



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

Optimizing Database Performance: Database Tuning

Optimizing Database Performance: Database Tuning

The performance of the Genesys Info Mart jobs is greatly affected by the performance of the RDBMS server. Before you create the target database, tune your RDBMS server for optimal performance.

This page provides recommended parameter settings and tuning guidelines that you can use to improve the Info Mart ETL execution time. It also includes suggestions for your database administrator for managing the target database after it is deployed.

Important

If your IDBs and Info Mart database are not hosted by the same RDBMS instance, ensure that you tune your IDB RDBMS server(s) to provision sufficient connections for the Genesys Info Mart extraction job. See [Note about connections for extraction](#), below.

Tuning the Info Mart Database

There are several database parameters that enable the ETL jobs to load the Info Mart database successfully. This section provides the recommended parameter settings for each RDBMS:

- [Tuning the Info Mart database on Microsoft SQL Server](#)
- [Tuning the Info Mart database on Oracle](#)
- [Tuning the Info Mart database on PostgreSQL](#)

See also [Additional Considerations](#).

Procedure: Tuning the Info Mart database on Microsoft SQL Server

Purpose: To set Microsoft SQL Server RDBMS parameters so that the ETL jobs load the Info Mart database successfully.

Prerequisites

- Create the Info Mart database schema, as instructed in [Preparing the Info Mart Database](#).
- Use Microsoft SQL Server Enterprise Manager to review the settings of the Microsoft SQL Server properties.

Steps

1. Allocate sufficient memory.

The ETL jobs issue many complex SQL queries against several Info Mart database tables. The amount of memory that you allocate to the database server is critical to the performance of these SQL queries. Allocate as much memory as possible to Microsoft SQL Server without causing paging. Optimal settings for your environment depend on the hardware and data volumes. As a quick approximation:

- a. Start with the total real memory on the database server.
- b. Subtract the memory that is required by the operating system and any other applications.
- c. Set the maximum memory that is allocated to Microsoft SQL Server to the result.

2. Select relevant server settings.

Some of the SQL commands that are issued by the ETL jobs are long running. To ensure that their cost is not limited by Microsoft SQL Server, disable **Use query governor to prevent queries exceeding specified cost**.

3. Configure connections.

The ETL jobs use many concurrent database connections. Set **Maximum concurrent user connections** to 0 (unlimited), but see also [Note about connections for extraction](#), below.

4. Review the settings of the following database properties:

- Data Files: Select **Automatically grow file** and **Unrestricted file growth**.
- Transaction Log: Select **Automatically grow file** and **Unrestricted file growth**.
- Options — Settings: Select **Auto update statistics** and **Auto create statistics**.
- Options — Recovery Model: Take into account the following considerations.

The ETL jobs generate large amounts of database activity, with a correspondingly large database log space. The amount of log space that is needed depends on the recovery model and on the frequency of log file backups. The full-recovery model provides the most protection against data loss due to failures, but it requires the most log space. Performing daily log file backups can help limit the size of the transaction logs.

In general, you control the size of a transaction log in one of the following ways:

- If you are maintaining a log backup sequence for full or bulk-logged recovery models, schedule **BACKUP LOG** statements to occur at intervals that will keep the transaction log from growing past the desired size.
- If you are not maintaining a log backup sequence, specify the simple recovery model.

For a more complete description of the recovery models, and for information about how to manage transaction logs, see the Microsoft SQL Server documentation.

5. (Optional, but recommended) Configure the database to use the **READ COMMITTED** isolation level.

In Microsoft SQL Server deployments, the transformation job might record duplicate rows in Info Mart fact tables because

of the inability of Microsoft SQL Server to finish rollback correctly when deadlock happens during execution of rollback. To minimize locking contention, Genesys recommends that you use the **READ COMMITTED** isolation level, with the **READ_COMMITTED_SNAPSHOT** database option set to ON.

To configure this database setting, execute the following SQL statement:

```
ALTER DATABASE <name_of_gim_db> set READ_COMMITTED_SNAPSHOT ON;
```

6. Consult with your database administrator to further fine-tune these and other parameters if you find the performance of the ETL jobs in your environment unacceptable. See also [Additional Considerations](#).

