Any valid combination of month and day, in M-d format
Changes Take Effect: At the next run of Job_MaintainGIM
Dependencies: fiscal-year-week-pattern is set to a valid pattern
Introduced: 8.1.1

Specifies the month and day that the fiscal year starts. For example, 1-1 means January 1; 10-1 means October 1. This functionality was introduced in release 8.1.1.

- 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.

fiscal-year-week-pattern

Default Value: none
Valid Values: none, 544, 454, 445
Changes Take Effect: At the next run of Job_MaintainGIM
Dependencies: None

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 will not be a fiscal one.

min-days-in-first-week

Default Value: 1
Valid Values: 1-6
Changes Take Effect: At the next run of Job_MaintainGIM
Dependencies: simple-week-numbering=false

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 will count 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, will be assigned to the previous year, and the calendar will count the first week of the new year as starting on Monday, January 4.

simple-week-numbering

Default Value: true
Valid Values: true, false
Changes Take Effect: At the next run of Job_MaintainGIM
Dependencies: None

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 will 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.

elasticsearch-<data-source-id> Section

- client
- g:index-interval

- g:tenant-prefix
- rest:max-content-length

- sources:extra

Data source–specific configuration sections enable you to specify Elasticsearch connections for different types of data that could be stored in an Elasticsearch database and retrieved by Genesys Info Mart.

By default, the Genesys Info Mart application template includes the **[elasticsearch-sdr0]** section and options, for specifying the Elasticsearch cluster that Genesys Info Mart uses to retrieve Session Detail Record (SDR) data in environments with Genesys Designer. (Support for Genesys Designer is available in certain Genesys Engage cloud and on-premises implementations.)

If you want Genesys Info Mart to process data that a particular data source stored in an Elasticsearch database and make this data available in the Info Mart database for downstream reporting applications, add the applicable configuration section(s) and options on the **Options** tab of the Genesys Info Mart Application object. Name the section **[elasticsearch-<data-source-id>],** where **<data-source-id>** matches the name of the data source that Genesys Info Mart supports.

Genesys Info Mart supports the following **elasticsearch-*** configuration sections for the specified features and data sources:

Configuration Section	Data Source	Feature	Introduced
elasticsearch-bgs0	Bot Gateway Server (BGS) -- restricted releases prior to BGS 9.0.004.08	Chat bot activity	8.5.011.04
elasticsearch-ldr0	CX Contact (CXC)	CX Contact campaigns - Unattempted records	8.5.012.15
elasticsearch-sdr0	Designer applications	Session Detail Records (SDRs)	8.5.001.12
elasticsearch-sdr1	Designer applications	SDR survey transcriptions	8.5.005.20

Each data source–specific configuration section can contain the following options.

client

Default Value: off
Valid Values: off or any valid location of the cluster node(s) of the Elasticsearch cluster, properly formatted
Changes Take Effect: On the next ETL cycle
Dependencies: None
Introduced: 8.5.009.20

This option specifies one or more nodes in the Elasticsearch cluster that Genesys Info Mart uses to retrieve data from an Elasticsearch database version 5.0 or higher. Genesys Info Mart uses the REST API client to communicate with the Elasticsearch cluster. You must specify the REST API URL address(es) for the REST client in the following format:

- `rest(http://<es-node>:<port>[,http://<es-node>:<port>]*)`

g:index-interval

Default Value: No default value
Valid Values: Duration in days or ISO8601 duration format
Changes Take Effect: On the next ETL cycle
Dependencies: None
Introduced: 8.5.013.06

Specifies the maximum expected interval, or range of data, stored in a single Elasticsearch index. The option enables you to override the default Elasticsearch index interval.

If the data source uses indices that can have different intervals, set the value of this option to the largest possible interval. For example, if the data source uses an index that sometimes contains three days of data and sometimes contains five days of data, set `g:index-interval=5`.

g:tenant-prefix

Default Value: No default value
Valid Values: A string identifying the tenant on a shared Elasticsearch cluster
Changes Take Effect: On the next ETL cycle
Dependencies: None
Introduced: 8.5.011.15

In Genesys Engage cloud deployments, the option defines a cloud tenant prefix for Elasticsearch indexes on an Elasticsearch cluster shared across multiple cloud tenants. The tenant prefix enables Genesys Info Mart to identify Elasticsearch indexes related to the particular cloud tenant.

If specified, the option value overrides the **index-pattern** and **index-regexp** values from the XML source metadata, and the tenant prefix is included in index pattern and regexp strings.

Example

The following table illustrates the effect of specifying a tenant prefix, where the source type is sdr and the source ID is sdr0.

[elasticsearch-sdr0].g:tenant-prefix	index-pattern	index-regexp
Not defined	'sdr'-yyyy.MM.dd	sdr-*
-my-tenant	'sdr-my-tenant'-yyyy.MM.dd	sdr-my-tenant-*

rest:max-content-length

Default Value: 100 MB
Valid Values: A positive integer, with or without a suffix specifying the unit of measure (B|KB|MB|GB)
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None
Introduced: 8.5.014.26

Specifies the maximum size of the buffer holding the response to an Elasticsearch REST request. If the body of the response exceeds **max-content-length**, the HTTP Client throws `org.apache.http.ContentTooLongException`.

When you define the buffer size, use one of the following suffixes to specify the unit of measure:

- B = bytes

- KB = kilobytes

- MB = megabytes

- GB = gigabytes

- If you do not specify a unit suffix, bytes are used.

sources:extra

Default Value: None
Valid Values: A comma-separated list of any valid identifiers (IDs) of the data sources for the Elasticsearch database
Changes Take Effect: On the next ETL cycle
Dependencies: None
Introduced: 8.5.011.14

In an environment where a single Elasticsearch database stores data from multiple sources, this option enables you to configure a number of data sources in a single section. Specify the ID of the main data source as the suffix in the section name, such as `sdr1` in `elasticsearch-sdr1`. Specify IDs of any additional data sources as a comma-separated list in the option value. Do not list an ID for the same data source more than once, whether explicitly or implicitly. For example, for Genesys Info Mart to retrieve the data that sdr1, ocs1, and ocs2 data sources store on node1 of the ElasticSearch cluster, add the following configuration:

Example:

```
[elasticsearch-sdr1]
```

```
sources:extra=ocs1,ocs2

client=rest(http://node1:9200)
```

When multiple `elasticsearch-*` sections are configured, ensure unique reference to a given data source ID across all the sections, to avoid an `InvalidConfiguration` exception and a failure of the ElasticSearch step and transformation job. For instance, specifying the sdr2 data source twice, as in the example below, one time in a dedicated section and the other as an additional data source, is treated as incorrect configuration.

Incorrect configuration:

```
[elasticsearch-sdr1]

sources:extra=sdr2

client=rest(http://node1:9200)

[elasticsearch-sdr2]

client=rest(http://node2:9200)
```

error-policy Section

- error-policy-call-mergecall-missing
- error-policy-campaign-group-missing
- error-policy-cfg-check-backup-data-source
- error-policy-ipurpose-numberformat
- error-policy-irf-exception

- error-policy-irf-exception-resumable
- error-policy-islink-dangling
- error-policy-islink-multiple-sources
- error-policy-islink-multiple-targets
- error-policy-islink-multiple-vertices

- error-policy-islink-source-party-missing
- error-policy-party-created-duplicated
- error-policy-party-created-missing
- error-policy-party-parent-missing

Use this configuration section to specify options that are related to error handling during transformation.

> ## Important
>
> - By default, all of the **[error-policy]** options — except for **error-policy-irf-exception, error-policy-irf-exception-resumable, error-policy-campaign-group-missing,** and **error-policy-cfg-check-backup-data-source** — are set so as not to generate an exception when the transformation job encounters data inconsistencies. The default settings mean that Genesys Info Mart will attempt to recover from inconsistencies in the source data and continue processing. The implications for data quality depend on the particular call flow and environment. The STATUS field in the INTERACTION_FACT record indicates the type of error that was encountered.
>
> - The default value for error-policy-irf-exception is log_db_resume. If you set the value of this option to exception, Genesys Info Mart will fail the transformation job when it encounters error-policy exceptions. You can set an alarm on the log event that is generated when the job fails.

error-policy-call-mergecall-missing

Default Value: resume
Valid Values: exception, resume
Changes Take Effect: On the next ETL cycle
Dependencies: None

Policy on handling the situation when the MERGECALLID field in the GIDB_G_CALL_V table refers to missing records in the table.

- **exception**—Instructs the transformation logic to interrupt transformation of the interaction with an exception, which is handled as specified by the error-policy-irf-exception option.

- **resume**—Instructs the transformation logic to ignore references to the missing data and continue with transformation. The transformation job logs the following error message: Interaction(...):call(...): merge call(...) is missing.

error-policy-campaign-group-missing

Default Value: exception
Valid Values: exception, resume
Changes Take Effect: On the next ETL cycle
Dependencies: None

Policy on handling the situation when an Outbound Contact campaign record refers to a campaign group, but group records that have the referenced GROUPID do not exist.

- **exception**—Instructs the transformation logic to fail the job.

- **resume**—Instructs the transformation logic to ignore the missing data and continue processing. In all campaign-related records that are associated with the missing group(s), the tenant is identified as unknown (the TENANT_KEY field in campaign-related fact tables is populated with -1).

error-policy-cfg-check-backup-data-source

Default Value: error
Valid Values: error, warning
Changes Take Effect: At the next configuration check
Dependencies: None
Introduced: 8.5.014.26

Specifies the severity level of the error Genesys Info Mart generates in high availability (HA) deployments if it detects that ICON connections to the HA pair(s) of data sources (T-Server, Outbound Contact Server, or Interaction Server) have not been configured correctly. Misconfiguration occurs when the primary data-source application is not specified for both ICON connections in the HA pair.

When misconfiguration is detected, Genesys Info Mart behavior depends on the defined severity level:

- warning — Error message 55-20162 (GIM_ETL_CFG_OBJ_INVALID) is logged, describing the illegal connection.

- error — The configuration check fails, preventing the ETL from running. Two error messages are logged: 55-20162 (GIM_ETL_CFG_OBJ_INVALID), which describes the illegal connection, and 55-20037 (GIM_ETL_CONFIG_CHECK_FAILED), which indicates that the configuration check failed.

error-policy-ipurpose-numberformat

Default Value: resume
Valid Values: exception, resume
Changes Take Effect: On the next ETL cycle
Dependencies: None

Policy on handling the situation when the IPurpose attached data key-value-pair (KVP) is present and the value of IPurpose is not a number. The error usually arises because of incorrect configuration.

- **exception**—Instructs the transformation logic to interrupt transformation of the interaction with an exception, which is handled as specified by the error-policy-irf-exception option.

- **resume**—Instructs the transformation logic to process the data as if the IPurpose KVP were not attached. In this case, whether the IVR is treated as a handling resource or a mediation resource depends on the value that is configured for the default-ivr-to-self-service option.

error-policy-irf-exception

Default Value: log_db_resume
Valid Values: log_db_resume, resume, exception
Changes Take Effect: On the next ETL cycle
Dependencies: None

Policy on handling the situation when an exception is encountered during transformation of some interaction thread.

- **log_db_resume** (default)—Instructs the transformation logic to discard the problematic interaction thread, write corresponding information into the STG_TRANSFORM_DISCARDS table, and resume processing.

- **resume**—Instructs the transformation logic to discard the problematic interaction thread and resume processing, without writing corresponding information into the database.

- **exception**—Instructs the transformation logic to fail the job.

error-policy-irf-exception-resumable

Default Value: Exception
Valid Values: Any valid Java regular expression
Changes Take Effect: On the next ETL cycle
Dependencies: error-policy-irf-exception=log_db_resume or resume

The value defines a filter, which enables you to fine-tune the job level behavior (as specified by the **error-policy-irf-exception** option) by controlling which exceptions that might be triggered during interaction transformation can be considered to be discardable. If the specified regular expression matches the name of the exception class or the name of the exception super classes, then the exception is considered to be noncritical; the results of the interaction transformation (IRFs and MSFs) will be discarded, but Job_TransformGIM will continue. If the specified regular expression does not match the name of the exception class or the exception super class, the job will be aborted.

For example, if **error-policy-irf-exception**=log_db_resume or resume, **error-policy-call-mergecall-missing**=exception, and the transformation job encounters that particular data inconsistency, the transformation job will generate an InteractionTransformException. If **error-policy-irf-exception-resumable** is set to:

- Exception—Genesys Info Mart will behave as described for **error-policy-irf-exception**= log_db_resume or resume.

- InteractionTransformException—Genesys Info Mart will behave as described for **error-policy-irf-exception**=log_db_resume or resume.

- NullPointerException—The transformation job will fail.

- IllegalStateException—The transformation job will fail.

error-policy-islink-dangling

Default Value: resume
 Valid Values: exception, resume
 Changes Take Effect: On the next ETL cycle
 Dependencies: None

Policy on handling the situation when information for only one side of an IS_LINK is available.

- **exception**—Instructs the transformation logic to interrupt transformation of the interaction with an exception, which is handled as specified by the error-policy-irf-exception option.

- **resume**—Instructs the transformation logic to process the interaction as if the missing IS-Link information were for a remote site that is not monitored by ICON. For example, an internal transfer will be transformed as an inbound or outbound interaction.

error-policy-islink-multiple-sources

Default Value: resume
 Valid Values: exception, resume
 Changes Take Effect: On the next ETL cycle
 Dependencies: None

Policy on handling the situation when there are multiple (>1) source IS_LINKs that have the same LINKID.

- **exception**—Instructs the transformation logic to interrupt transformation of the interaction with an exception, which is handled as specified by the error-policy-irf-exception option.

- **resume**—Instructs the transformation logic to choose one of the source records randomly and ignore the other source records.

error-policy-islink-multiple-targets

Default Value: resume
 Valid Values: exception, resume
 Changes Take Effect: On the next ETL cycle
 Dependencies: None

Policy on handling the situation when there are multiple (>1) target IS_LINKs that have the same LINKID.

- **exception**—Instructs the transformation logic to interrupt transformation of the interaction with an exception, which is handled as specified by the error-policy-irf-exception option.

- **resume**—Instructs the transformation logic to choose one of the target records randomly and ignore the other target records.

error-policy-islink-multiple-vertices

Default Value: resume
 Valid Values: exception, resume
 Changes Take Effect: On the next ETL cycle
 Dependencies: None

Policy on handling the situation when there are more than two bidirectional IS_LINKs that have the same LINKID. The option is similar to error-policy-islink-multiple-targets and error-policy-islink-multiple-sources, but it applies to bidirectional links. This data inconsistency occasionally occurs with older T-Servers.

error-policy-islink-source-party-missing

Default Value: resume
Valid Values: exception, resume
Changes Take Effect: On the next ETL cycle
Dependencies: None

Policy on handling the situation when the source call for the IS_LINK for a dial-out attempt does not have a remote dialed party. As a result, the transformation job does not have sufficient information to build the order for Interaction Resource Facts (IRFs).

- **exception**—Instructs the transformation logic to interrupt transformation of the interaction with an exception, which is handled as specified by the error-policy-irf-exception option.

- **resume**—Instructs the transformation logic to build the order for IRFs randomly as it processes the interaction.

error-policy-party-created-duplicated

Default Value: resume
Valid Values: exception, resume
Changes Take Effect: On the next ETL cycle
Dependencies: None

Policy on handling the situation when G_PARTY_HISTORY contains multiple records that have ChangeType=1(party_created) for some party.

- **exception**—Instructs the transformation logic to interrupt transformation of the interaction with an exception, which is handled as specified by the error-policy-irf-exception option.

- **resume**—Instructs the transformation logic to treat the first record that it reads as the party created record and to ignore the other party created records.

error-policy-party-created-missing

Default Value: resume
Valid Values: exception, resume
Changes Take Effect: On the next ETL cycle
Dependencies: None

Policy on handling the situation when G_PARTY_HISTORY does not contain a record that has ChangeType=1(party_created) for some party.

- **exception**—Instructs the transformation logic to interrupt transformation of the interaction with an exception, which is handled as specified by the error-policy-irf-exception option.

- **resume**—Instructs the transformation logic to construct a party created record, based on assumptions from the first party history record that it reads.

error-policy-party-parent-missing

Default Value: resume
Valid Values: exception, resume
Changes Take Effect: On the next ETL cycle
Dependencies: None

Policy on handling the situation when the party refers to a parent, but party records that have the referenced PARTYID do not exist.

- **exception**—Instructs the transformation logic to interrupt transformation of the interaction with an

exception, which is handled as specified by the error-policy-irf-exception option.

- **resume**—Instructs the transformation logic to ignore the missing data and continue processing.

gim-etl Section

- aggregation-engine-class-name
- days-to-keep-active-facts
- days-to-keep-cfg-facts
- days-to-keep-deleted-annex
- days-to-keep-discards-and-job-history
- days-to-keep-gdpr-history
- days-to-keep-gidb-facts
- days-to-keep-gim-facts
- delayed-data-threshold
- etl-start-date
- extract-data-cfg-facts-chunk-size
- extract-data-chunk-size
- extract-data-max-conn
- extract-data-stuck-threshold
- extract-data-thread-pool-size

- extract-last-second
- link-msf-userdata-mm
- link-msf-userdata-voice
- max-call-duration
- max-camp-group-session-duration-in-hours
- max-camp-group-state-duration-in-hours
- max-chain-processing-duration-in-hours
- max-chunks-per-job
- max-msfs-per-irf
- max-parties-per-call
- max-session-duration-in-hours
- max-state-duration
- max-thread-duration-after-inactive-in-days
- max-time-deviation

- memory-threshold
- merge-chunk-size
- merge-failed-is-link-timeout
- partitioning-ahead-range
- partitioning-interval-size-gidb
- partitioning-interval-size-gidb-mm
- partitioning-interval-size-gidb-ocs
- partitioning-interval-size-gim
- purge-thread-pool-size
- purge-transaction-size
- q-answer-threshold-voice
- q-short-abandoned-threshold-voice
- short-abandoned-threshold
- sm-resource-state-priority
- user-event-data-timeout

Use this configuration section to set general options.

aggregation-engine-class-name

Default Value: none
Valid Values: Any string
Changes Take Effect: On restart
Dependencies: None

Specifies the class name of the aggregation package, if it is installed. If your deployment uses the Genesys historical reporting presentation layer—Genesys CX Insights (GCXI)—or the separately installed Reporting and Analytics Aggregates (RAA) package, specify the following value: "GIMAgg.GimInterfaceImpl.AggregationImpl"

For more information, see the *Reporting and Analytics Aggregates Deployment Guide.*

days-to-keep-active-facts

Default Value: 30
Valid Values: Any positive integer
Changes Take Effect: At the next run of Job_TransformGIM or Job_MaintainGIM
Dependencies: None in releases 8.1.0 and 8.1.1; starting with release 8.1.2, days-to-keep-gim-facts
Modified: 8.5.015.19 (plays no role in determining purge thresholds); 8.1.2 (dependencies introduced); 8.1.1 (behavior and default value changed)

Starting with release 8.1.1, 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 artificial termination.

