



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

Genesys Info Mart Deployment Guide

Database Considerations

12/17/2025

Database Considerations

This page describes database issues that you must consider before you deploy Genesys Info Mart.

Databases in Your Genesys Info Mart Deployment

Genesys Info Mart has predefined jobs that process data in extract, transform, and load (ETL) cycles. These ETL jobs access several databases, as described in [Genesys Info Mart Jobs](#). There are also several supported database schemas (Genesys Info Mart and tenant user schemas), as described in [Info Mart Database](#).

In addition to the topics that are discussed on this page, there are several deployment-specific Genesys Info Mart database design considerations that are outside the scope of this *Deployment Guide*, including partitioning, indexing and storage. To develop a suitable physical database design and implementation for your environment, consult your database administrator or data warehousing specialist.

Important

- When you install Genesys Info Mart, you select a single RDBMS type: Microsoft SQL Server, Oracle, or PostgreSQL. All source IDBs and the target Info Mart database must reside in databases of this same RDBMS type. The only possible exception is the IDB from which Genesys Info Mart extracts configuration history.
- Genesys Info Mart does not support database compression, except for compression that RAA may implement. Do not remove any compression that has been defined by RAA.
- Be sure to consult the [Genesys Info Mart Product Alerts](#) for important information about known RDBMS issues and potential solutions to those issues.

Database Capacity

As described in [Genesys Info Mart Jobs](#), Genesys Info Mart reads from and writes to the Info Mart database. To determine the database capacity requirements for the Info Mart database in your environment, answer the following questions:

- How much space does the database require?
- How much space is needed for future growth?
- How powerful should the database server be?
- How do you plan to use the Genesys Info Mart database?

- How do you plan to manage the Genesys Info Mart indexes?
- What will your purging strategy be?

Use the following information to help answer these questions.

Storage Capacity

The Genesys Info Mart database must have the capacity to store the row data extracted from the required number of IDBs, as well as the facts and dimensions that the ETL jobs load. The Genesys Info Mart database grows over time, because the ETL jobs load new facts and dimension values each day.

If your deployment includes the Genesys historical reporting presentation layer (GCXI) or the RAA package, you will also need to provide storage for the aggregate tables and indexes that you create in order to improve query performance, and to provide storage for an extended period of time that suits your requirements. See the [RAA documentation set](#) for more information.

Genesys Info Mart provides a maintenance job that purges data in the Info Mart database. The maintenance job automatically purges data in accordance with configurable data-retention policies. For a list of the tables the maintenance job purges, see [Info Mart Tables Purged by the Maintenance Job](#) in the *Genesys Info Mart Operations Guide*. For information about the configuration options that control data retention, see the **days-to-keep-*** option descriptions on the [gim-etl Section](#) page in the *Genesys Info Mart Configuration Options Reference*.

In PostgreSQL deployments, a job to update statistics (**Job_UpdateStats**) performs important aspects of database maintenance that improve query performance.

Genesys provides an interactive tool, the [Genesys Info Mart 8.5 Database Size Estimator](#), to help you estimate the size of your Info Mart database.

Important

Note for Oracle deployments

Genesys Info Mart 8.5 makes extensive use of temporary tables during data collection and processing, executing frequent CREATE/DROP TABLE statements. As a result, the SYSTEM tablespace increases over time. Ensure that you allocate sufficient space for the SYSTEM tablespace to accommodate long-term operations. Monitor the growth of the SYSTEM tablespace, and extend it as required.

Processing Capacity

The ETL jobs perform many intensive SQL operations against extracted data in the Info Mart database, including SELECT, INSERT, and UPDATE. These operations require significant resources, such as disk (for tables, indexes, and logs), memory, and CPU capacity.

The ETL jobs load data in the Info Mart database at the end of each ETL cycle. The amount of time that the ETL jobs run varies, depending on how often you schedule them and on the volume of data that they process. The ETL jobs do not create or update statistics on the Genesys Info Mart fact tables.

If you intend to use the Genesys Info Mart database as the database that your business applications query, provide additional capacity so that many users can query the data.