The following description applies to behavior starting with release 8.1.1. For the benefit of Genesys Info Mart 8.1.0 customers, Appendix E in the *Genesys Info Mart 8.1 Deployment Guide* provides information about purge behavior in previous 8.x releases.

Artificial Termination of Long-Living Interactions

This functionality is intended primarily to prevent problems in deployments that include aggregation—for example, when updates to long-living multimedia interactions trigger unnecessary re-aggregation, which results in overflow errors.

If an interaction is still active when the **days-to-keep-active-facts** period expires, the transformation job terminates the interaction artificially in fact tables. Any activity that is related to this interaction that occurs after termination is discarded.

Genesys Info Mart generates log message 55-20123 when it terminates an interaction artificially. You can set an alarm on this message.

Purging

Starting with release 8.5.015.19, this option plays no role in determining purge thresholds. In releases earlier than 8.5.015.19, the relationship of **days-to-keep-active-facts** to **days-to-keep-gidb-facts** affects the purge threshold for multimedia fact data in GIDB. In releases earlier than 8.5.015.19:

- If **days-to-keep-active-facts** is smaller than **days-to-keep-gidb-facts**, Genesys Info Mart terminates the interaction artificially when the **days-to-keep-active-facts** interval expires, then subsequently purges the terminated interaction from GIDB when the **days-to-keep-gidb-facts** interval expires and from the dimensional model when the **days-to-keep-gim-facts** interval expires.

- If **days-to-keep-active-facts** is greater than or equal to **days-to-keep-gidb-facts** but smaller than **days-to-keep-gim-facts**, Genesys Info Mart terminates the interaction artificially when the **days-to-keep-active-facts** interval expires, then purges the terminated interaction from GIDB in the next run of the maintenance job and from the dimensional model when the **days-to-keep-gim-facts** interval expires.

Starting with release 8.1.2, Genesys Info Mart requires **days-to-keep-active-facts** to be smaller than **days-to-keep-gim-facts**. For earlier releases, Genesys strongly recommends that you configure **days-to-keep-active-facts** to be smaller than **days-to-keep-gim-facts**.

days-to-keep-cfg-facts

Default Value: 400 or days-to-keep-gim-facts
Valid Values: Greater than, or equal to, days-to-keep-gim-facts
Changes Take Effect: At the next run of Job_MaintainGIM
Dependencies: days-to-keep-gim-facts
Introduced: 8.5.003

Specifies the number of days to retain deleted configuration fact data in GIDB and relevant fact tables. Facts that have an end time that is earlier than the retention period are eligible to be purged. The value of the **days-to-keep-cfg-facts** configuration options must always be greater than, or equal to, **days-to-keep-gim-facts**.

For a list of the tables for which this option controls the purge threshold, see Info Mart Tables Purged by the Maintenance Job in the *Genesys Info Mart Operations Guide.*

days-to-keep-deleted-annex

Default Value: 2
Valid Values: Any positive integer
Changes Take Effect: At the next run of Job_MaintainGIM
Dependencies: None

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

For example, with the default value (2 days), if Job_MaintainGIM is running on August 4, 2014, the job will purge *_ANNEX records that terminated (in other words, the configuration setting on the object's **Annex** tab was deleted) before August 2, 2014.

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.

days-to-keep-discards-and-job-history

Default Value: 600
Valid Values: Any positive integer
Changes Take Effect: At the next run of Job_MaintainGIM
Dependencies: None

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 a list of the tables for which this option controls the purge threshold, see Info Mart Tables Purged by the Maintenance Job in the *Genesys Info Mart Operations Guide.*

For example, if Job_MaintainGIM is running on August 23, 2011 (day 235 of the year) and **days-to-keep-discards-and-job-history**=600, Job_MaintainGIM will purge all records in the discard, audit, and history tables that were written by instances of the ETL jobs that ran before January 1, 2010 (day 1 of the previous year).

days-to-keep-gdpr-history

Default Value: 15
Valid Values: An integer in the range 0-30
Changes Take Effect: At the next run of Job_MaintainGIM
Dependencies: None
Introduced: 8.5.010

Specifies the number of days to retain data in the CTL_GDPR_HISTORY table.

The table stores the actual personally identifiable information (PII) that relates to a General Data Protection Regulation (GDPR) request. For Right of Access (export) requests, the table stores the data that was requested for export. For Right of Erasure ("forget me") requests, the table stores the data that was redacted.

To ensure that this data is ephemeral, the retention period for records in the CTL_GDPR_HISTORY table is restricted to a maximum of 30 days.

days-to-keep-gidb-facts

Default Value: 14
Valid Values: Any positive integer
Changes Take Effect: At the next run of Job_MaintainGIM
Dependencies: None in releases 8.1.0 and 8.1.1; starting with release 8.1.2, days-to-keep-gim-facts
Modified: 8.5.015.19, 8.1.2 (behavior changed with respect to multimedia interactions)
Related Options: days-to-keep-active-facts

Starting with release 8.1.1, 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.

Starting with release 8.5.015.19, the retention period for multimedia facts in GIDB is determined solely by the value of **days-to-keep-gidb-facts**. Be aware that, if **days-to-keep-gidb-facts** is less than **days-to-keep-active-facts**, user data for active multimedia interactions older than **days-to-keep-gidb-facts** might not be transformed correctly.

For multimedia interactions in releases earlier than 8.5.015.19, this option, together with **days-to-keep-active-facts**, specifies the retention period for facts in GIDB: The retention period is the greater of the **days-to-keep-gidb-facts** and **days-to-keep-active-facts** option values.

In 8.1.1 releases, **days-to-keep-gidb-facts** specifies the retention period for completed multimedia facts, provided that there are no active facts that have the same (or earlier) start timestamps. If there are active facts with the same (or earlier) start timestamps, the eligible completed facts will not actually be purged until these active interactions have been terminated and, therefore, also become

eligible to be purged, as described in **days-to-keep-active-facts**.

> ### Important
>
> Genesys expects the value of **days-to-keep-gidb-facts** to be less than 30. Be aware that, 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.

For a list of the tables for which this option controls the purge threshold, see Info Mart Tables Purged by the Maintenance Job in the *Genesys Info Mart Operations Guide.*

days-to-keep-gim-facts

Default Value: 400
Valid Values: Any positive integer
Changes Take Effect: At the next run of Job_MaintainGIM
Dependencies: None

Starting with release 8.1.1, 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).

Starting with release 8.1.2, Genesys Info Mart enforces **days-to-keep-active-facts** < **days-to-keep-gim-facts** and **days-to-keep-gidb-facts** < **days-to-keep-gim-facts**.

For voice and multimedia interactions in release 8.1.1, this option specifies the retention period for completed facts, provided that there are no active facts that have the same (or earlier) start timestamps. If there are active facts with the same (or earlier) start timestamps, the eligible completed facts will not actually be purged until these active interactions have been terminated and, therefore, also become eligible to be purged, as described in **days-to-keep-active-facts**.

For a list of the tables for which this option controls the purge threshold, see Info Mart Tables Purged by the Maintenance Job in the *Genesys Info Mart Operations Guide.*

> ### Important
>
> - 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. In release 8.1.1, as the **Dimensional Model (releases earlier than 8.1.2)** example shows, active multimedia interactions can delay the purging of completed interactions if the period of time for which multimedia interactions can remain active in your deployment (**days-to-keep-active-facts**) is greater than **days-to-keep-gim-facts**.
  Starting with release 8.1.2, Genesys Info Mart no longer permits this configuration. In 8.1.1 deployments, Genesys strongly recommends against this configuration, particularly in deployments that include RAA.

How Purge Options Work

To illustrate how various **days-to-keep-*** options combine to determine when data will be purged, the following tables provide examples of different scenarios for GIDB and the dimensional model, respectively. There are separate examples for releases before and after 8.1.2, because of significant behavioral differences.

- Purge examples for 8.1.2 and later releases
- Purge examples for releases earlier than 8.1.2

Purge examples for 8.1.2 and later releases

GIDB (release 8.1.2 and later)

Scenario	Start Time of Facts to Be Purged
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: **days-to-keep-gidb-facts** threshold is November 21, 2020 **days-to-keep-active-facts** threshold is November 5, 2020	Voice interactions, agent activity, Outbound Contact facts: November 21, 2020, or earlier Multimedia interactions: - Starting with release 8.5.015.19, November 21, 2020, or earlier - In releases earlier than 8.5.015.19, November 5, 2020, or earlier

Dimensional Model (release 8.1.2 and later)

Scenario	Start Time of Facts to Be Purged
days-to-keep-gim-facts=400. **days-to-keep-active-facts**=30.	All facts: November 1, 2019, or earlier

Scenario	Start Time of Facts to Be Purged

• Job_MaintainGIM is running on December 4, 2020, after the last ETL cycle for the day. In other words:
 	
  • **days-to-keep-gim-facts** threshold is November 1, 2019
 	
  • **days-to-keep-active-facts** threshold is November 5, 2020
	The oldest timestamp for artificially terminated multimedia interactions is November 5, 2020.

Purge examples for releases earlier than 8.1.2

GIDB (releases earlier than 8.1.2)

Scenario		Start Time of Facts to Be Purged

• **days-to-keep-gidb-facts**=14.
• **days-to-keep-active-facts**=30.
• Job_MaintainGIM is running on June 4, 2012, after the last ETL cycle for the day. In other words:
  • **days-to-keep-gidb-facts** threshold is May 21, 2012
  • **days-to-keep-active-facts** threshold is May 5, 2012		

1	A previously active e-mail interaction, which started on May 5, 2012, was artificially terminated on June 4, 2012. Other e-mail interactions, which started on May 6, 2012, and later, remain active.	Voice interactions, agent activity, Outbound Contact facts: May 21, 2012, or earlier Multimedia interactions: May 5, 2012, or earlier
2	The earliest active fact relates to an e-mail interaction that started on May 27, 2012.	All facts: May 5, 2012, or earlier

Dimensional Model (releases earlier than 8.1.2)

Scenario	Start Time of Facts to Be Purged

• **days-to-keep-gim-facts**=400.
• **days-to-keep-active-facts**=30.
• Job_MaintainGIM is running on June 4, 2012, after the last ETL cycle for the day. In other words:
  • **days-to-keep-gim-facts** threshold is May 1, 2011
  • **days-to-keep-active-facts** threshold is May 5, 2012	

Scenario		Start Time of Facts to Be Purged
1	The oldest previously active e-mail interaction, which started on May 5, 2012, was artificially terminated on June 4, 2012.	All facts: May 1, 2011, or earlier

	Important
	 The following configuration is not supported starting with release 8.1.2 and is strongly discouraged in earlier releases.
	
	 • **days-to-keep-gim-facts**=400.
	 • **days-to-keep-active-facts**=600.
	 • Job_MaintainGIM is running on June 4, 2012, after the last ETL cycle for the day. In other words:
	      • **days-to-keep-gim-facts threshold** is May 1, 2011
	      • **days-to-keep-active-facts** threshold is October 12, 2010	

1	The oldest previously active e-mail interaction, which started on October 12, 2010, was artificially terminated on June 4, 2012.	Agent activity, Outbound Contact facts: May 1, 2011, or earlier Voice and multimedia interactions: October 12, 2010, or earlier
2	The earliest active fact relates to an e-mail interaction that started on June 20, 2011.	All facts: May 1, 2011, or earlier

delayed-data-threshold

Default Value: 900
Valid Values: 60-3600
Changes Take Effect: On the next ETL cycle
Dependencies: None

When expected data is delayed, specifies the amount of time, in seconds, before Genesys Info Mart logs message 55-20110. The log message includes detailed information about the data sources and IDB tables from which data was expected.

Starting with release 8.1.3, Genesys Info Mart applies the delayed-data-threshold to extraction delays, when there is no data available to be extracted from an active data source. In previous releases, Genesys Info Mart applied the delayed-data-threshold to transformation delays, when Genesys Info Mart was unable to process a data chunk because dependent data was not available (in other words, had not been extracted yet).

Genesys recommends that you set an alarm on log message 55-20110, so that you can investigate the reasons for data delays in a timely manner and take appropriate action.

For more information about appropriate actions, see the Troubleshooting information in the *Genesys*

Info Mart Operations Guide.

etl-start-date

Default Value: Initialization date minus 30 days
Valid Values: Any date after 1970 in the format yyyy-mm-dd hh:mm:ss
Changes Take Effect: On database initialization
Dependencies: None
Introduced: 8.1.103.07
Modified: 8.1.4

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, which was introduced in release 8.1.103.07, 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.

Starting with release 8.1.4, 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**.

extract-data-cfg-facts-chunk-size

Default Value: 90000
Valid Values: Any integer
Changes Take Effect: On the next ETL cycle for the DAP with role=ICON_CFG
Dependencies: None

Specifies the size of the time interval, in seconds, for which configuration relationship data is committed in one transaction. The data from this period is considered to be one data chunk, for the purpose of extract and transform.

For example, if you set the value of this option to 90000, Genesys Info Mart extracts 25 hours (90,000 seconds) of available configuration relationship fact data. When any nonpositive value is set (for example, 0 or -1), data for all available time intervals is extracted in one chunk.

> Important

Nonpositive values should not be used in production, but they can be useful for lab testing.

This option enables you to configure a larger extraction window for configuration relationship data than for other types of data. This increases the likelihood that all available configuration relationship data will be extracted and transformed in one extraction cycle, so that the configuration relationship data is available to support the transformation of other data.

This option does not affect extraction and transformation of configuration object data. All available configuration object data is extracted in a single extraction cycle.

extract-data-chunk-size

Default Value: 900
Valid Values: Any integer
Changes Take Effect: On the next ETL cycle
Dependencies: None

Specifies the size of the time interval, in seconds, for which data is committed in one transaction. The data from this period is considered to be one data chunk, for the purpose of extract and transform. For example, if you set the value of extract-data-chunk-size to 900, Genesys Info Mart extracts 15 minutes (900 seconds) of available data.

The extraction job processes only one chunk of data during each extraction cycle. In other words, the value of this option sets the batch size for an iteration of the ETL cycle.

extract-data-max-conn

Default Value: 128
Valid Values: Any positive integer
Changes Take Effect: On the next ETL cycle
Dependencies: None
Discontinued: 8.5.014.26

For each DAP that Genesys Info Mart uses to extract data, specifies the maximum size of the connection pool for connections between the Genesys Info Mart Server and the IDB. Ensure that you configure your RDBMS to handle a sufficient number of concurrent connections.

For example, if Genesys Info Mart uses DAP1 to access IDB1 and DAP2 to access IDB2, and if **extract-data-max-conn**=128, then the extraction job will open up to 128 connections through DAP1 and up to 128 connections through DAP2 (in other words, up to a total of 256 connections), as required, for concurrent extraction of data from both IDBs.

Increasing the value of this option reduces the amount of time that is required for the extraction job, but it increases CPU and memory requirements, especially if database links are not used. The optimal value of this option depends on the operating system, hardware (such as RAM and the number of CPUs), and the number of IDBs in the environment. You must also consider ICON requirements for RDBMS resources when you set this limit.

extract-data-stuck-threshold

Default Value: 28860
Valid Values: Any integer equal to or greater than 1800
Changes Take Effect: On the next ETL cycle
Dependencies: None

Specifies the time, in seconds, that Genesys Info Mart waits for stuck data to become available in IDB. For example, if new data is available until 3:00 PM and the value for extract-data-stuck-threshold is set to 28860 (8 hrs 1 min.), Genesys Info Mart extracts data that has a timestamp from 6:59 AM to 3:00 PM. IDB data from the time intervals before 6:59 AM is not, and will never be, extracted.

The value of this option sets the stuck data timeout for the extraction windows for the Voice, Multimedia, and Outbound Contact data domains. However, the actual time ranges of the extraction windows and HWMs for each data domain might be different.

Setting this parameter to too high a value can affect performance and cause delay of data that is available for reporting. For example, if data from a T-Server is not available—for example, because of network delays in a multi-site deployment or, even in a single-site deployment, if there is a problem within ICON so that, say, user data from the gud provider is behind call data from the gcc provider—Genesys Info Mart does not transform data until data from that T-Server becomes available or until this timeout has expired.

On the other hand, setting this parameter to too low a value might cause data to be skipped if, after the stuck data threshold expires, the extraction high-water mark (HWM) advances past data that was delayed.

While delayed data from an active data source does not delay extraction of data from other data sources in that domain, missing data from a delayed data source can delay transformation. It is important, therefore, to identify and fix extraction delays as soon as possible. See the description of the delayed-data-threshold option for related information.

extract-data-thread-pool-size

Default Value: 32
Valid Values: Any positive integer, as appropriate for your environment
Changes Take Effect: On the next ETL cycle
Dependencies: None

Specifies the maximum number of worker threads that are used to extract data concurrently. This option does not set a strong limit on the total number of threads that will be used in extraction processing, because certain extraction algorithms create additional helper threads. Instead, this option specifies the maximum number of logical partitions for concurrent extraction of subsets of data.

Increasing the value of this option reduces the amount of time that is required for the extraction job, but it increases CPU and memory requirements. (At a rough estimate, each additional worker thread requires an additional 180 MB of RAM.) The optimal value of this option depends on the operating system, hardware (such as RAM and the number of CPUs), and the number of IDBs in the environment.

extract-last-second

Default Value: false
Valid Values: true, false
Changes Take Effect: On the next ETL cycle
Dependencies: None
Discontinued: 8.1.4

Warning! This option must be set to false in a production environment.

Specifies whether Genesys Info Mart extracts data from the last second of the specified time interval into the extracted data chunk. In a lab environment, setting this option to true speeds up data validation. If you are extracting the last second of data in a lab environment, make sure that all data from the executed test scenario is stored into IDB before Genesys Info Mart runs the extraction job.

In a production environment, setting this option to true might result in lost data. For example, the ETL might be extracting data at time t1 for the time interval [t0–t1], while the data source is still producing events that have a timestamp of t1; at the next extraction, Genesys Info Mart will consider that all data that has t1 timestamps has been extracted, and the last t1 data will be lost.

The option was discontinued in release 8.1.4 because of issues with missing data. To speed up data validation in a lab environment, you can achieve a similar result by setting the ICON NoData timestamp (specified by the ICON **dss-no-data-tout** option) to a small value, thus minimizing the processing delays that Genesys Info Mart provides to account for data-source inactivity.

link-msf-userdata-mm

Default Value: false
Valid Values: true, false
Changes Take Effect: On the next ETL cycle
Introduced: 8.5.003
Related Options: link-msf-userdata (DN) and link-msf-userdata (Script)

Specifies whether associated user data will be stored in mediation segment facts (MSFs) for multimedia interactions that are in mediation.

- **true**—MSFs for multimedia interactions will store associated user data.

- **false**—MSFs for multimedia interactions will not store associated user data.

Setting this option to true enables user data storage for all mediation resources for multimedia interactions. The value of the **link-msf-userdata** option, if specified on individual DN or Script objects, overrides this application-level **link-msf-userdata-mm** option.

Important

Setting a single option to control user data storage for all multimedia mediation resources simplifies configuration, but processing and storing excessive amounts of user data can degrade performance. Therefore, Genesys recommends that you enable the **link-msf-userdata** configuration option at the DN level or Script level, for the

virtual queues or interaction queues or workbins that are significant for your reporting, instead of enabling **link-msf-userdata-mm** at the application level.

link-msf-userdata-voice

Default Value: false
Valid Values: true, false
Changes Take Effect: On the next ETL cycle
Dependencies: None
Introduced: 8.5.003
Related Options: link-msf-userdata (DN)

Specifies whether associated user data will be stored in mediation segment facts (MSFs) for voice interactions that are in mediation.

- **true**—MSFs for voice interactions will store associated user data.

- **false**—MSFs for voice interactions will not store associated user data.

Setting this option to t rue enables user data storage for all mediation resources for voice interactions. The value of the **link-msf-userdata** option, if specified on individual DN objects, overrides this application-level **link-msf-userdata-voice** option.

Important

Setting a single option to control user data storage for all voice mediation resources simplifies configuration, but processing and storing excessive amounts of user data can degrade performance. Therefore, Genesys recommends that you enable the **link-msf-userdata** configuration option at the DN level, for the ACD or virtual queues that are significant for your reporting, instead of enabling **link-msf-userdata-voice** at the application level.

max-call-duration

Default Value: 3600
Valid Values: 0 or any positive integer
Changes Take Effect: On the next ETL cycle
Dependencies: None

Specifies the maximum duration, in seconds, of voice interactions in the deployed environment, as well as a number of important timeouts for both voice and multimedia interactions.

This option controls the following timeouts:

- The stuck-link threshold for merge processing of voice interactions — The option specifies the amount of

time that the merge operation will wait for IS-Link information from another site, before it considers the IS-Link to be stuck. An unpaired link is considered to be stuck if the link initiation timestamp in the G_IS_LINK merge table record exceeds the stuck-link threshold calculated from the earliest extraction high-water mark.

- The stuck data threshold for merge processing of voice interactions — The option determines the amount of time that the merge operation will wait for G_IR and G_CALL records. To ensure that the merge procedure does not finalize merge processing prematurely with missing data, the stuck data threshold for merge is double the stuck-link threshold (2 * **max-call-duration**).

> ### Important
>
> For voice interactions, Genesys Info Mart extracts Interaction Records (IRs) only for completed calls. Resolution of stuck calls is controlled by the ICON application.

- The limit for waiting for After Call Work (ACW) before transforming voice interaction data.

- For user data with propagation rules other than CALL, the limit for linking user data to IRFs for multimedia interactions — The option specifies the amount of time, starting from the beginning of the IRF, that the transformation job allows for user data to link to the IRF.

- In an HA deployment in which geolocation is a consideration, the amount of time that Genesys Info Mart will wait after a good ICON provider session has been re-established on the local IDB following a data disruption, before the ETL switches back to extract Voice details from the local IDB.

Because this option affects a number of aspects of Genesys Info Mart functioning during extraction and transformation of voice and multimedia interactions, be aware that there are potential data implications if you change the option value.

max-camp-group-session-duration-in-hours

Default Value: 168
Valid Values: 1-10000
Changes Take Effect: On the next ETL cycle
Dependencies: None

Specifies the amount of time, in hours, after which Genesys Info Mart ends active campaign group sessions if the transformation process encounters a campaign session row in IDB's GO_Campaign table that has no terminated time.

If you change the value of this option, the new option value is not applied to previously loaded facts.

max-camp-group-state-duration-in-hours

Default Value: 168
Valid Values: 1-10000
Changes Take Effect: On the next ETL cycle
Dependencies: None

Specifies the amount of time, in hours, after which Genesys Info Mart ends campaign group states if

the transformation process has not extracted a stopped row that brackets a previously extracted started row from the IDB GO_CampaignHistory table.

If you change the value of this option, the new option value is not applied to previously loaded facts.

max-chain-processing-duration-in-hours

Default Value: 8
Valid Values: 0 or any positive integer
Changes Take Effect: On the next ETL cycle
Dependencies: None. Applies only when the Info Mart database is partitioned.

Specifies the maximum expected duration, in hours, of the processing of a single chain by Outbound Contact Server (OCS). Chain processing starts when the chain is loaded from the OCS database and ends when the chain is unloaded. For more information, see the section about the Chain Model in the chapter about integrating with Outbound Contact in the *Interaction Concentrator 8.x User's Guide.*

The option affects Genesys Info Mart behavior only when the Info Mart database is partitioned.

On partitioned databases, Genesys Info Mart will not create CONTACT_ATTEMPT_FACT (CAF) records for activity that occurs after the maximum expected duration of chain processing. If the value of this option is 0 (zero), Genesys Info Mart will use the value of **max-camp-group-session-duration-in-hours**.

The **max-chain-processing-duration-in-hours** option was introduced in release 8.1.2, and the default value was set to be consistent with Genesys Info Mart behavior in previous 8.x releases. You can safely increase the value of this option up to the maximum length of campaigns in your deployment (the value of **max-camp-group-session-duration-in-hours**). However, increasing the value means that Genesys Info Mart will have more Outbound Contact data to process in each ETL cycle. Therefore, for performance reasons, the optimal value for **max-chain-processing-duration-in-hours** is the smallest value that matches actual patterns of activity in your deployment (in other words, the smallest value that does not result in missing CAF records).

max-chunks-per-job

Default Value: 10
Valid Values: 1-100
Changes Take Effect: On the next ETL cycle
Dependencies: None

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.

max-msfs-per-irf

Default Value: 50
Valid Values: 10-10000
Changes Take Effect: On the next ETL cycle

Dependencies: None
Introduced: 8.1.401.02

Specifies the limit for the number of MSF records that are associated with a given multimedia interaction that will be represented in the Info Mart database. When the number of MSF records associated with a single IRF record exceeds the limit, the transformation job processes only the first n mediation DNs and the last n mediation DNs in the interaction, where n = max-msfs-per-irf/2. In this way, ETL performance avoids being degraded by huge numbers of MSF records for unsuccessful routing attempts for "stuck strategy" scenarios.

The first n mediation records for a given interaction are processed by the transformation job and populated in the MSF table as they occur. The last n records are postponed until an associated IRF record has been created.

The option does not affect the mediation durations that are reported in the IRF (`QUEUE_DURATION`, `ROUTING_POINT_DURATION`, `MEDIATION_DURATION`, and `PREVIOUS_MEDIATION_DURATION`); all these metrics correctly report the full overall mediation time.

Log message number 55-20120 is generated when the **max-msfs-per-irf** limit is exceeded, triggering the behavior to abbreviate the representation of unsuccessful routing attempts in the Info Mart database. You can set an alarm on the log message, to prompt you to investigate strategies that might be inappropriate for your deployment.

max-parties-per-call

Default Value: 100
Valid Values: 50-10000
Changes Take Effect: On the next ETL cycle
Dependencies: None
Introduced: 8.1.103.03
Discontinued: 8.1.4

Specifies the limit for the number of parties or virtual queues that were associated with the same multimedia interaction that will be represented in the Info Mart database. The option limits the amount of data that will be selected for transformation from the GIDB party and virtual-queue tables. When the number of parties or virtual queues in an interaction exceeds the limit, the transformation job processes only the first n parties or virtual queues and the last n parties or virtual queues in the interaction, where n = max-parties-per-call/2. In this way, the transformation job avoids being overwhelmed by huge numbers of party and virtual-queue records for unsuccessful routing attempts for "stuck strategy" scenarios.

The option was discontinued in release 8.1.4 because improvements in incremental transformation removed the need to enforce a limit on the number of records to be read in each chunk.

The option enables you to control behavior that was not configurable before release 8.1.103.03. Log events (message numbers 55-20120, 55-20121, or 55-20122) identify when excessive numbers of Party, Party History, or Virtual Queue records, respectively, trigger the behavior to limit the number of parties or virtual queues to process. The default value of **max-parties-per-call** triggers the behavior at lower limits than in releases earlier than 8.1.103.03. In the earlier 8.1 releases, the limits were 1000 parties and 500 virtual queues. (The maximum number of Party History records is a much higher, theoretical limit.)

Genesys recommends that you set an alarm on the log messages, to prompt you to locate and fix

inappropriate routing strategies that result in huge numbers of records in IDB tables.

Starting with release 8.1.4, while Genesys Info Mart continues to abbreviate the representation of unsuccessful routing attempts in the Info Mart database, log message numbers 55-20120, 55-20121, and 55-20122 are no longer generated.

> ### Tip
>
> For all Genesys Info Mart releases, if you cannot change the behavior of inappropriate strategies, Genesys recommends that you utilize Interaction Server and Interaction Routing Designer (IRD) functionality to hide strategy activity, as described in the *Genesys Info Mart Deployment Guide*. Hiding strategy activity is the preferred method of suppressing event data that is not important for historical reporting, and will result in better data quality in Genesys Info Mart.

max-session-duration-in-hours

Default Value: 24
Valid Values: 1-10000
Changes Take Effect: On the next ETL cycle
Dependencies: None
Modified: 8.1.3 (minimum valid value changed)

Specifies the maximum duration, in hours, for a resource session in the SM_RES_SESSION_FACT table. Genesys Info Mart will end a resource session for an agent/media type after this timeout elapses if the session end has not been extracted from the IDB GX_SESSION_ENDPOINT table.

For related information, see the description of **max-state-duration**. See also the discussion about long-duration sessions or states on the Populating Agent Activity Data page in the *Genesys Info Mart User's Guide*.

If you change the value of this option, the new option value is not applied to previously loaded facts.

In releases earlier than 8.1.3, the minimum valid value of this option was 0.

max-state-duration

Default Value: 14400
Valid Values: Any integer greater than 900
Changes Take Effect: On the next ETL cycle
Dependencies: No direct dependency, but it would not be meaningful to set max-state-duration > (value of max-session-duration-in-hours expressed in seconds)
Introduced: 8.1.3

Specifies the maximum duration, in seconds, for an agent state in the SM_RES_STATE_FACT table. Genesys Info Mart ends the agent state if the IDB data does not show the agent transitioning to a different state by the time that the maximum duration (the state timeout) expires. The option enables you to:

- Improve performance of agent-activity transformation by reducing the amount of active data that Genesys Info Mart must maintain and by limiting the period of time for which lookups must be performed for agent-activity data.

- Recognize in a timely manner when an agent session has gone inactive.

The option was introduced in release 8.1.3.

In releases earlier than 8.1.3, when agents forgot to log out, the reported durations of sessions and states might be misleadingly high. While the **max-session-duration-in-hours** option puts a limit on the overall session duration, **max-state-duration** provides the ability to detect when a session has gone inactive. If a state times out and there are no other active states, Genesys Info Mart can conclude that there really is no agent activity within the session.

The default value of 14400 seconds (4 hours) for **max-state-duration** strikes a balance between considering long periods of idleness (no agent state transitions) to be normal behavior and realizing the performance benefits mentioned above.

Consider increasing the value of **max-state-duration** if it is normal for many of your agents to remain idle (no agent state transitions at all) for significant periods of time, and you do not want long periods of idleness to indicate that the agent session has gone inactive.

Consider decreasing the value of **max-state-duration** if it is not normal for your agents to remain idle for significant periods of time, and you want to recognize that a session has gone inactive after a shorter period of idleness. Additionally, decrease the value if you prefer to realize more of the performance benefits.

Particularly in large contact centers with hundreds or thousands of agents, Genesys recommends that you set the value of **max-state-duration** to the smallest value that is consistent with patterns of agent behavior and session-reporting requirements in your deployment.

For related information, see the discussion about long-duration sessions or states in the chapter about populating Genesys Info Mart data in the *Genesys Info Mart User's Guide*.

If you change the value of this option, the new option value is not applied to previously loaded facts.

max-thread-duration-after-inactive-in-days

Default Value: 31
Valid Values: Any positive integer greater than days-to-keep-active-facts
Changes Take Effect: On the next ETL cycle
Dependencies: populate-thread-facts=true
Introduced: 8.1.1

Specifies the maximum duration, in days, of an interaction thread after all of the interactions that belong to the thread have terminated. After the timeout expires, Genesys Info Mart considers the thread to be closed.

If a new, related interaction (for example, an InboundCustomerReply to an agent's OutboundReply) begins after the timeout has expired, the interaction is considered to be the root interaction for a new thread. If a new, related interaction begins before the timeout has expired, the thread remains active; Genesys Info Mart will not consider the thread to be closed until the configured amount of time passes after all interactions in the thread (including this new interaction) have terminated.

This option affects how Genesys Info Mart identifies which interactions are associated with the same threads and, therefore, affects the thread-related data that is reported in IF and IRF records. For performance reasons, Genesys recommends that you set this option to the smallest value that is consistent with interaction patterns in your deployment, so that Genesys Info Mart does not have to maintain memory of inactive threads for unnecessarily long periods of time.

This option does not affect the chat thread reporting feature introduced in release 8.5.014.09 for Genesys Engage cloud deployments with Advanced Chat.

For more information about how interaction threads are identified, see the subsection about interaction threads, in the section about populating interaction data, in the *Genesys Info Mart User's Guide* for your release.

max-time-deviation

Default Value: 30
Valid Values: 1-120
Changes Take Effect: On the next ETL cycle
Dependencies: None
Modified: 8.5.015.19 (scope extended); 8.1.4 (minimum valid value changed)

Specifies the maximum time deviation, in seconds, to take into account for time synchronization inaccuracy between host system clocks. Starting with release 8.5.015.19, this option also specifies the maximum acceptable delay the transformation job will accommodate 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*)
<= **max-time-deviation**

In HA deployments, **max-time-deviation** is used 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. For more information about the ICON NoData indicator, see the section about determining IDB availability in the *Interaction Concentrator User's Guide*.

In releases earlier than 8.1.4, the minimum valid value of this option was 0.

memory-threshold

Default Value: 0
Valid Values: 0-99
Changes Take Effect: Immediately
Dependencies: None

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 set to 0, the feature will be disabled.

merge-chunk-size

Default Value: 200000
Valid Values: Any positive integer
Changes Take Effect: On the next ETL cycle
Dependencies: None

Specifies the maximum number of root G_IR rows in a chunk of merged data, for the purpose of transformation. By limiting the size of the data chunks produced by the merge process, this option enables you to manage situations in which there is a very large amount of merged data. The optimal value of this option depends on the characteristics of your deployment.

Consider the following guidelines:

- The option value should be large enough that it does not require transformation of an unnecessarily large number of data chunks and, therefore, does not interfere with normal processing of one chunk of extracted data. As a rule of thumb, allow 200000 rows for a call rate of 100 calls per second (cps).

- The option value should be small enough that it protects the transformation job from running out of memory. For example, for an environment with 8 GB of RAM dedicated to Genesys Info Mart, allow a number of rows that corresponds to 1 million root G_IR records. (Every root G_IR record corresponds to one interaction fact.)

merge-failed-is-link-timeout

Default Value: 0
Valid Values: 0 or any positive integer
Changes Take Effect: On the next ETL cycle
Dependencies: None
Introduced: 8.1.1

Specifies the time interval, in seconds, for which the merge of failed Inter Server Call Control (ISCC) links (IS-Links) will be delayed to enable Genesys Info Mart to receive both sides of the links. If the value of this option is greater than the value of max-call-duration, which specifies the timeout for stuck links, then max-call-duration controls the timeout for the merge of failed IS-Links as well.

The default value (0), which preserves legacy behavior, means that the merge operation will process unpaired failed links immediately, without waiting for the other side of the failed IS-Link. As a result of the partial merge, the transformation job might encounter dangling links.

partitioning-ahead-range

Default Value: 14
Valid Values: Any positive integer
Changes Take Effect: At the next run of Job_MaintainGIM
Dependencies: None

Specifies, in terms of number of days, how far ahead Job_InitializeGIM (in the first instance) and Job_MaintainGIM (on an ongoing basis) will create partitions for GIDB, Control, and Info Mart fact

tables that are partitioned. Starting with release 8.1.2, 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.) In earlier releases, Job_MaintainGIM adds partitions for the number of days ahead of the extraction high-water mark (extractHWM). 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` (1 day), `partitioning-interval-size-gim=604800` (1 week), and `partitioning-ahead-range=14,` Job_MaintainGIM will create as many additional partitions as necessary to provide partitions for GIDB, Control (in release 8.1.2 and later), and Info Mart fact tables up to 14 days ahead. If Job_MaintainGIM runs daily, this means that:

- For GIDB tables, each run of Job_MaintainGIM will create one new partition of size 1 day, for the fourteenth day. (Previous runs will have created partitions for the other days.)

- For Info Mart fact tables, as well as for Control tables, the maintenance job will create 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 will create a new partition only when the last day of the partitioning ahead range falls in a week for which a partition has not yet been created. For example, if `partitioning-interval-size-gim=604800 (7 days)` but `partitioning-ahead-range=10,` two new partitions will be created on the very first run of Job_MaintainGIM, and the next partition will be created on the fifth run.

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

This option applies only in deployments that use partitioning.

partitioning-interval-size-gidb

Default Value: 86400
Valid Values: Any positive integer
Changes Take Effect: At the next run of Job_MaintainGIM
Dependencies: None

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.

partitioning-interval-size-gidb-mm

Default Value: 86400 or partitioning-interval-size-gidb
Valid Values: Any positive integer
Changes Take Effect: At the next run of Job_MaintainGIM
Dependencies: None
Introduced: 8.1.402.07

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, on the one hand, partitions are large enough that data that is being actively used is in the fewest number of partitions while, on the other hand, the partitions are small enough that their indexes can fit into the cache.

partitioning-interval-size-gidb-ocs

Default Value: 86400 or partitioning-interval-size-gidb
Valid Values: Any positive integer
Changes Take Effect: At the next run of Job_MaintainGIM
Dependencies: None
Introduced: 8.1.402.07

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, on the one hand, partitions are large enough that data that is being actively used is in the fewest number of partitions while, on the other hand, the partitions are small enough that their indexes can fit into the cache.

partitioning-interval-size-gim

Default Value: 86400
Valid Values: Any positive integer
Changes Take Effect: At the next run of Job_MaintainGIM
Dependencies: None

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. Starting with release 8.1.1, the default size of Info Mart fact table partitions is 1 day (86400 seconds). In release 8.1.0, the default size was 7 days (604800 seconds).