If you intend to upload Genesys Info Mart data to a data warehouse, instead of having users query the data directly, you do not need to have the capacity to support many users who query the data. You will probably require fewer indexes, and will probably store the data for less than a year.

Important

Adding multiple indexes to those fact tables in the Info Mart database that contain data relating to interactions or resources can have a significant negative effect on the performance of **Job_TransformGIM**. Genesys recommends that you first test the impact of additional indexes in a non-production environment.

Database Partitioning

Genesys Info Mart supports the use of partitioning in the Info Mart database in Oracle (range partitioning only), Microsoft SQL Server, and PostgreSQL deployments. In general, if partitioning is used, fact tables and associated indexes in GIDB and the dimensional model are partitioned, as well as certain Control tables. Configuration object tables and configuration relationship fact tables are not partitioned, and all dimension tables are also not partitioned.

If your contact center is large and operates at high volumes, Genesys strongly recommends that you implement partitioning. Otherwise, maintenance of a nonpartitioned database can significantly affect performance.

For each applicable RDBMS, Genesys Info Mart provides a separate script to create the partitioned database schema. The purpose of the script is to identify the tables that are partitioned. The script creates the partitions in the primary filegroup or tablespace. The script creates a single, throwaway partition for each partitioned database object; these partitions are purged during the first run of **Job_MaintainGIM**. During initialization, **Job_InitializeGIM** creates the first set of partitions to be populated during the first ETL cycle, and **Job_MaintainGIM** subsequently creates additional partitions as required.

Configuration options enable you to specify the size of the partitions in GIDB (see partitioning-interval-size-gidb, partitioning-interval-size-gidb-mm, and partitioning-interval-size-gidb-ocs) as well as in the dimensional model and Control tables (see partitioning-interval-size-gim). Another configuration option (partitioning-ahead-range) enables you to control how far ahead the Genesys Info Mart jobs will create partitions, in preparation for future ETL cycles. (**Job_InitializeGIM** creates the partitions in the first instance, then **Job_MaintainGIM** creates them on an ongoing basis.)

Important

Maintenance job scheduling — By default, Genesys Info Mart runs the maintenance job daily. If your database is partitioned, ensure that you do not jeopardize routine maintenance of the partitions by inappropriately changing the configuration options that control scheduling of the ETL cycle and the maintenance job. For more information about the scheduling-related configuration options, see [schedule Section](#) in the *Genesys Info Mart Options Reference*.

PostgreSQL partitioning recommendations — For PostgreSQL, Genesys recommends that you review and modify the values of the **partitioning-interval-size-*** options to increase partition sizes, as described in the option descriptions (see links above).

Multi-Language Support

Genesys Info Mart supports using Unicode characters to store data in multiple languages, starting with release 8.5.003 in Oracle and PostgreSQL deployments and starting with release 8.5.007 in Microsoft SQL Server deployments.

To support Unicode characters in the Info Mart database:

- In Oracle and PostgreSQL deployments, the Info Mart database must be created with UTF-8 encoding. No other special configuration is required. Genesys Info Mart does not support the NVARCHAR2 data type on Oracle. The Oracle schema-creation scripts use explicit CHAR character length semantics in fields with VARCHAR2 data types, to enable consistent storage of Unicode characters.
For more information about National Language Support (NLS) in Oracle, see the *Database Globalization Support Guide* for your Oracle release, available from the [Oracle Help site](#).
- In Microsoft SQL Server deployments, you must use the multi-language versions of the database-creation scripts (**make_gim_multilang.sql** or **make_gim_multilang_partitioned.sql**) to create the Info Mart schema. These scripts use nvarchar instead of varchar data types for fields that store Unicode characters.

For information about creating the Info Mart database schema, see [Info Mart Database Scripts](#).

In addition, there are important requirements for Configuration Layer components and Interaction Concentrator. For full information, see:

- [Configuring for Multi-Language Support](#) in the *Interaction Concentrator Deployment Guide*
- [Framework Database Connectivity Reference Guide](#)
- [Deploying Genesys for Key Mixed Language Scenarios](#)

After the Info Mart database schema has been initialized in Microsoft SQL Server deployments, a flag in the CTL_SCHEMA_INFO table (CTL_SCHEMA_INFO.ML_FLAG=1) identifies whether the database supports multiple languages.