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.

purge-thread-pool-size

Default Value: 32
Valid Values: Any positive integer
Changes Take Effect: At the next run of Job_MaintainGIM
Dependencies: None

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.

purge-transaction-size

Default Value: 100000
Valid Values: Any positive integer
Changes Take Effect: At the next run of Job_MaintainGIM
Dependencies: None

Specifies the number of deleted records per table that will be 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 will delete and commit 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 will delete and commit 90,000 records from the first table in one transaction and 10,000 records from the second table in a separate transaction.

q-answer-threshold-voice

Default Value: 60
Valid Values: 1-10000
Changes Take Effect: On the next ETL cycle

Dependencies: None

Specifies the default duration, in seconds, that is used on all configured queues as a target time to answer voice interactions that were distributed by virtual queues or ACD queues.

Genesys Info Mart uses this value unless, in the interface you use for configuration, you configure an option that has the same name on an individual Virtual Queue or ACD Queue DN object or on the Switch object.

For more information about setting queue-specific thresholds, see Configuring a DN for ICON and Genesys Info Mart reporting in the *Genesys Info Mart Deployment Guide*.

If you change the value of this option, the new option value is not applied to previously loaded facts.

For similar options that control equivalent thresholds for multimedia interactions, see the gim-etl-media-chat Section and the gim-etl-media-email Section.

q-short-abandoned-threshold-voice

Default Value: 10
Valid Values: 1-1000
Changes Take Effect: On the next ETL cycle
Dependencies: None
Modified: 8.5.003

Specifies the maximum duration of mediation, in seconds, that is used on all configured queues to indicate that an interaction that was abandoned while in a queue should be considered a "short" abandon. Genesys Info Mart uses this value to determine the state of SHORT_ABANDONED_FLAG in the MEDIATION_SEGMENT_FACT (MSF) row for voice interactions that are abandoned in a virtual queue or ACD queue.

Genesys Info Mart uses this value unless, in the interface you use for configuration, you configure an option that has the same name on a Switch object or, starting with release 8.5.003, an individual virtual queue or ACD queue DN object. In releases earlier than 8.5.003, this value could not be set on an individual DN object.

If you change the value of this option, the new option value is not applied to previously loaded facts.

For similar options that control equivalent thresholds for multimedia interactions, see the gim-etl-media-chat Section and the gim-etl-media-email Section.

short-abandoned-threshold

Default Value: 10
Valid Values: 0-60
Changes Take Effect: On the next ETL cycle
Dependencies: None

Specifies the minimum duration, in seconds, of an abandoned voice interaction in order for it to be considered truly abandoned. Genesys Info Mart uses this value to determine the state of SHORT_ABANDONED_FLAG in the IRF row.

If you change the value of this option, the new option value is not applied to previously loaded facts.

For similar options that control equivalent thresholds for multimedia interactions, see the gim-etl-media-chat Section and the gim-etl-media-email Section.

sm-resource-state-priority

Default Value: ACW, NOT_READY, BUSY, READY
Valid Values: ACW, BUSY, NOT_READY, READY (all four in any order)
Changes Take Effect: On the next ETL cycle
Dependencies: populate-sm-resource-activity
Modified: 8.5.002

Specifies a list of the state names—BUSY, ACW, NOT_READY, or READY—in order of decreasing priority. When an agent simultaneously has different states on different DNs for a given media type, Genesys Info Mart uses this list to determine which state has the highest priority when determining a summarized state to store in the SM_RES_STATE_FACT or SM_RES_STATE_REASON_FACT table. Starting with release 8.5.002, Genesys Info Mart uses this list to evaluate the priority for media-neutral states (the agent's state across all media types), as well, if the **populate-media-neutral-sm-facts** option is set to *true*.

If you change the value of this option, you must specify all four values, in the changed order. If you do not specify all the values, Genesys Info Mart considers the option value to be invalid and uses the default value instead.

If you change the value of this option, the new option value is not applied to previously loaded facts.

> ### Important
> The list does not include the LOGGED_IN state, which always has the lowest priority.

user-event-data-timeout

Default Value: 3600
Valid Values: 0 or any positive integer
Changes Take Effect: On the next ETL cycle
Dependencies: None
Modified: 8.1.1 (behavior changed)

Specifies the maximum time, in seconds, after the end of a call, during which an agent who handled that call can send UserEvent-based key-value pair (KVP) data. If the call has ended and the UserEvent-based KVP data is sent after this timeout, the transformation job does not process the UserEvent-based KVP data.

The option value also has a role in:

- Enabling complete ACW information in IRF records.

- Starting with release 8.5.014, controlling the amount of time for which Genesys Info Mart will wait for Focus Time information for multimedia interactions after the agent leaves the interaction.

Therefore, consider your interaction topologies carefully before modifying the value of this option.

The behavior of this option changed between release 8.1.0 and release 8.1.1. In release 8.1.0, the timeout was calculated from the start of the call.

This option is not applicable to UserEvent data related to Genesys Callback.

gim-etl-media-chat Section

- q-answer-threshold
- q-short-abandoned-threshold
- short-abandoned-threshold

Media-specific configuration sections enable you to specify separate thresholds for different types of multimedia interactions (Genesys eServices interactions as well as 3rd Party Media interactions).

By default, the Genesys Info Mart application template includes the **[gim-etl-media-chat]** section and options, for specifying thresholds for Genesys eServices chat interactions.

If you want to customize thresholds for other media types, add the applicable media-specific configuration section(s) and options on the **Options** tab of the Genesys Info Mart Application object. Name the section **[gim-etl-media-<media type>],** where **<media type>** exactly matches the predefined name of the Media Type Business Attribute, including case (for example, **[gim-etl-media-fax]**).

Each media-specific configuration section can contain any of the following options.

q-answer-threshold

Default Value: 60
Valid Values: 1-600000
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None

Specifies the default duration, in seconds, that is used on all configured queues as a target time to accept a chat interaction that entered a queue. Genesys Info Mart uses the option setting in the Genesys Info Mart Application object unless you override this value by configuring an option that has the same name in the interface you use for configuration:

- On a Media Type Business Attribute for a particular tenant, for interactions of that media type for that tenant.

- On a Switch object, for all virtual queues that belong to that switch, or on an individual Virtual Queue DN object, for chat interactions that come through that virtual queue.

- On an individual Script object that corresponds to a particular Multimedia Interaction Queue or Multimedia Interaction Workbin, for chat interactions that come through an interaction queue or workbin.

q-short-abandoned-threshold

Default Value: 10
Valid Values: 1-1000
Changes Take Effect: On the next ETL cycle
Dependencies: None
Modified: 8.5.003

Specifies the maximum duration of mediation, in seconds, that is used on all configured queues to indicate that an interaction that was abandoned while in a queue should be considered a "short" abandon. Genesys Info Mart uses this value to determine the state of SHORT_ABANDONED_FLAG in the MSF row for multimedia interactions that are abandoned in a Virtual Queue, Interaction Queue, or Workbin.

Genesys Info Mart uses the option setting in the Genesys Info Mart Application object unless you override this value by configuring an option that has the same name in the interface you use for configuration:

- On a Media Type Business Attribute for a particular tenant, for interactions of that media type for that tenant.

- Starting with release 8.5.003, on a Switch object, for all virtual queues that belong to that switch, or on an individual Virtual Queue DN object, for chat interactions that come through that virtual queue.

- Starting with release 8.5.003, on an individual Script object that corresponds to a particular Multimedia Interaction Queue or Multimedia Interaction Workbin, for chat interactions that come through an interaction queue or workbin.

In releases earlier than 8.5.003, this value could not be set on an individual Switch, DN, or Script object.

If you change the value of this option, the new option value is not applied to previously loaded facts.

short-abandoned-threshold

Default Value: 10
Valid Values: 0-60
Changes Take Effect: On the next ETL cycle
Dependencies: None

Specifies the minimum duration, in seconds, of an abandoned chat interaction in order for it to be considered truly abandoned. Genesys Info Mart uses this value to determine the state of SHORT_ABANDONED_FLAG in the IRF row. Genesys Info Mart uses the option setting in the Genesys Info Mart Application object unless you override this value by configuring an option that has the same name in the interface you use for configuration:

- On a Media Type Business Attribute for a particular tenant, for interactions of that media type for that tenant.

If you change the value of this option, the new option value is not applied to previously loaded facts.

> **Tip**
> The same option in the **[gim-etl]** section controls the equivalent threshold for voice interactions.

gim-etl-media-email Section

- q-answer-threshold
- q-short-abandoned-threshold
- short-abandoned-threshold

Media-specific configuration sections enable you to specify separate thresholds for different types of multimedia interactions (Genesys eServices interactions as well as 3rd Party Media interactions).

By default, the Genesys Info Mart application template includes the **[gim-etl-media-email]** section and options, for specifying thresholds for Genesys eServices email interactions.

If you want to customize thresholds for other media types, add the applicable media-specific configuration section(s) and options on the **Options** tab of the Genesys Info Mart Application object. Name the section **[gim-etl-media-<media type>],** where **<media type>** exactly matches the predefined name of the Media Type Business Attribute, including case (for example, **[gim-etl-media-fax]**).

Each media-specific configuration section can contain any of the following options.

q-answer-threshold

Default Value: 60
Valid Values: 1-600000
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None

Specifies the default duration, in seconds, that is used on all configured queues as a target time to accept an email interaction that entered a queue. Genesys Info Mart uses the option setting in the Genesys Info Mart Application object unless you override this value by configuring an option that has the same name in the interface you use for configuration:

- On a Media Type Business Attribute for a particular tenant, for interactions of that media type for that tenant.

- On a Switch object, for all virtual queues that belong to that switch, or on an individual Virtual Queue DN object, for email interactions that come through that virtual queue.

- On an individual Script object that corresponds to a particular Multimedia Interaction Queue or Multimedia Interaction Workbin, for email interactions that come through an interaction queue or workbin.

q-short-abandoned-threshold

Default Value: 10
Valid Values: 1-1000
Changes Take Effect: On the next ETL cycle
Dependencies: None

Specifies the maximum duration of mediation, in seconds, that is used on all configured queues to indicate that an interaction that was abandoned while in a queue should be considered a "short" abandon. Genesys Info Mart uses this value to determine the state of SHORT_ABANDONED_FLAG in the MSF row for multimedia interactions that are abandoned in a Virtual Queue, Interaction Queue, or Workbin.

Genesys Info Mart uses the option setting in the Genesys Info Mart Application object unless you override this value by configuring an option that has the same name in the interface you use for configuration:

- On a Media Type Business Attribute for a particular tenant, for interactions of that media type for that tenant.

- Starting with release 8.5.003, on a Switch object, for all virtual queues that belong to that switch, or on an individual Virtual Queue DN object, for email interactions that come through that virtual queue.

- Starting with release 8.5.003, on an individual Script object that corresponds to a particular Multimedia Interaction Queue or Multimedia Interaction Workbin, for email interactions that come through that interaction queue or workbin.

In releases earlier than 8.5.003, this value could not be set on an individual Switch, DN, or Script object.

If you change the value of this option, the new option value is not applied to previously loaded facts.

short-abandoned-threshold

Default Value: 10
Valid Values: 0-60
Changes Take Effect: On the next ETL cycle
Dependencies: None

Specifies the minimum duration, in seconds, of an abandoned email interaction in order for it to be considered truly abandoned. Genesys Info Mart uses this value to determine the state of SHORT_ABANDONED_FLAG in the IRF row. Genesys Info Mart uses the option setting in the Genesys Info Mart Application object unless you override this value by configuring an option that has the same name in the interface you use for configuration:

- On a Media Type Business Attribute for a particular tenant, for interactions of that media type for that tenant.

If you change the value of this option, the new option value is not applied to previously loaded facts.

Tip

The same option in the **[gim-etl]** section controls the equivalent threshold for voice interactions.

gim-etl-populate Section

- populate-irf-asm-engage-duration
- populate-media-neutral-sm-facts
- populate-mm-ixnqueue-facts
- populate-mm-workbin-facts
- populate-sip-im-facts
- populate-sm-busy-from-mm-ixns
- populate-sm-chat-resource-activity
- populate-sm-email-resource-activity
- populate-sm-resource-session-facts
- populate-sm-resource-state-facts
- populate-sm-resource-state-reason-facts
- populate-sm-voice-resource-activity
- populate-thread-facts
- populate-workbin-as-hold

Use this configuration section to specify optional population of certain Info Mart tables and how certain interaction flows are reported.

populate-irf-asm-engage-duration

Default Value: false
Valid Values: true, false
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None
Introduced: 8.5.004

Enables or disables reporting on Active Switching Matrix (ASM) engage duration (for voice interactions only). Applies only to deployments with Outbound Contact in a VoIP environment where campaigns are running in an ASM dialing mode.

- **false**—Genesys Info Mart does not differentiate between ASM engage duration and regular talk time.

- **true**—Genesys Info Mart populates the ASM_COUNT and ASM_ENGAGE_DURATION columns in the INTERACTION_RESOURCE_FACT (IRF) table as follows:

  - ASM_COUNT is set to 1 to indicate that an attempt was made to engage an agent. If no attempt was made, then the column shows a value of 0.

  - ASM_ENGAGE_DURATION indicates the amount of time that a successfully-engaged agent waited to be connected to a customer. This duration will be excluded from regular talk time.

populate-media-neutral-sm-facts

Default Value: false

Valid Values: true, false
Changes Take Effect: On the next ETL cycle
Dependencies: None
Introduced: 8.5.002

Enables or disables population of the summarized resource states across all media (including voice, e-mail, chat, and 3rd Party Media media types) in the SM_MEDIA_NEUTRAL_STATE_FACT table. The value of the applicable **populate-sm-*-resource-activity** option must be set to true for agent states on those media to be reflected in the media-neutral agent states. The media-neutral state is the highest-priority state in effect for any of the agent's media, where the priority is determined by the **sm-resource-state-priority** option.

populate-mm-ixnqueue-facts

Default Value: false
Valid Values: true, false
Changes Take Effect: On the next ETL cycle
Dependencies: None

Enables or disables the population of eServices/Multimedia Interaction Queue activity to the MSF table. Genesys Info Mart uses this value for all configured eServices/Multimedia Interaction Queues, unless you configure an option that has the same name on an individual Script object for a specific Interaction Queue in the interface you use for configuration.

- **false**—The placement of an interaction in an Interaction Queue will not be represented in the MSF table.

- **true**—The placement of an interaction in an Interaction Queue will be represented in the MSF table.

Starting with release 8.5.003, Genesys Info Mart always creates an MSF record for the first Interaction Queue that an inbound interaction enters, regardless of the setting of this option.

Strategy time is not included in the mediation duration, with the following exceptions:

- Scenarios in which an interaction is bounced between an Interaction Queue and a strategy while an available agent is being identified

- Between releases 8.5.002 and 8.5.006, when the **expand-mediation-time-for-gapless** option is `true`

- Starting with release 8.5.007, when the **show-non-queue-mediation-mm** option is set to `true`

For more information, see Populating Mediation Segments in the *Genesys Info Mart User's Guide*.

> ### Tip
> Ideally, **populate-mm-ixnqueue-facts** should be set to `true` only for Interaction Queues for which mediation reporting is desired, so that the MSF table does not become cluttered with unnecessary information. Some Interaction Queues are not useful for reporting, such as the Interaction Queues that are associated with

Interaction Workbin objects only through a configured Interaction Queue View.

populate-mm-workbin-facts

Default Value: true
Valid Values: true, false
Changes Take Effect: On the next ETL cycle
Dependencies: populate-workbin-as-hold

Enables or disables the population of eServices/Multimedia Interaction Workbin activity to the MSF table. For workbins that are associated with handling resources of type Agent or Place, this option comes into effect only if Genesys Info Mart has not been configured to consider workbin time as hold (populate-workbin-as-hold=false). For the circumstances under which Genesys Info Mart considers workbin time as hold, see the description of the populate-workbin-as-hold option.

Genesys Info Mart uses this value for all configured eServices/Multimedia Interaction Workbins, unless you configure an option that has the same name on an individual Script object for a specific Interaction Workbin in the interface you use for configuration.

- `false`—The placement of an interaction in an Interaction Workbin will not be represented in the MSF table.

- `true`—Provided that Genesys Info Mart does not consider the workbin time as hold, the placement of an interaction in an Interaction Workbin will be represented in the MSF table. The MSF record references a WORKBIN dimension that identifies the type of resource that is associated with the workbin and the specific resource that is associated with the mediation. Strategy time is not included in the mediation duration, except for scenarios in which an interaction is bounced between an Interaction Workbin and a strategy, as the strategy retries busy agents repeatedly. For more information, see the description of the MEDIATION_DURATION field in the section about populating MSFs in the *Genesys Info Mart User's Guide*.

populate-sip-im-facts

Default Value: false
Valid Values: true, false
Changes Take Effect: On the next ETL cycle
Introduced: 8.5.003.16

Enables or disables the population of both interaction data and agent activity data from SIP Instant Messaging (IM) interactions. The default value of this option is `false`, indicating that Genesys Info Mart will not transform SIP IM data.

populate-sm-busy-from-mm-ixns

Default Value: false
Valid Values: true, false
Changes Take Effect: On the next ETL cycle

Dependencies: None
Introduced: 8.5.016.04

Enables you to align reporting of agent states more correctly in multimedia deployments with Interaction Server Cluster where:

- The Interaction Servers are configured as an Interaction Server Cluster, but each ICON connects directly to an individual Interaction Server in the cluster, instead of to the proxy.

- Only one ICON (or ICON HA pair) is configured to write agent state information (the ICON **role** option value includes gls).

Genesys Info Mart introduced support for this scenario in release 8.5.016.04, to resolve reporting inconsistencies that would otherwise occur if more than one ICON (or ICON HA pair) recorded agent state information. However, when only one ICON (or ICON HA pair) is configured with the gls role, only that ICON (or ICON HA pair) will report a BUSY state for an agent when the agent is engaged in an interaction handled by the Interaction Server to which that ICON (or HA pair) is connected. Other ICONs in this configuration will write interaction data when the respective Interaction Server to which the ICON is connected handles an interaction, but these other ICONs will not provide BUSY states.

To enable Genesys Info Mart to support the scenario without losing any BUSY state information, set **populate-sm-busy-from-mm-ixns** to true. With this setting, Genesys Info Mart generates BUSY states for agents based on multimedia interaction data, resulting in more accurate agent state reporting.

However, be aware that state durations might still be affected if time is not synchronized among the Interaction Servers in the cluster—for example, a state's duration might be reported as one second less than actual.

populate-sm-chat-resource-activity

Default Value: true
Valid Values: true, false
Changes Take Effect: On the next ETL cycle
Dependencies: None for release 8.1.1 and later; in earlier releases, has an effect only on the tables that are populated when the values of the following options are set to true: populate-sm-resource-session-facts, populate-sm-resource-state-facts, and populate-sm-resource-state-reason-facts.

Enables or disables the population of Genesys eServices/Multimedia chat resource activity in the SM_RES_SESSION_FACT, SM_RES_STATE_FACT, and SM_RES_STATE_REASON_FACT tables.

populate-sm-email-resource-activity

Default Value: true
Valid Values: true, false
Changes Take Effect: On the next ETL cycle
Dependencies: None for release 8.1.1 and later; in earlier releases, has an effect only on the tables that are populated when the values of the following options are set to true: populate-sm-resource-session-facts, populate-sm-resource-state-facts, and populate-sm-resource-state-reason-facts.

Enables or disables the population of Genesys eServices/Multimedia e-mail resource activity in the

SM_RES_SESSION_FACT, SM_RES_STATE_FACT, and SM_RES_STATE_REASON_FACT tables.

populate-sm-resource-session-facts

Default Value: true
Valid Values: true, false
Changes Take Effect: On the next ETL cycle
Dependencies: None
Discontinued: 8.1.1

In release 8.1.0 and earlier, enables or disables population of the SM_RES_SESSION_FACT table for voice, e-mail, or chat media types if the value of the applicable populate-sm-*-resource-activity option is set to true.

Starting with release 8.1.1, this option has been discontinued, and Genesys Info Mart always populates the SM_RES_SESSION_FACT table. Because of performance improvements in SM_* transformation, Genesys Info Mart no longer needs to control population of this table.

populate-sm-resource-state-facts

Default Value: true
Valid Values: true, false
Changes Take Effect: On the next ETL cycle
Dependencies: populate-sm-resource-session-facts
Discontinued: 8.1.1

In release 8.1.0 and earlier, enables or disables population of the SM_RES_STATE_FACT table for voice, e-mail, or chat media types if the value of the applicable populate-sm-*-resource-activity option is set to true. Because state facts cannot be populated without corresponding session facts, you must set the value of the populate-sm-resource-session-facts option to true in order for the populate-sm-resource-state-facts option to have an effect.

Starting with release 8.1.1, this option has been discontinued, and Genesys Info Mart always populates the SM_RES_STATE_FACT table. Because of performance improvements in SM_* transformation, Genesys Info Mart no longer needs to control population of this table.

populate-sm-resource-state-reason-facts

Default Value: true
Valid Values: true, false
Changes Take Effect: On the next ETL cycle
Discontinued: 8.1.1

In release 8.1.0 and earlier, enables or disables population of the SM_RES_STATE_REASON_FACT table for voice, e-mail, or chat media types if the value of the applicable populate-sm-*-resource-activity option is set to true. Because reason facts cannot be populated without corresponding session and state facts, you must set the values of the populate-sm-resource-session-facts and populate-sm-resource-state-facts options to true in order for the populate-sm-resource-state-reason-facts option to have an effect.

Starting with release 8.1.1, this option has been discontinued, and Genesys Info Mart always populates the SM_RES_STATE_REASON_FACT table. Because of performance improvements in SM_* transformation, Genesys Info Mart no longer needs to control population of this table.

populate-sm-voice-resource-activity

Default Value: true
Valid Values: true, false
Changes Take Effect: On the next ETL cycle

Enables or disables the population of voice resource activity in the SM_RES_SESSION_FACT, SM_RES_STATE_FACT, and SM_RES_STATE_REASON_FACT tables. The option also affects population of INTERACTION_RESOURCE_FACT columns related to agent states—namely, AFTER_CALL_WORK_* and *_ACW_* metrics and RES_PREVIOUS_SM_STATE_* and RES_PREV_SM_STATE_* references.

populate-thread-facts

Default Value: false
Valid Values: true, false
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None
Introduced: 8.5.001

Enables or disables the population of thread-related metrics in the ANCHOR_FLAGS dimension in multimedia deployments.

When populate-thread-facts = false (the default value):

- The following fields in the ANCHOR_FLAG dimension will be ignored for the purposes of populating the IRF.ANCHOR_FLAGS_KEY metric:

  - FIRST_ENGAGE_FOR_AGENT_THRD

  - FIRST_REPLY_FOR_AGENT_THRD

  - FIRST_ENGAGE_THRD

- In deployments that use Reporting and Analytics Aggregates (RAA) or Genesys CX Insights (GCXI), the following agent thread metrics are also not populated:

  - AG2_ID_*.ACCEPTED_THREAD

  - AG2_AGENT_*.ACCEPTED_THREAD

  - AG2_AGENT_GRP_*.ACCEPTED_THREAD

Before release 8.5.001, the behavior was not configurable, and Genesys Info Mart populated thread-related metrics by default. However, interaction thread processing can consume significant resources. Therefore, to improve performance in multimedia deployments that do not need to track interaction threads, Genesys Info Mart no longer populates thread-related metrics by default. The option was introduced to enable you to explicitly activate this functionality, if desired.

This option does not affect population of the chat thread tables and metrics introduced in release

8.5.014.09 for Genesys Engage cloud deployments with Advanced Chat.

populate-workbin-as-hold

Default Value: false
Valid Values: true, false
Changes Take Effect: On the next ETL cycle
Dependencies: None

Specifies whether the time that an interaction is in an Interaction Workbin is considered to be hold time or mediation.

- **false**—Workbin time is considered to be mediation. Whether or not the workbin activity will be represented by an MSF record depends on the value of the populate-mm-workbin-facts option.

- **true**—Workbin time is considered to be hold if the handling resource places the interaction into its own personal workbin. The hold ends when any resource takes the interaction out of the workbin. Various hold metrics are reported in the IRF record that is associated with that handling resource.

For time in an Interaction Workbin to be considered hold time, the Agent or Place resource must place the interaction into its own personal workbin. For example, if Agent1 places an e-mail interaction into its personal workbin, the workbin time is considered to be hold; however, if another resource, such as another agent or a strategy, places the interaction into the workbin for Agent1 to handle, the workbin time is not considered to be hold.

If a handling resource places an interaction into more than one of its own personal workbins, all of the personal workbin time is considered to be hold, and there is no distinction between the hold time in the various workbins. For example, if Agent1 places an e-mail interaction into a Drafts workbin, then pulls the interaction to continue handling it and subsequently places it into a FollowUp workbin, the hold duration that is reported in the IRF record for Agent1's association with the interaction combines the time that is spent in both workbins.

The following table summarizes the results of different permutations of the **populate-workbin-as-hold** and **populate-mm-workbin-facts** options, in combination with the specifics of the interaction flow. If the values of the **populate-workbin-as-hold** and **populate-mm-workbin-facts** options are both set to `false`, the workbin activity is not represented in the dimensional model.

Workbin Reporting Results Matrix

Scenario	Reporting Result for Workbin Activity		
populate-workbin-as-hold=false **populate-mm-workbin-facts=true**	**populate-workbin-as-hold=true** **populate-mm-workbin-facts=true**	**populate-workbin-as-hold=true** **populate-mm-workbin-facts=false**	
An Agent/Place resource places the interaction into its personal workbin.	MSF record.* This is the default behavior.	Hold metrics in the IRF record for the association of that Agent/Place with the interaction.	
Another resource places the interaction into the personal workbin of an Agent/Place.		MSF record.*	Not represented in the dimensional model.

Scenario	Reporting Result for Workbin Activity		
Any resource places the interaction into an AgentGroup or PlaceGroup workbin.			
*Starting with release 8.1.1, the MSF record includes a WORKBIN_KEY value that identifies the workbin instance that is associated with the mediation.			

gim-export Section

- chunk-size-seconds
- days-to-keep-output-files
- max-retries
- output-directory

- output-files-encoding
- retry-delay-seconds
- start-date
- thread-pool-size

- use-export-views
- encrypt-certificate

Use this configuration section to control the behavior of the job that exports data from the Info Mart database (Job_ExportGIM). This Data Export functionality, which was originally introduced for Genesys Engage cloud deployments, is supported for on-premises deployments starting with release 8.5.011.22. For more information about the Data Export feature (also known as BI Data Feed), see "About Data Export" in the *Reporting in the cloud* document or the *Physical Data Model* for your RDBMS.

chunk-size-seconds

Default Value: 86400
Valid Values: Any positive integer
Changes Take Effect: On the next run of Job_ExportGIM
Dependencies: None
Introduced: 8.5.005

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

days-to-keep-output-files

Default Value: 14
Valid Values: Any positive integer
Changes Take Effect: On the next run of Job_ExportGIM
Dependencies: None
Introduced: 8.5.005

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

max-retries

Default Value: 3
Valid Values: Any non-negative integer

Changes Take Effect: On the next run of Job_ExportGIM
Dependencies: None
Introduced: 8.5.005

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

output-directory

Default Value: output
Valid Values: Valid directory path (might not exist)
Changes Take Effect: On the next run of Job_ExportGIM
Dependencies: None
Introduced: 8.5.005

Specifies the directory where exported files are stored.

output-files-encoding

Default Value: utf8
Valid Values: Character encoding supported by Java
Changes Take Effect: On the next run of Job_ExportGIM
Dependencies: None
Introduced: 8.5.005

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

retry-delay-seconds

Default Value: 30
Valid Values: Any non-negative integer
Changes Take Effect: On the next run of Job_ExportGIM
Dependencies: None
Introduced: 8.5.005

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

start-date

Default Value: No default value
Valid Values: Any date after 1970 in the format yyyy-mm-dd hh:mm:ss
Changes Take Effect: On the next run of Job_ExportGIM
Dependencies: None
Introduced: 8.5.005

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.

Important: This option applies only when the directory in which the exported files will be stored — in other words, the directory specified by the **output-directory** option — is empty.

thread-pool-size

Default Value: 10
Valid Values: Any positive integer
Changes Take Effect: On the next run of Job_ExportGIM
Dependencies: None
Introduced: 8.5.005

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

use-export-views

Default Value: false
Valid Values: true, false
Changes Take Effect: On the next run of Job_ExportGIM
Dependencies: None
Introduced: 8.5.009.20

Controls whether Genesys Info Mart will export 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 will export 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 will export data using export views. The export will include the same tables and columns as the previous export(s), regardless of any schema changes resulting from migration(s) that may have occurred in the interim.

encrypt-certificate

Default Value: No default value
Valid Values: Valid path and file name.
Changes Take Effect:
Dependencies: None
Introduced: 8.5.116.20

Specifies the full path and file name of the certificate Genesys Info Mart uses to encrypt exported ZIP files.

gim-transformation Section

- adjust-vq-time-by-strategy-time
- canceled-queues
- cb-virtual-queue-pattern
- chunk-size
- completed-queues
- default-ivr-to-self-service
- expand-mediation-time-for-gapless
- fix-missing-party-links

- ignored-reason-codes
- introduced-transfer-threshold
- irf-io-parallelism
- ixn-data-limit
- kafka-idle-timeout
- link-vrp-vq-msf-to-irf
- msf-target-route-thru-queue
- ocs-caf-aggregates-calls
- ocs-chain-history-limit

- pipeline-timeout-in-hours
- routing-target-regular-dn-fold-external
- show-non-queue-mediation-mm
- stop-ixn-queues
- ud-io-parallelism
- ocs-allowed-lateness

Use this configuration section to specify options that are related to transformation.

adjust-vq-time-by-strategy-time

Default Value: false
Valid Values: true, false
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None

Starting with release 8.1.1, specifies whether Genesys Info Mart will adjust the mediation duration in MSFs for virtual queues to include time that was spent in strategies but not in associated virtual queues—for example, if an interaction spends 3 minutes in a strategy and is in a virtual queue for 2 of those minutes, whether the MSF for the virtual queue will report the duration as 3 minutes or 2 minutes.

- **false**—Genesys Info Mart never includes strategy time that is outside the virtual queue in the mediation duration. This setting means that there might be gaps between the end time of the MSF for a virtual queue that is used by a strategy and the IRF that follows the strategy's routing, or between the start time of the MSF for a virtual queue and the end time of a previous MSF (for example, the MSF for an Interaction Workbin).

- **true**—Genesys Info Mart includes strategy time that is outside the virtual queue in the mediation duration in MSFs for virtual queues.

In multimedia scenarios in which an interaction is bounced between a mediation resource (for example, an Interaction Queue or a Workbin) and a strategy, as the strategy retries busy agents

repeatedly, this option comes into effect only for the virtual queue that is associated with the last strategy party, before the interaction is routed successfully or else terminated. Genesys Info Mart does not report on virtual-queue activity that overlaps the repeated interim mediations.

canceled-queues

Default Value: iWD_Canceled
Valid Values: A comma-separated list of queue names
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None
Introduced: 8.1.3

In multimedia deployments that use archive queues in their business processes, specifies the Interaction Queues that are used as archives for canceled interactions. When an interaction is placed into one of these queues, Genesys Info Mart considers the interaction to be terminated. The transformation job assigns the technical result/reason combination of COMPLETED/CANCELED in the IRF of the handling resource that placed the interaction in the queue, and Genesys Info Mart excludes the interaction from further processing.

The option was introduced in release 8.1.3. In earlier 8.x releases, Genesys Info Mart handled these situations as transfers to a queue.

The option parallels the **completed-queues** configuration option available in Interaction Server. The default value of the Genesys Info Mart option matches the archive queue for canceled interactions in the default business process for the Genesys intelligent Workload Distribution (iWD) solution.

cb-virtual-queue-pattern

Default Value: .*
Valid Values: Any Java regular expression
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None
Introduced: 8.5.015.19

Specifies a pattern for the names of virtual queues used for callbacks. The option enables you to fine-tune Genesys Info Mart behavior with respect to excluding callback virtual queues from mediation reporting. Use any Java regular expression to specify the pattern.

Examples

For this result in accepted callback scenarios...	...use this option value
Exclude all callback-related virtual queue activity that ended after termination of the original call. (This is the default value, which is consistent with legacy behavior.)	.*
Exclude callback-related activity in the virtual queue named *CallbackQueue* if it ended after termination of the original call.	CallbackQueue
Exclude callback-related activity in virtual queues that start with *CallbackQueue* (for example, CallbackQueue1, CallbackQueue2, and so on), if the virtual queue activity ended after termination of the original call.	CallbackQueue.*

For this result in accepted callback scenarios...	...use this option value
Do not exclude any virtual queues.	x^ or any pattern that won't match any virtual queue

chunk-size

Default Value: 7200
Valid Values: Any nonnegative integer
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None
Introduced: 8.5.013.06
Modified: .5.015.19 (behavior changed with respect to Elasticsearch data)

Specifies the maximum size of the time interval, in seconds, for which transformed data is committed in one transaction.

- For data that comes through ICON and IDB, the option applies to transformation of Multimedia and Outbound Contact details.

    - A value of 0 means that the transformation job takes chunks from GIDB based on extraction high-water marks (HWMs) and audit keys generated by the extraction job. In this case, the transformation chunk size does not exceed the chunk size for extraction, as specified by the **extract-data-chunk-size** option.

    - A value greater than 0 means that the transformation job takes chunks from GIDB based on the value of **chunk-size** and the previous transformation HWM, without regard to audit keys set by the extraction job.

- For data that comes through Elasticsearch:

    - In releases earlier than 8.5.015.19, the transformation job always uses the value of **extract-data-chunk-size** to set the chunk size for transformation.

    - Starting with release 8.5.015.19, the transformation job uses the smaller of **extract-data-chunk-size** and **chunk-size** to determine the chunk size to use for Elasticsearch data transformation.

- For data that comes through Kafka, starting with release 8.5.015.19 the chunk size logic is similar to the logic applied to Elasticsearch data.

For data that comes from IDB, Genesys expects that the value of **chunk-size** (default value 2 hours) will usually be greater than the value of **extract-data-chunk-size** (default value 15 minutes). When Genesys Info Mart extraction and transformation are running normally with frequent ETL cycles, the amount of data available for transformation is smaller than the **chunk-size** option, and all available data is transformed. In cases where there is a transformation backlog, a **chunk-size** value greater than **extract-data-chunk-size** will enable transformation to catch up with extraction quickly.

However, attempting to transform a large amount of data in one chunk can lead to performance issues or OutOfMemory errors. In cases where there is a very large backlog or where a very large quantity of data has been extracted (for example, because **extract-data-chunk-size** has been set to a large value and/or "runaway strategy" scenarios have occurred), temporarily reducing the value of the **chunk-size** option, even to a value smaller than **extract-data-chunk-size**, enables Genesys Info Mart to process abnormally large amounts of data more efficiently, eventually catching up with extraction.

completed-queues

Default Value: iWD_Completed
Valid Values: A comma-separated list of queue names
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None
Introduced: 8.1.3

In multimedia deployments that use archive queues in their business processes, specifies the Interaction Queues that are used as archives for completed interactions. When an interaction is placed into one of these queues, Genesys Info Mart considers the interaction to be terminated. The transformation job assigns the technical result/reason combination of COMPLETED/ARCHIVED in the IRF of the handling resource that placed the interaction in the queue, and Genesys Info Mart excludes the interaction from further processing.

The option was introduced in release 8.1.3. In earlier 8.x releases, Genesys Info Mart handled these situations as transfers to a queue.

The option parallels the **completed-queues** configuration option available in Interaction Server. The default value of the Genesys Info Mart option matches the archive queue for completed interactions in the default business process for the Genesys iWD solution.

default-ivr-to-self-service

Default Value: false
Valid Values: true, false
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None

Specifies how Genesys Info Mart will treat IVRs when the IPurpose attached data KVP is not defined, or when it has an incorrect value.

* **false**—The IVR is treated as a nonself-service IVR (in other words, as a mediation device).