Database Connections

Genesys Info Mart processing is highly multi-threaded. The Genesys Info Mart jobs open multiple connections to the Info Mart database and, for extraction, to IDBs to enable the jobs to process data from multiple tables concurrently.

For example, the extraction job extracts concurrently from multiple IDBs and from multiple tables from each IDB. The execution plan for extracting data from a particular IDB is not executed in strict sequence. Instead, worker threads open connections as required to process execution plan items concurrently. The connection for extracting from a table with a lot of data may still be active after the work item for a table with very little data has been executed and the second connection is no longer active, even if it was opened later than the first one.

Starting with release 8.5.009.14, Genesys Info Mart uses connection pooling to improve extraction performance, by reusing idle connections from the pool to extract from additional tables in the same IDB. Paradoxically, the use of a connection pool might appear to increase the number of connections used during extraction, because connections are not closed when extraction from a particular IDB is finished. Instead, all connections opened during a run of the extraction or transformation job remain available in the connection pool for the life of the Genesys Info Mart instance. Subsequent runs of the extraction job can reuse pooled connections to IDB, and the extraction and transformation jobs can both reuse connections to the Info Mart database.

Note, however, that a large connection pool does not usually imply significantly increased demand on RDBMS resources, because many of the connections will be idle for much of the time the job takes to execute. Starting with release 8.5.015.07, you can use the `cp-reuse-count` option to limit the number of times each connection can be reused if you find that Genesys Info Mart connection-pooling behavior results in issues such as excessive memory growth on the DBMS side.

To enable Genesys Info Mart to open sufficient connections during job execution, it is important that you tune your RDBMS instance(s) for both the Info Mart database and your IDBs (see [Optimizing Database Performance: Database Tuning](#)). In addition, ensure that the Genesys Info Mart configuration options that control numbers of connections and threads (see [Performance tuning](#) for a summary list) are suitable for your deployment.

Database Security

Genesys Info Mart supports the following database security features:

- Secure Socket Layer (SSL) connections to the Info Mart database and IDB(s) to encrypt communications between Genesys Info Mart Server and its source and target databases. (PostgreSQL deployments require Genesys Info Mart release 8.5.011.18 to support SSL.) For a summary of the steps to enable this feature, see [Enabling Secure Communications](#).
- Encryption of data in the Info Mart database in Microsoft SQL Server and Oracle deployments. For more information, see [Encrypted Data in Databases](#) in the *Genesys Security Deployment Guide*.

Source Data Retention and Purging

Genesys Info Mart does not automatically purge source data in IDB. However, Genesys does provide stored procedures that are recommended for source database purging. When you use these procedures, be sure to:

- Avoid deleting data that has not yet been extracted.
- Retain enough data to allow for error recovery and problem determination.

The amount of data that you should retain in your source databases depends on both the database server's hardware resources — such as memory and disk space — and the performance of its disk subsystems.

Generally, you should aim to achieve a balance that enables you to retain enough data in your Interaction Concentrator databases without affecting either the operating performance of your source database or the extraction process of Genesys Info Mart.

For Voice details and Outbound Contact details, Genesys recommends that you consider using IDB partitioning, which is supported for Oracle deployments starting with Interaction Concentrator release 8.1.1. Partitioning improves overall IDB performance and streamlines maintenance. However, given the way that partitioning support has been implemented, with only a limited number of partitions expected to be used, and given the long-living nature of multimedia interactions, Genesys recommends that you not use partitioned IDBs for Multimedia details. For more information about IDB partitioning, see the [Interaction Concentrator Deployment Guide](#).

Genesys provides specific recommendations regarding source data retention and purging frequency. For more information, see [Managing Interaction Concentrator and Data Sources](#) in the *Genesys Info Mart Operations Guide*.

Database User Authentication

The ETL jobs make many database connections as they extract, transform, and load data. To ensure that connections are authenticated quickly, review the authentication policy that is configured in your database software. Authentication timeouts can greatly increase the amount of time that it takes for the ETL jobs to run to completion.

For more information about database users, see [Database Privileges](#).