* **true**—The IVR is treated as a self-service IVR (in other words, as a handling resource).

expand-mediation-time-for-gapless

Default Value: true
Valid Values: true, false
Changes Take Effect: On the next ETL cycle
Dependencies: None
Introduced: 8.5.002
Modified: 8.5.003 (default value changed from *false* to *true*)
Discontinued: 8.5.007 (replaced by show-non-queue-mediation-mm)

In eServices deployments in which routing activities are performed without the use of virtual queues, enables or disables reporting of the time that a multimedia

interaction spends in a routing strategy as mediation time. To prevent a reporting gap between the end of mediation at the Interaction Queue and the start of handling by the agent, set this option to true. If routing involves more than one Interaction Queue, this configuration also removes a reporting gap between the end of the MSF for one Interaction Queue and the start of the MSF for another Interaction Queue.

- **false**—The time that an interaction spends in a routing strategy will not be reflected in the MSF record for the Interaction Queue.

- **true**—The time that an interaction spends in a routing strategy will be included in the MSF record for the Interaction Queue. If virtual queues are configured to route certain interactions, an additional, separate MSF record represents the interaction's placement in a virtual queue, and the duration of the virtual-queue MSF might also be adjusted.

In release 8.5.003, the default value of the option was changed to true, to ensure there is no gap during user data collection for mediations of active multimedia interactions that have not yet been handled. With the new default behavior, in releases 8.5.003 through 8.5.006 you must overtly disable **expand-mediation-time-for-gapless** if your multimedia deployment uses virtual queues and you do not want the durations of MSFs for virtual queues to be adjusted to eliminate gaps.

In release 8.5.007, when the Genesys Info Mart implementation of gapless mediation reporting changed, **expand-mediation-time-for-gapless** was replaced by **show-non-queue-mediation-mm**.

fix-missing-party-links

Default Value: false
Valid Values: true, false
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None
Introduced: 8.5.015.23

In multimedia deployments, specifies whether Genesys Info Mart will fill data gaps caused by missing party information in ICON data, to reduce the number of redundant records in the MSF table.

- **true**—Genesys Info Mart fills in missing party information and consolidates mediation activity into a reduced number of MSF records, which more meaningfully reflect mediation intent.

- **false**—Genesys Info Mart does not apply the logic to infer missing party information. As a result, Genesys Info Mart might create unnecessary MSF records and might report misleading mediation durations. (This is the default value, which preserves previous behavior.)

The option affects Genesys Info Mart reporting in scenarios in which parent party or other party information is missing when an interaction repeatedly enters the same interaction queue or virtual queue—for example, because of strategy behavior following routing timeouts, or because ICON or Universal Routing Server (URS) was restarted while the interaction was in the queue.

ignored-reason-codes

Default Value: INTERACTION_WORKSPACE
Valid Values: A comma-separated list of reason codes
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None
Introduced: 8.1.4

Specifies agent-state reason codes that will be ignored by reporting. Each reason code that you specify must exactly match the key name for that reason code. Any hardware or software reason code keys specified by this option will not appear in the RESOURCE_STATE_REASON and SM_RES_STATE_REASON_FACT tables.

The option was introduced in release 8.1.4.

The default value means that Genesys Info Mart will ignore reason codes with a key name of INTERACTION_WORKSPACE, which Genesys License Reporting Manager (LRM) attaches to indicate that Genesys Workspace Desktop Edition—formerly known as Interaction Workspace (IWS)—is being used. This reason code is seldom useful for business reporting.

The option value applies only to agent-state reason codes that have not already begun transformation. Changing the option value does not affect agent-state reason data that has already been transformed, even if the agent-state reason was not yet finished when it was transformed.

introduced-transfer-threshold

Default Value: 0
Valid Values: Any nonnegative integer
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None
Introduced: 8.5.001

Specifies a time threshold, in seconds, for Genesys Info Mart to treat a conference as an introduced transfer. If the conference initiator's participation is less than the threshold, while the receiving agent continues on the call, Genesys Info Mart treats this call flow as a special case of transfer. The option applies only to voice calls.

In deployments with business processes that require a transferring agent to introduce the customer to another agent before transferring the call, the option enables you to identify this call flow as an introduced transfer. The IRFs for the receiving and introducing agent in an introduced transfer have technical descriptor combinations for transfers, but with a role reason or technical result reason of `IntroducedTransfer`. Similarly, IRF metrics for the agents accrue as they do for transfers.

The default value means that Genesys Info Mart will treat this call flow as a short conference, which the initiating agent happens to leave first. In this case, the IRFs for the conference part of the receiving and introducing agents' activity will have the usual technical descriptor combinations and metrics for conferences.

Genesys Info Mart supports both single-step and two-step introduced transfers, but support for single-step introduced transfers is limited to deployments in which ICON 8.1.500.04 or higher supports single-step conference (see the *Interaction Concentrator 8.1.x Release Note*).

irf-io-parallelism

Default Value: 4
Valid Values: Any positive integer
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None

Specifies the number of parallel reading processes for IRF transformation. The optimal value of this option depends on DBMS tuning and available resources.

ixn-data-limit

Default Value: 10000
Valid Values: 0 or any positive integer, where 0 means there is no limit imposed
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None
Introduced: 8.5.014.19

If positive, the value imposes a limit on the number of input records—such as calls, links, virtual queues, or party histories—that can be associated with a particular interaction. If the actual number of records exceeds the configured limit, error message 55-20106 is logged, and the interaction is discarded. The option currently applies only to voice interactions.

> ### Important
>
> Do not change this value unless directed to do so by Genesys Customer Care.

Consider setting an alarm on log message 55-20106, so that you can correct problematic scenarios that result in excessive numbers of records.

kafka-idle-timeout

Default Value: 10 (seconds)
Valid Values: A number (of seconds) or duration in ISO 8601 format
Changes Take Effect: On the next ETL cycle
Dependencies: None
Introduced: 8.5.011.23

Specifies the idle timeout for polling Kafka records. If polling does not return any records within the timeout, Genesys Info Mart stops polling Kafka until the next ETL cycle.

In releases earlier than 8.5.011.23, the timeout was hard-coded, with a value of 2 seconds. In high-latency networks, the hard-coded value sometimes caused Genesys Info Mart to skip transformation of Kafka records, even if data was available.

link-vrp-vq-msf-to-irf

Default Value: false
Valid Values: true, false
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None
Introduced: 8.5.004.09

Enables linking of the MSF record for a virtual queue to the IRF record for the target agent, with a technical result of `Diverted/AnsweredByAgent`, in scenarios where a nonself-service IVR port uses a virtual routing point for routing operations, and the strategy includes a virtual queue.

The default value of `false` preserves the existing behavior of not linking the MSF record to the agent's IRF record, and assigning a technical result of `Diverted/Unspecified`.

> ### Important
> Do not change this value unless directed to do so by Genesys Customer Care.

msf-target-route-thru-queue

Default Value: false
Valid Values: true, false
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None

Specifies which party is recorded as the target of mediation in the MSF table in "route-thru-queue" scenarios—scenarios in which a call is routed from a Routing Point through an ACD queue to an agent, using Direct Agent Call functionality (such as Avaya's).

- **true**—Genesys Info Mart considers the next handling resource to be the target. In other words, Genesys Info Mart considers the party to which the ACD queue eventually diverts the call to be the target of the virtual-queue distribution as well. In the previously described scenario, the target would be the agent to whom the ACD queue diverts the call.

- **false**—Genesys Info Mart uses the party immediately following the Routing Point as the target. In the previously described scenario, the target would be the ACD queue.

This option affects the technical results and targets that are reported in the MSF records for virtual queues and ACD queues, as well as the mediation segments and resources that are referenced in the IRF record. For more information, see the section about populating MSFs in the *Genesys Info Mart User's Guide*.

ocs-caf-aggregates-calls

Default Value: true
Valid Values: true, false
Changes Take Effect: At the next run of Job_TransformGIM

Dependencies: None
Introduced: 8.5.015.07

Specifies whether Genesys Info Mart will create separate CONTACT_ATTEMPT_FACT (CAF) records or a single, aggregated CAF record for calls dialed in the context of the same CALL_ATTEMPT_GUID.

- **true**—Genesys Info Mart will create a single CAF record for all outbound calls dialed as part of the same attempt to reach a customer. If there are multiple calls, CAF.CALLID refers to the last dialed call.

- **false**—Genesys Info Mart will create a separate CAF record for each outbound call dialed as part of the same attempt to reach a customer. This nondefault value reflects Genesys Info Mart behavior prior to release 8.5.015.07.

ocs-chain-history-limit

Default Value: 5000
Valid Values: 0 or any positive integer, where 0 means there is no limit imposed
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None
Introduced: 8.5.014.14

If positive, the value imposes a limit on the number of GIDB_GO_FIELDHIST or GIDB_GO_CHAINREC_HIST records that can be associated with a particular CHAINGUID. If the actual number of records exceeds the configured limit, error message 55-20176 is logged, and the chain is ignored.

The option was introduced to prevent OutOfMemory errors during transformation in scenarios where suboptimal SCXML logic results in an excessive number of redial attempts. For example, when an internal case gets reopened from the closed state and the result of the dialed call to the customer is NO_ANSWER, if SCXML keeps adding a record to the calling list and Outbound Contact Server (OCS) keeps dialing the number until the customer answers the call, there can be a very large number of GIDB_GO_FIELDHIST and GIDB_GO_CHAINREC_HIST records for this one CHAINGUID.

Consider setting an alarm on log message 55-20176, so that you can correct problematic scenario logic if necessary.

pipeline-timeout-in-hours

Default Value: 1
Valid Values: Any positive integer
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None
Introduced: 8.1.3

Specifies the maximum expected duration, in hours of single execution transformation pipeline. If timeout is exceeded, GIM will try to abort pipeline and fail transformation job.

Single transformation pipeline refers to the separate pipelines for in-memory transformation of separate data streams, such as for voice, multimedia, Outbound Contact, and agent data.

The option was introduced in release 8.1.3, with the default value set to preserve legacy behavior.

For performance reasons, Genesys recommends that you retain the default value unless Genesys Customer Care advises you to change it.

routing-target-regular-dn-fold-external

Default Value: true
Valid Values: true, false
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None
Introduced: 8.5.008.25

Controls whether Genesys Info Mart populates ROUTING_TARGET records for each distinct regular external DN, or folds them into a single record.

- true — Genesys Info Mart creates a single ROUTING_TARGET record with TARGET_OBJECT_SELECTED = EXTERNAL, to represent the routing target dimension for all regular external DNs used as routing targets.

- false — Genesys Info Mart creates separate ROUTING_TARGET records for each distinct regular external DN used as a routing target. (This is the legacy behavior.)

In deployments that include Callback reporting, Genesys strongly recommends retaining the default value of true. Otherwise, the large number of callbacks to external phone numbers, each of which is identified as a separate target selected for routing, results in an explosion of records in the ROUTING_TARGET dimension that are not significant for reporting.

show-non-queue-mediation-mm

Default Value: false
Valid Values: true, false
Changes Take Effect: On the next ETL cycle
Dependencies: None
Introduced: 8.5.007 (replaces expand-mediation-time-for-gapless)

In eServices deployments, controls whether all mediation time for multimedia interactions, even mediation time that does not occur within a queue, is represented by an MSF.

- **false** — MSFs for multimedia interactions are focused only on the portion of the mediation time that occurs in a queue, whether it is an Interaction Queue or a virtual queue. There may be gaps in time between MSFs, because the multimedia interaction may not be in a queue, or may not be in a queue that is represented in Genesys Info Mart with an MSF.

- **true** — Provided that there is an MSF for the first Interaction Queue in a mediation (which is true for mediation time of unhandled interactions), all mediation time for multimedia interactions that occurs after the first Interaction Queue, even mediation time that does not occur within a queue, is represented by one or more MSFs. Additional, non-queue MSFs are created for multimedia interactions to represent mediation time that occurs outside an Interaction Queue MSF — for example, mediation time that occurs after an MSF for an Interaction Queue, when a routing strategy is attempting to find a routing target without the use of a virtual queue.

For multimedia interactions, mediation that occurs in a virtual queue is always represented in Genesys Info Mart by an MSF. However, depending on the setting for **populate-mm-ixnqueue-**

facts, mediation that occurs in an Interaction Queue might not be represented by an MSF; in fact, often it is not. When **show-non-queue-mediation-mm** is set to `true`, the additional MSFs that are created occur between MSFs for Interaction Queues (in other words, when an interaction moves during mediation from one Interaction Queue that is represented by an MSF to another that is represented by an MSF), or between the MSF for an Interaction Queue and a routing target (agent). The additional, non-queue MSFs may overlap with MSFs for virtual queues, since an interaction may also be in a virtual queue for some (or all) of the mediation time that occurs outside of an Interaction Queue MSF. Furthermore, the additional MSFs may include time that the interaction spent in Interaction Queues that are not represented by an MSF.

The **show-non-queue-mediation-mm** option was introduced in release 8.5.007, when the Genesys Info Mart implementation of gapless mediation reporting changed. In an earlier implementation, Genesys Info Mart used the **expand-mediation-time-for-gapless** configuration option to control whether to expand the durations of queue MSFs to eliminate reporting gaps. Starting with release 8.5.007, Genesys Info Mart no longer adjusts the durations of queue MSFs, and **expand-mediation-time-for-gapless** was discontinued.

Example

Consider the scenario in which a multimedia interaction enters the contact center at time t0 and, after mediation involving various queues and routing strategies, is routed for handling at time t4. Following first handling, there is additional mediation before the interaction is routed for further handling.
t0–t1: InteractionQueue1
t1–t2: Strategy1
t2–t3: InteractionQueue2 (not the same queue as InteractionQueue1)
t3–t4: Strategy2
t4–t5: Agent1
t5–t6: InteractionQueue3
t6–t7: Strategy3
t7–t8: InteractionQueue4
t8–t9: Strategy4
t9–t10: Agent2

The following table summarizes the mediation reporting results, depending on configuration option settings.

populate-mm-ixnqueue-facts=false, show-non-queue-mediation-mm=false	populate-mm-ixnqueue-facts=true, show-non-queue-mediation-mm=false	populate-mm-ixnqueue-facts=false, show-non-queue-mediation-mm=true	populate-mm-ixnqueue-facts=true, show-non-queue-mediation-mm=true
Without virtual queues			
• MSF1 (InteractionQueue1)* • IRF1 (Agent1) • IRF2 (Agent2) *If the second InteractionQueue in the	• MSF1 (InteractionQueue1) • MSF2 (InteractionQueue2) • IRF1 (Agent1) • MSF3	• MSF1 (InteractionQueue1) • MSF2 (Strategy1)** • IRF1 (Agent1) • IRF2 (Agent2)	• MSF1 (InteractionQueue1) • MSF2 (Strategy1) • MSF3 (InteractionQueue2) • MSF4 (Strategy2)

populate-mm-ixnqueue-facts=false, show-non-queue-mediation-mm=false	populate-mm-ixnqueue-facts=true, show-non-queue-mediation-mm=false	populate-mm-ixnqueue-facts=false, show-non-queue-mediation-mm=true	populate-mm-ixnqueue-facts=true, show-non-queue-mediation-mm=true
scenario is actually the same as InteractionQueue1, then MSF1 would cover t0–t3, and there would be a gap of t3–t4.	(InteractionQueue3) • MSF4 (InteractionQueue4) • IRF2 (Agent2)	**Spans t1–t4, including time for InteractionQueue2 and Strategy2.	• IRF1 (Agent1) • MSF5 (InteractionQueue3) • MSF6 (Strategy3) • MSF7 (InteractionQueue4) • MSF8 (Strategy4) • IRF2 (Agent2)
Gaps: t1–t4, t5–t9	**Gaps:** t1–t2, t3–t4, t6–t7, t8–t9	**Gaps:** None up to first handling; t5–t9	**Gaps:** None
If the strategies use virtual queues, there will also be separate MSFs for the virtual queues, which might eliminate gaps, even when **show-non-queue-mediation-mm**=false, and which will overlap with the MSF for the Strategy party when **show-non-queue-mediation-mm**=true.			

> ## Important
>
> As discussed elsewhere (see the description of the IRF.MEDIATION_DURATION column in the *Physical Data Model* for your RDMBS), there are other scenarios, not controlled by this option, where non-queue time might be incorporated into an MSF for a queue, such as when a multimedia interaction bounces back and forth between the same queue and a strategy. In such cases, the repeated bounce-backs are collapsed into a single MSF, which might include non-queue time when the routing strategy was attempting to find a routing target before returning the interaction to the same queue.

stop-ixn-queues

Default Value: No default value
Valid Values: A comma-separated list of queue names
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None
Introduced: 8.1.402

In multimedia deployments that use stop-interaction queues in their business processes, this option specifies the Interaction Queues that are used to handle stopping an interaction (for example, Twitter_StopIxn). When an interaction is placed into one of these queues, Genesys Info Mart considers the interaction to be terminated. The transformation job assigns the technical result/reason combination of COMPLETED/UNSPECIFIED in the IRF of the handling resource that placed the interaction in the queue, and Genesys Info Mart excludes the interaction from further processing. The

agent who placed the interaction in the queue is represented as the party that stopped the interaction, and the strategy that actually stops the interaction and performs any associated post-processing is not represented in Genesys Info Mart reporting.

The option was introduced in release 8.1.402. In earlier 8.x releases, Genesys Info Mart handled these situations as transfers to a queue.

ud-io-parallelism

Default Value: 5
Valid Values: Any positive integer
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None

Starting with release 8.1.1, specifies the number of parallel threads for user-data transformation. The optimal value of this option depends on DBMS tuning and available resources.

ocs-allowed-lateness

Default Value: PT0S
Valid Values: ISO 8601 duration format
Changes Take Effect: On the next ETL cycle
Dependencies: None
Introduced: 8.5.116.20

Specifies how long Genesys Info Mart waits for late-arriving Outbound Contact data. Outbound Contact data (such as GO_METRICS) that arrives late, but within the specified limit, is processed during transformation.

kafka-<cluster-name> Section

- bootstrap.servers
- g:topic:<topic-name>

Kafka cluster–specific configuration sections enable you to specify connection and topic information that enables Genesys Info Mart to consume data from Apache Kafka, in order to make data from producer applications available in the Info Mart database for downstream reporting applications.

If you want Genesys Info Mart to process data that a particular data source published to a Kafka topic, add the applicable configuration section(s) and options on the **Options** tab of the Genesys Info Mart Application object. Name the section **kafka-<cluster-name>**, where <cluster-name> is any alphanumeric string that identifies the Kafka cluster in your deployment, such as kafka-1.

In addition to the **bootstrap.servers** and **g:topic:<topic-name>** options described below and the **kafka-idle-timeout** option in the **[gim-transformation]** section, Genesys Info Mart supports native Kafka client options. Any options in the **kafka-<cluster-name>** section whose name does not start with the g: prefix are treated as Kafka client options. In particular, for a Kafka cluster that uses SASL_SSL authentication, consider configuring the following security options:

- **sasl.mechanism** = SCRAM-SHA-512

- **security.protocol** = SASL_SSL

- **sasl.jaas.config** = org.apache.kafka.common.security.scram.ScramLoginModule required username="<USERNAME>" password="<PASSWORD>";

> ### Important
> To avoid connection failures in deployments where your Kafka cluster uses SASL_SSL authentication and you are using Kafka 2.0.0 or later, you might need to adjust your security setup or else add the **ssl.endpoint.identification.algorithm** option and set the value to an empty string. For more information, see the Known Issue in the *Genesys Info Mart 8.5.x Release Note*.

If you are using SSL connections with a self-signed certificate, consider configuring the following option to specify the path to the trust store file that contains the self-signed certificate and that is located outside of the Genesys Info Mart installation directory:

- **ssl.truststore.location** = <PATH>

For descriptions of native Kafka configuration options, refer to Apache Kafka documentation at https://kafka.apache.org/documentation#consumerapi.

Each cluster-specific configuration section can contain the following options.

bootstrap.servers

Default Value: No default value
Valid Values: Any valid host:port combination that identifies a Kafka server in the cluster
Changes Take Effect: On the next ETL cycle
Dependencies: None
Introduced: 8.5.011.18

Specifies the location of the Kafka broker(s) in the cluster, in the form of host:port. If there are multiple servers, use a comma-separated list.

bootstrap.servers is a standard Kafka consumer option. The option is mandatory for Genesys Info Mart, as a Kafka consumer, to know where to connect for the initial connection to the Kafka cluster. In high availability (HA) deployments, Genesys recommends that you list all the brokers in the cluster.

For more information about the **bootstrap.servers** option, see the Apache Kafka documentation (https://kafka.apache.org/documentation#consumerapi).

g:topic:<topic-name>

Default Value: No default value
Valid Values: Any Genesys Info Mart–defined <mapping-id>, as listed in the option description
Changes Take Effect: On the next ETL cycle
Dependencies: None
Introduced: 8.5.011.18

Specifies a Kafka topic to consume and how messages in this topic will be mapped. The <topic-name> in the option name identifies the topic Genesys Info Mart will look for in Kafka message headers and matches the topic name configured in the producer application. The option value identifies the ID in the CTL_XML_CONFIG table that Genesys Info Mart will use to map that topic into the database schema.

Configure a separate **g:topic:<topic-name>** option for each topic to which the Kafka producers publish reporting data. A particular **kafka-<cluster-name>** section might contain multiple **g:topic:*** options, some of which might have the same <mapping-id> value, depending on the data source.

Genesys Info Mart supports the following <mapping-id> values:

Producer application	<mapping-id>	Supported since Genesys Info Mart release
Bot Gateway Server (BGS)	BGS_K	8.5.011.18
Genesys Co-browse (GCB)	COBROWSE	8.5.011.18

log Section

- standard
- verbose

Use this configuration section to specify the Genesys Info Mart logging options. You can use the following options to enable centralized logging.

> ### Important
> The following two options and their specified valid value(s) are the only common log options that Genesys Info Mart supports.

standard

Default Value: network
Valid Values: network
Changes Take Effect: Immediately
Dependencies: verbose

Specifies the location of the log output.

verbose

Default Value: standard
Valid Values: none, standard, trace
Changes Take Effect: Immediately
Dependencies: standard

Specifies the minimum level of logging.

log4j Section

- console-pattern-layout
- file-pattern-layout
- log-file-name
- log4j.appender.ConsoleLogger.Threshold
- logging-level
- max-backup-index
- max-log-file-size

Use this configuration section to specify the Genesys Info Mart ETL options for logging events to a local file and to STDOUT. These options are separate from, and independent of, the Genesys Central Logger options that you specify in the **[log] section**.

console-pattern-layout

Default Value: %d{ISO8601} %-5p %-12t %m%n
Valid Values: Any valid pattern string
Changes Take Effect: Immediately
Dependencies: None

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

file-pattern-layout

Default Value: %d{ISO8601} %-5p %-12t %m%n
Valid Values: Any valid pattern string
Changes Take Effect: Immediately
Dependencies: None

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

log-file-name

Default Value: gim_etl.log
Valid Values: filespec
Changes Take Effect: Immediately
Dependencies: None

Specifies the path and file name of the log file. If you do not specify a path, the log files will be

created in the installation directory.

log4j.appender.ConsoleLogger.Threshold

Default Value: info
Valid Values: trace, debug, info, warn, error, all, off
Changes Take Effect: Immediately
Dependencies: None

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

logging-level

Default Value: info
Valid Values: debug, info, warn, error, none, all, off
Changes Take Effect: Immediately
Dependencies: None

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.

In releases earlier than 8.1.3, the valid values were DEBUG, INFO, WARN, and NONE.

max-backup-index

Default Value: 10
Valid Values: Any positive integer
Changes Take Effect: Immediately
Dependencies: None

Specifies the maximum number of backup log files that are kept in addition to the active log file.

max-log-file-size

Default Value: 50MB
Valid Values: Any number, followed by a scale (KB for kilobytes, MB for megabytes, or GB for gigabytes)
Changes Take Effect: Immediately
Dependencies: None

Specifies the maximum size of the active log file before it is considered full and renamed as a backup file.

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.

DAP Options

gim-etl Section

- agg-jdbc-url
- cp-reuse-count
- default-schema

- geo-location
- jdbc-host
- jdbc-port

- jdbc-sid
- jdbc-url
- role

Use this configuration section to set Genesys Info Mart–related options on database access points (DAPs) that Genesys Info Mart uses:

- The *extraction DAPs,* which enable access to the databases from which Genesys Info Mart extracts data

- The *Info Mart DAP,* which enables access to the Info Mart database, in which Genesys Info Mart stores reporting data

agg-jdbc-url

Default Value: No default value
Valid Values: Any valid JDBC URL
Changes Take Effect: On restart of the Genesys Info Mart Server.

In an Oracle RAC configuration in which you want aggregation to use a separate node, specifies the connection parameters for the JDBC connection to the Info Mart database for Reporting & Analytics Aggregates (RAA). If this option is defined, the aggregation process uses the connection string specified by **agg-jdbc-url** instead of the **jdbc-url** connection string. For the required syntax for this option, consult the vendor documentation for your JDBC driver.

For performance reasons, Genesys recommends that you divide processing by function, allocating separate functions to specific nodes in the cluster. Use the Oracle RAC Server Control Utility (SRVCTL) to configure named services, including a named service for RAA, and associate those services with particular nodes. For more information, see "Oracle RAC Configuration" and "DAP Objects for Genesys Info Mart" in the Genesys Info Mart 8.x chapter in the *Genesys Hardware Sizing Guide*.

Example for Oracle
Oracle requires the following format to specify connection parameters for the Oracle thin client:

```
jdbc:oracle:thin:@(DESCRIPTION = (LOAD_BALANCE=OFF)(ADDRESS = (PROTOCOL = TCP)(HOST =
<database_host_name>)(PORT = 1521))(CONNECT_DATA = (SERVER =
```

```
DEDICATED)(SERVICE_NAME=<named service>)))
```

To enable SSL or IPv6 over the JDBC connection, you must specify additional parameters, as described in your RDBMS vendor documentation. For an indication of the additional parameters you must provide, extrapolate from the examples shown in the Notes in the **jdbc-url** option description.

cp-reuse-count

Default Value: -1
Valid Values: Any integer
Changes Take Effect: On restart of the Genesys Info Mart Server
Dependencies: None
Introduced: 8.5.015.07

Specifies the maximum number of times a database connection can be reused.

- -1 (default) or any negative integer — The number of times a connection can be reused is unlimited.

- 0 — Connections cannot be reused. Instead of returning to the pool, they will be closed after being used once.

- n>0 — A connection will be returned to the pool up to n times and then closed.

This option is typically used to periodically recycle connections in order to eliminate issues such as excessive memory growth on the DBMS side.

default-schema

Default Value: No default value
Valid Values: Any valid database owner or schema name
Changes Take Effect: For an extraction DAP, at the next run of the extraction job for the particular data domain; for the Info Mart DAP, on restart of the Genesys Info Mart Server.

Specifies the database schema or owner name, if it is different from the database user ID. If you do not specify this option, the database tables and other database objects are assumed to be owned by the user name that you specify on the **DB Info** tab.

Note for PostgreSQL: Genesys recommends that you use lower case for all schema names. If you cannot do so, specify the value in case-sensitive mode by surrounding the value with a set of double quotes.

geo-location

Default Value: ""
Valid Values: Any string
Changes Take Effect: For an extraction DAP, at the next run of the extraction job for the particular data domain; for the Info Mart DAP, on restart of the Genesys Info Mart Server.
Introduced: 8.1.2

On the extraction DAP, specifies the location of the IDB. On the Info Mart DAP, specifies the location of the Info Mart database. Genesys Info Mart compares the string value of the option on the Info Mart DAP against the value of the **geo-location** option in redundant extraction DAPs. If the values are the same, the IDB for which the extraction DAP provides the connection information is considered to be local; if the values are not the same, the IDB is considered to be remote.

In an HA environment, Genesys Info Mart uses geolocation as a tie-breaker to determine the best IDB from which to extract data: If data-quality criteria do not identify one IDB as the best, Genesys Info Mart gives preference to the local IDB.

Example
If the value of the **geo-location** option on the Info Mart DAP is Site 1, the value of the option on the extraction DAP for IDB-1 is Site 1, and the value on the extraction DAP for IDB-2 is Site 2, Genesys Info Mart will consider IDB-1 to be local and IDB-2 to be remote, regardless of where the hosts are physically located.

jdbc-host

Default Value: No default value
Valid Values: Any valid host name
Changes Take Effect: For an extraction DAP, at the next run of the extraction job for the particular data domain; for the Info Mart DAP, on restart of the Genesys Info Mart Server.

Specifies the host on which the DBMS is running. Configure this option if the DBMS server for the IDB is not on the same host as the DB Server.

Note: If the **jdbc-url** option is defined, the host name that is specified in the URL overwrites the value of **jdbc-host.**

jdbc-port

Default Value: 1433 (for Microsoft SQL Server), 1521 (for Oracle), or 5432 (for PostgreSQL)
Valid Values: Any valid port number
Changes Take Effect: For an extraction DAP, at the next run of the extraction job for the particular data domain; for the Info Mart DAP, on restart of the Genesys Info Mart Server.

Specifies the port of the DBMS listener. Configure this option if the DBMS uses a nondefault listener port.

Note: If the **jdbc-url** option is defined, the port name that is specified in the URL overwrites the value of **jdbc-port.**

jdbc-sid

Default Value: No default value

Valid Values: Any valid SID name
Changes Take Effect: For an extraction DAP, at the next run of the extraction job for the particular data domain; for the Info Mart DAP, on restart of the Genesys Info Mart Server.

(For Oracle only) Specifies the SID name, which Genesys Info Mart uses as a connection parameter to connect to Oracle databases. DB Server uses the Transparent Network Substrate (TNS) name, which is specified in the **DBMS Name** field on the **DB Info** tab of the DAP Application object. Configure this option if the SID name is not the same as the TNS name.

Note: If the **jdbc-url** option is defined, the SID name that is specified in the URL overwrites the value of **jdbc-sid.**

jdbc-url

Default Value: No default value
Valid Values: Any valid JDBC URL
Changes Take Effect: For an extraction DAP, at the next run of the extraction job for the particular data domain; for the Info Mart DAP, on restart of the Genesys Info Mart Server.

Specifies all of the connection parameters for JDBC connection to the IDB or the Info Mart database. For the required syntax for this option, consult the vendor documentation for your JDBC driver.

- **Example for Microsoft SQL Server**
  Microsoft SQL Server requires the following format to specify connection parameters for the Microsoft JDBC driver: `jdbc:sqlserver://<dbserver>;DatabaseName=<dbname>;SelectMethod=cursor`

- **Example for Oracle**
  Oracle requires the following format to specify connection parameters for the Oracle thin client:
  `jdbc:oracle:thin:@(DESCRIPTION = (ADDRESS = (PROTOCOL = TCP)(HOST = <database_host_name>)(PORT = 1521))(CONNECT_DATA = (SERVER = DEDICATED)(SID = <Oracle_SID>)))`

- **Example for PostgreSQL**
  PostgreSQL requires the following format to specify connection parameters for the PostgreSQL JDBC driver: `jdbc:postgresql://<database_host_name>:<port>/<dbname>`

Notes:

- If this option is defined, it overwrites the values of any other **jdbc-*** options that might be configured for this DAP.

- To enable SSL over the JDBC connection, you must specify additional parameters, as described in your RDBMS vendor documentation. The following are sample URLs for Microsoft SQL Server, Oracle, and PostgreSQL, respectively:

  - `jdbc:sqlserver://host:1433;databaseName=NAME;integratedSecurity=true;encrypt=true;trustServerCertificate`

  - `jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=tcps)(HOST=servername)(PORT=2484))(CONNECT_DATA=(SERVI`

  - `jdbc:postgresql://<database_host_name>:<port>/<dbname>?ssl=true`

- To enable IPv6 over the JDBC connection, you must specify additional parameters, as described in your RDBMS vendor documentation. The following are sample URLs for Microsoft SQL Server, Oracle, and PostgreSQL, respectively:

  - `jdbc:sqlserver://;serverName=<database_ipv6_address>;DatabaseName=<dbname>;SelectMethod=cursor`

  - `jdbc:oracle:thin:@(DESCRIPTION = (ADDRESS = (PROTOCOL = TCP)(HOST = [<database_ipv6_address>])(PORT = 1521))(CONNECT_DATA = (SERVER = DEDICATED)(SERVICE_NAME = <oracle_SID>)))`

  - `jdbc:postgresql://[<database_ipv6_address>]:<port>/<dbname>`

- Genesys does not recommend enabling load-balancing in Oracle RAC configurations, because continual node resynchronization degrades performance. If your deployment includes RAA, Genesys recommends that you use the **agg-jdbc-url** option to enable aggregation to use a separate node. For more recommendations about Oracle RAC configuration, see "Oracle RAC Configuration" and "DAP Objects for Genesys Info Mart" in the Genesys Info Mart 8.x chapter in the *Genesys Hardware Sizing Guide*.

role

Default Value: No default value
Valid Values: ICON_CFG, ICON_CORE, ICON_MM, ICON_OCS, INFO_MART
Changes Take Effect: For an extraction DAP, the next time that extraction jobs are launched; for the Info Mart DAP, on restart of the Genesys Info Mart Server.

Extraction DAPs

Specifies which data domain Genesys Info Mart will extract through the DAP:

- `ICON_CFG` — The DAP specifies the connection information for an IDB from which Genesys Info Mart will extract ICON Configuration details. This role must be associated with one and only one DAP, or with redundant DAPs that constitute one HA set.

- `ICON_CORE` — The DAP specifies the connection information for an IDB from which Genesys Info Mart will extract ICON Voice details, including agent activity details. This role is optional, but you need at least one DAP with a role of `ICON_CORE` or `ICON_MM`.

- `ICON_MM` — The DAP specifies the connection information for an IDB from which Genesys Info Mart will extract ICON Multimedia details, including agent activity details. This role is optional, but you need at least one DAP with a role of `ICON_CORE` or `ICON_MM`.

> ## Warning
> You cannot reconfigure the DAP role from `ICON_CORE` to `ICON_MM`, or vice versa, once data extraction occurs.

- `ICON_OCS` — The DAP specifies the connection information for an IDB from which Genesys Info Mart will extract ICON Outbound Contact details.

- For the supported topologies that combine data domains (see Interaction Concentrator Topologies in the *Genesys Info Mart Deployment Guide*), specify all the applicable roles, separated by commas — for example: `ICON_CFG,ICON_CORE,ICON_OCS`.

> ### Important
>
> Ensure that the extraction role you specify is consistent with the role that is specified for ICON (in the [callconcentrator] section in both the ICON application and DAP). For more information about the corresponding ICON roles, see ICON role in the *Deployment Guide.*

Info Mart DAP

- INFO_MART — The DAP specifies connection information for the Genesys Info Mart Server to access the Info Mart database to read and write data. This role is required; it can be associated with only one DAP.

Switch Options

Settings on the switch affect both ICON and Genesys Info Mart reporting. Switch options that directly affect Genesys Info Mart functioning are contained in the following configuration sections, configured on the **Annex** of Switch objects:

- gim-etl
- gim-etl-media-chat
- gim-etl-media-email

There are additional ICON-related Switch options that affect reporting. For full information about configuring switches for ICON and Genesys Info Mart, see Configuring Switch Objects.

> ### Important
> The ***-threshold** options described on this page do not control the thresholds used for aggregation. For information about the options that do control aggregation-specific thresholds, see the *Reporting and Analytics Aggregates Deployment Guide.*

gim-etl Section

Use this configuration section to set Genesys Info Mart–related options on the Switch object that handles voice or multimedia interactions.

- factor-dnd-into-sm-resource-states
- network-switch

- q-answer-threshold-voice
- q-short-abandoned-threshold-voice

factor-dnd-into-sm-resource-states

Default Value: false (for voice-handling and SIP switches); true (for multimedia-handling switches)
Valid Values: true, false
Changes Take Effect: On the next ETL cycle

Specifies whether DND (do-not-disturb) status for a resource on a given switch is factored into resource states and reasons.

- **false** (default for voice and SIP switches)—Genesys Info Mart does not factor DND into summarized resource states and reasons.

- **true** (default for eServices/Multimedia switches)—Genesys Info Mart does factor DND into summarized resource states and reasons.

When the value is true, DND will be reported only if it occurs within an agent's session relative to a given media type. DND that is set in a Multimedia login session prior to the addition of any media will not be reported.

If you change the value of this option, the new option value is not applied to previously loaded facts.

network-switch

Default Value: false
Valid Values: true, false
Changes Take Effect: At the next run of Job_ExtractICON
Discontinued: 8.x for Genesys Info Mart purposes (the option has no effect)

Specifies whether the switch is a network switch.

q-answer-threshold-voice

Default Value: 60
Valid Values: 1-10000
Changes Take Effect: On the next ETL cycle
Dependencies: None

Specifies the default duration, in seconds, that is used as a target time to answer voice interactions that were distributed by virtual queues or ACD queues configured on the switch.

This setting overrides the value of the option with the same name configured on the Genesys Info Mart Application object. The switch-level value can, in turn, be overridden by an option with the same name configured on the individual DN queue objects.

If you change the value of this option, the new option value is not applied to previously loaded facts.

For similar options that control equivalent thresholds for multimedia interactions at switch level, see the gim-etl-media-chat Section and the gim-etl-media-email Section.

q-short-abandoned-threshold-voice

Default Value: 10
Valid Values: 1-1000
Changes Take Effect: On the next ETL cycle
Dependencies: None

Introduced: 8.5.003

For all queues that are configured on this switch, specifies the maximum duration of mediation, in seconds, that is used to indicate that an interaction that was abandoned while in a queue should be considered a "short" abandon. Genesys Info Mart uses this value to determine the state of SHORT_ABANDONED_FLAG in the MEDIATION_SEGMENT_FACT (MSF) row for voice interactions that are abandoned in a virtual queue or ACD queue.

This setting overrides a value that is set in the Genesys Info Mart Application object. The switch-level value can, in turn, be overridden by an option with the same name configured on the individual DN queue objects.

If you change the value of this option, the new option value is not applied to previously loaded facts.

In releases earlier than 8.5.003, this value could not be set on an individual Switch object.

For similar options that control equivalent thresholds for multimedia interactions, see the gim-etl-media-chat Section and the gim-etl-media-email Section.

gim-etl-media-chat Section

Use this configuration section to set Genesys Info Mart–related options for Genesys Chat on the Switch object that handles multimedia interactions. Media-specific configuration sections enable you to specify separate thresholds for different types of multimedia interactions.

- q-answer-threshold
- q-short-abandoned-threshold

q-answer-threshold

Default Value: 60
Valid Values: 1-600000
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None

For all Virtual Queue DNs configured on the switch, specifies the default duration, in seconds, that is used as a target time to accept a chat interaction that entered a queue.

This setting overrides a value that is set in the Genesys Info Mart Application object or the tenant-level Media Type Business Attribute. The switch-level value can, in turn, be overridden for individual

virtual queues by an option with the same name configured on Virtual Queue DN objects.

q-short-abandoned-threshold

Default Value: 10
Valid Values: 1-1000
Changes Take Effect: On the next ETL cycle
Introduced: 8.5.003

For all Virtual Queue DNs configured on the switch, specifies the maximum duration of mediation, in seconds, that is used to indicate that an interaction that was abandoned while in a virtual queue should be considered a "short" abandon. Genesys Info Mart uses this value to determine the state of SHORT_ABANDONED_FLAG in the MSF row for multimedia interactions that are abandoned in a virtual queue.

This setting overrides a value that is set in the Genesys Info Mart Application object or the tenant-level Media Type Business Attribute. The switch-level value can, in turn, be overridden for individual virtual queues by an option with the same name configured on Virtual Queue DN objects.

In releases earlier than 8.5.003, this value could not be set on an individual Switch object.

If you change the value of this option, the new option value is not applied to previously loaded facts.

gim-etl-media-email Section

Use this configuration section to set Genesys Info Mart–related options for Genesys E-mail on the Switch object that handles multimedia interactions. Media-specific configuration sections enable you to specify separate thresholds for different types of multimedia interactions.

- q-answer-threshold
- q-short-abandoned-threshold

q-answer-threshold

Default Value: 60
Valid Values: 1-600000
Changes Take Effect: At the next run of Job_TransformGIM

For all Virtual Queue DNs configured on the switch, specifies the default duration, in seconds, that is used as a target time to accept an email interaction that entered a queue.

This setting overrides a value that is set in the Genesys Info Mart Application object or the tenant-level Media Type Business Attribute. The switch-level value can, in turn, be overridden for individual virtual queues by an option with the same name configured on Virtual Queue DN objects.

q-short-abandoned-threshold

Default Value: 10
Valid Values: 1-1000
Changes Take Effect: On the next ETL cycle
Introduced: 8.5.003

For all Virtual Queue DNs configured on the switch, specifies the maximum duration of mediation, in seconds, that is used to indicate that an interaction that was abandoned while in a virtual queue should be considered a "short" abandon. Genesys Info Mart uses this value to determine the state of SHORT_ABANDONED_FLAG in the MSF row for multimedia interactions that are abandoned in a virtual queue.

This setting overrides a value that is set in the Genesys Info Mart Application object or the tenant-level Media Type Business Attribute. The switch-level value can, in turn, be overridden for individual virtual queues by an option with the same name configured on Virtual Queue DN objects.

In releases earlier than 8.5.003, this value could not be set on an individual Switch object.

If you change the value of this option, the new option value is not applied to previously loaded facts.

Media Type Business Attribute Options

A separate Media Type Business Attribute is configured for each media type for each tenant. The Genesys Info Mart–related options described on this page enable you to specify separate thresholds for different types of multimedia interactions (Genesys eServices interactions as well as 3rd Party Media interactions) for different tenants.

For more information about configuring Genesys Info Mart–related options for Media Types, see Configuring Media Type Business Attributes.

> ## Tip
> For Genesys Info Mart–related options that you configure in a Media Type Business Attribute object, changes take effect on the next ETL cycle. The new option value is not applied to previously loaded facts.

gim-etl-media Section

An option configured at the tenant level overrides the setting of that same option at the application level (in other words, on the Application object), implementing media-specific thresholds for multimedia interactions for a specific tenant.

- q-answer-threshold
- q-short-abandoned-threshold
- short-abandoned-threshold

q-answer-threshold

Default Value: 60
Valid Values: 1-600000
Changes Take Effect: At the next run of Job_TransformGIM

For multimedia interactions of the particular media type for the particular tenant, specifies the default duration, in seconds, that is used on all configured queues as a target time to accept a multimedia interaction that entered a queue.

This setting overrides a value that is set in the Genesys Info Mart Application object. The tenant-level value can, in turn, be overridden for virtual queues by an option with the same name configured on the Switch or individual Virtual Queue DN objects or, for interaction queues or workbins, by an option with the same name configured on the individual Script objects.

q-short-abandoned-threshold

Default Value: 10
Valid Values: 1-1000
Changes Take Effect: On the next ETL cycle
Introduced: 8.5.003

For multimedia interactions of the particular media type for the particular tenant, specifies the maximum duration of mediation, in seconds, that is used on all configured queues to indicate that an interaction that was abandoned while in an interaction queue or workbin should be considered a "short" abandon. Genesys Info Mart uses this value to determine the state of SHORT_ABANDONED_FLAG in the MSF row for multimedia interactions that are abandoned in a virtual queue, interaction queue, or workbin.

This setting overrides a value that is set in the Genesys Info Mart Application object. The tenant-level value can, in turn, be overridden for virtual queues by an option with the same name configured on the Switch or individual Virtual Queue DN objects or, for interaction queues or workbins, by an option with the same name configured on the individual Script objects.

In releases earlier than 8.5.003, this value could not be set on a DN or Script object.

short-abandoned-threshold

Default Value: 10
Valid Values: 0-60
Changes Take Effect: On the next ETL cycle

For interactions of the particular media type for the particular tenant, specifies the minimum duration, in seconds, of an abandoned multimedia interaction in order for it to be considered truly abandoned. Genesys Info Mart uses this value to determine the state of SHORT_ABANDONED_FLAG in the IRF row.

This setting overrides a value that is set in the Genesys Info Mart Application object.

DN Options

Settings on DNs affect both ICON and Genesys Info Mart reporting. The Genesys Info Mart–related options described on this page enable you to control reporting behavior at the level of individual virtual or ACD queues. The options are contained in the following configuration sections, configured on the **Annex** of DN objects:

- gim-etl
- gim-etl-media-chat
- gim-etl-media-email

There are additional ICON-related DN options that affect reporting. For full information about configuring DNs for ICON and Genesys Info Mart, see Configuring DN Objects.

Tip

For Genesys Info Mart–related options that you configure in a DN object, changes take effect on the next ETL cycle. The new option value is not applied to previously loaded facts.

Important

The ***-threshold** options described on this page do not control the thresholds used for aggregation. For information about the options that do control aggregation-specific thresholds, see the *Reporting and Analytics Aggregates Deployment Guide.*

gim-etl Section

- link-msf-userdata
- q-answer-threshold-voice
- q-short-abandoned-threshold-voice

link-msf-userdata

Default Value: false
Valid Values: true, false
Changes Take Effect: On the next ETL cycle
Introduced: 8.1.2
Related Options: link-msf-userdata-voice and link-msf-userdata-mm

Specifies whether associated user data will be stored in mediation segment facts (MSFs) for interactions that are in mediation at this DN (in other words, in this queue), including user data for interactions that are not distributed to a handling resource after the mediation.

- **true**—MSFs will store associated user data.

- **false**—MSFs will not store associated user data.

Setting **link-msf-userdata** at the DN level overrides the link-msf-userdata-voice and link-msf-userdata-mm options on the Genesys Info Mart Application object.

q-answer-threshold-voice

Default Value: 60
Valid Values: Any integer in the range of 1 to 10000
Changes Take Effect: On the next ETL cycle
Dependencies: None

On a DN of type ACD Queue or Virtual Queue, specifies the default duration, in seconds, that is used as a target time to answer voice interactions that were distributed by the virtual queue or ACD queue. This setting overrides the value of the option with the same name configured on the Application or Switch objects.

q-short-abandoned-threshold-voice

Default Value: 10
Valid Values: 1-1000
Changes Take Effect: On the next ETL cycle
Dependencies: None
Introduced: 8.5.003

On a DN of type ACD Queue or Virtual Queue, specifies the maximum duration of mediation, in seconds, that is used to indicate that an interaction that was abandoned while in a queue should be considered a "short" abandon. Genesys Info Mart uses this value to determine the state of SHORT_ABANDONED_FLAG in the MEDIATION_SEGMENT_FACT (MSF) row for voice interactions that are abandoned in this virtual queue or ACD queue.

This setting overrides a value that is set on the Application or Switch objects.

In releases earlier than 8.5.003, this value could not be set on an individual DN object.

For similar options that control equivalent thresholds for multimedia interactions, see the gim-etl-media-chat Section and the gim-etl-media-email Section.

gim-etl-media-chat Section

- q-answer-threshold
- q-short-abandoned-threshold

q-answer-threshold

Default Value: 60
Valid Values: 1-600000
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None

On a DN of type Virtual Queue, specifies the default duration, in seconds, that is used as a target time to accept a chat interaction that entered the queue.

This setting overrides the value of the option with the same name configured on the Application, tenant-specific Media Type Business Attribute, or Switch objects.

q-short-abandoned-threshold

Default Value: 10
Valid Values: 1-1000
Changes Take Effect: On the next ETL cycle
Dependencies: None
Introduced: 8.5.003

On a DN of type Virtual Queue, specifies the maximum duration of mediation, in seconds, that is used to indicate that a chat interaction that was abandoned while in the queue should be considered a "short" abandon. Genesys Info Mart uses this value to determine the state of SHORT_ABANDONED_FLAG in the MSF row for chat interactions that are abandoned in virtual queues.

This setting overrides the value of the option with the same name configured on the Application, tenant-specific Media Type Business Attribute, or Switch objects.

In releases earlier than 8.5.003, this value could not be set on an individual DN object.

gim-etl-media-email Section

- q-answer-threshold
- q-short-abandoned-threshold

q-answer-threshold

Default Value: 60
Valid Values: 1-600000
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None

On a DN of type Virtual Queue, specifies the default duration, in seconds, that is used as a target time to accept an email interaction that entered the queue. This setting overrides the value of the option with the same name configured on the Application, tenant-specific Media Type Business Attribute, or Switch objects.

q-short-abandoned-threshold

Default Value: 10
Valid Values: 1-1000
Changes Take Effect: On the next ETL cycle
Dependencies: None
Introduced: 8.5.003

On a DN of type Virtual Queue, specifies the maximum duration of mediation, in seconds, that is used to indicate that an email interaction that was abandoned while in the queue should be considered a "short" abandon. Genesys Info Mart uses this value to determine the state of SHORT_ABANDONED_FLAG in the MSF row for email interactions that are abandoned in virtual queues.

This setting overrides the value of the option with the same name configured on the Application, tenant-specific Media Type Business Attribute, or Switch objects.

In releases earlier than 8.5.003, this value could not be set on an individual DN object.

Script Options

Script objects specified in the contact center configuration define the interaction queues and interaction workbins used to manage multimedia interactions. The Genesys Info Mart–related options described on this page enable you to control reporting behavior at the level of individual interaction queues or workbins.

Script options that affect Genesys Info Mart reporting are contained in the following configuration sections, configured on the **Annex** of Script objects:

- gim-etl
- gim-etl-media-chat
- gim-etl-media-email
- gim-etl-populate

For more information about configuring Genesys Info Mart–related options for interaction queues and workbins, see Configuring Script Objects.

> ## Tip
> For Genesys Info Mart–related options that you configure in a Script object, changes take effect on the next ETL cycle. The new option value is not applied to previously loaded facts.

gim-etl Section

- link-msf-userdata

link-msf-userdata

Default Value: false
Valid Values: true, false
Changes Take Effect: On the next ETL cycle
Introduced: 8.1.2
Related Options: link-msf-userdata-mm

Specifies whether associated user data will be stored in mediation segment facts (MSFs) for interactions that are in mediation in this queue, including user data for interactions that are not distributed to a handling resource after the mediation.

- **true**—MSFs will store associated user data.

- **false**—MSFs will not store associated user data.

Setting **link-msf-userdata** at the Script level overrides the link-msf-userdata-mm option on the Genesys Info Mart Application object.

gim-etl-media-chat

- q-answer-threshold
- q-short-abandoned-threshold

q-answer-threshold

Default Value: 60
Valid Values: 1-600000
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None

For the interaction queue or interaction workbin defined by the Script object, specifies the default duration, in seconds, that is used as a target time to accept a chat interaction that entered the interaction queue or workbin.

This setting overrides the value of the option with the same name configured on the Application or tenant-specific Media Type Business Attribute objects.

q-short-abandoned-threshold

Default Value: 10
Valid Values: 1-1000
Changes Take Effect: On the next ETL cycle
Dependencies: None
Introduced: 8.5.003

For the interaction queue or interaction workbin defined by the Script object, specifies the maximum duration of mediation, in seconds, that is used to indicate that a chat interaction that was abandoned while in a queue should be considered a "short" abandon. Genesys Info Mart uses this value to determine the state of SHORT_ABANDONED_FLAG in the MSF row for chat interactions that are

abandoned in an interaction queue or workbin.

This setting overrides the value of the option with the same name configured on the Application or tenant-specific Media Type Business Attribute objects.

In releases earlier than 8.5.003, this value could not be set on an individual Script object.

gim-etl-media-email

- q-answer-threshold
- q-short-abandoned-threshold

q-answer-threshold

Default Value: 60
Valid Values: 1-600000
Changes Take Effect: At the next run of Job_TransformGIM
Dependencies: None

For the interaction queue or interaction workbin defined by the Script object, specifies the default duration, in seconds, that is used as a target time to accept an email interaction that entered the interaction queue or workbin.

This setting overrides the value of the option with the same name configured on the Application or tenant-specific Media Type Business Attribute objects.

q-short-abandoned-threshold

Default Value: 10
Valid Values: 1-1000
Changes Take Effect: On the next ETL cycle
Dependencies: None
Modified: 8.5.003

For the interaction queue or interaction workbin defined by the Script object, specifies the maximum duration of mediation, in seconds, that is used to indicate that an interaction that was abandoned while in a queue should be considered a "short" abandon. Genesys Info Mart uses this value to determine the state of SHORT_ABANDONED_FLAG in the MSF row for email interactions that are abandoned in an interaction queue or workbin.

This setting overrides the value of the option with the same name configured on the Application or tenant-specific Media Type Business Attribute objects.

In releases earlier than 8.5.003, this value could not be set on an individual Script object.

gim-etl-populate

- populate-mm-ixnqueue-facts
- populate-mm-workbin-facts

populate-mm-ixnqueue-facts

Default Value: false
Valid Values: true, false
Changes Take Effect: On the next ETL cycle

For the Interaction Queue defined by the Script object, enables or disables the population of eServices/Multimedia queue activity to the MSF table. This setting overrides a value that is set on the Genesys Info Mart Application object.

populate-mm-workbin-facts

Default Value: true
Valid Values: true, false
Changes Take Effect: On the next ETL cycle
Dependencies: populate-workbin-as-hold=false

For the Interaction Workbin defined by the Script object, enables or disables the population of eServices/Multimedia workbin activity to the MSF table. For workbins that are associated with handling resources of type Agent or Place, this option comes into effect only if Genesys Info Mart has not been configured to consider workbin time as hold. For the circumstances under which Genesys Info Mart considers workbin time as hold, see the description of the **populate-workbin-as-hold** option.

This setting overrides a value that is set on the Genesys Info Mart Application object.

Field Options

This page describes the Genesys Info Mart–related options you configure on the **Annex** tab of **Field** objects for Outbound Contact campaigns. Field options that affect Genesys Info Mart reporting are contained in the gim-etl-mapping Section and in the default Section.

gim-etl-mapping Section

- column-name
- table-name

column-name

Default Value: No default value
Valid Values: Any valid column name in the table specified by the table-name option
Changes Take Effect: On the next ETL cycle
Dependencies: table-name

Specifies the name of the column in the selected Info Mart table in which the ETL should store the value of this field. The option value is case insensitive.

Refer to the Mapping OCS Record Fields Worksheet in the *Genesys Info Mart Deployment Guide* for the list of column names.

The type of the OCS field must match the data type of the column.

Note: Configure this option only for the nonmandatory Field objects that you want to store in the Info Mart database.

Genesys recommends that you specify the name of a column (in the table that is specified by the table-name option) that is not used in the configuration of any other Field object.

table-name

Default Value: No default value
Valid Values: CONTACT_ATTEMPT_FACT, RECORD_FIELD_GROUP_1, RECORD_FIELD_GROUP_2
Changes Take Effect: On the next ETL cycle
Dependencies: The **icon_attribute** option has been set in the [default] section.

Specifies the name of the Info Mart table in which the ETL should store the value of this field. The option value is case insensitive.

You can specify storage in one of the following tables:

- CONTACT_ATTEMPT_FACT
- RECORD_FIELD_GROUP_1
- RECORD_FIELD_GROUP_2

Note: Configure this option only for the nonmandatory Field objects that you want to store in the Info Mart database.

default Section

In addition to the Genesys Info Mart–related Field options configured in the **[gim-etl-mapping]** section, options in the **[default]** section on the **Annex** tab of **Field** objects affect the storage of Outbound Contact data in ICON and, therefore, in Genesys Info Mart. For more information, see Configuring Field Objects in the *Genesys Info Mart Deployment Guide*.