Next Steps

Configure the required DAPs. For more information, see [Configuring Required DAPs](#).

Procedure: Tuning the Info Mart database on Oracle

Purpose: To set Oracle RDBMS parameters so that the ETL jobs load the Info Mart database successfully.

Prerequisites

- Create the Info Mart database schema, as instructed in [Preparing the Info Mart Database](#).

Steps

1. Set the Oracle initialization parameters:
 - `filesystemio_options = ASYNCH`
 - `processes = at least 1000`, but see also [Note about connections for extraction](#), below
 - `sessions = at least 1000`, but see also [Note about connections for extraction](#), below
 - `open_cursors = at least 1000`
 - (For use with database links) `open_links = at least the value of extract-data-thread-pool-size`.

If your deployment uses database links, configure the **open_links** parameter to allow a sufficient number of connections through the database links. For more information about how to configure your deployment to use database links, see [Optimizing Database Performance: Database Links](#).

2. Allocate sufficient memory.

The ETL jobs issue many complex SQL queries against several Info Mart database tables. The amount of memory that you allocate to the database server buffers is critical to the performance of these SQL queries. The optimal settings for your environment depend on the hardware and data volumes. As a quick approximation:

- a. Start with the total real memory on the database server.
 - b. Subtract the amount of memory that is required by the operating system and any other applications.
 - c. Split the result between the buffer cache and the PGA aggregate target.
3. Increase the size of the listener queue for TCP/IP connection requests.

The need to modify the listener configuration depends on deployment-specific factors such as hardware and the configured or desired degree of parallelism. See the Oracle documentation for information about the conditions under which you should increase the queue size in the listener configuration file, **listener.ora**, as well as the syntax for doing so. See also [Performance tuning](#) for a summary of Genesys Info Mart configuration options that control the degree of parallelism for ETL processing.

4. Genesys recommends that you use case-sensitive settings for **NLS_SORT** and **NLS_COMPARE**.
5. Consult with your database administrator to further fine-tune these and other parameters, including JVM startup parameters such as **queryParallelism** (see [Modifying JVM Startup Parameters](#)), if you find the performance of the ETL jobs in your environment unacceptable. See also [Additional Considerations](#).

Next Steps

Configure the required DAPs. For more information, see [Configuring Required DAPs](#).

Procedure: Tuning the Info Mart database on PostgreSQL

Purpose: To set PostgreSQL RDBMS parameters so that the ETL jobs load the Info Mart database successfully.

Prerequisites

- Create the Info Mart database schema, as instructed in [Preparing the Info Mart Database](#).

Steps

1. Set the database server configuration parameters in the **postgresql.conf** file:
 - `checkpoint_segments = 64`

- synchronous_commit = off
- wal_buffers = 256MB
- max_connections = 1000, but see also [Note about connections for extraction](#), below
- constraint_exclusion = partition
- default_statistics_target = 100
- enable_mergejoin = off
- temp_buffers = 128MB
- max_prepared_transactions = 1000
- work_mem = A value in the range 256MB-1GB*
- maintenance_work_mem = A value in the range 64MB-256MB*

*For smaller available memory and smaller workloads, use a lower value in the range.

2. Allocate sufficient memory.

The ETL jobs issue many complex SQL queries against several Info Mart database tables. The amount of memory that you allocate to the database server buffers is critical to the performance of these SQL queries. The optimal settings for your environment depend on the hardware and data volumes. As a quick approximation:

- a. Start with the total real memory on the database server.
- b. Subtract the amount of memory that is required by the operating system and any other applications. The result is the available memory for PostgreSQL.
- c. Set **effective_cache_size** to 75 percent of available memory and **shared_buffers** to 25 percent of available memory.

3. Configure autovacuum settings.

Autovacuum is a background process that performs several important tasks, including producing table statistics for the query planner, recovering disk space for changed rows, and preventing transaction ID wraparound. To balance these tasks with the needs of the ETL jobs, set the following combination of autovacuum parameters in the **postgresql.conf** file:

- autovacuum = on
- autovacuum_analyze_threshold = 10000
- autovacuum_freeze_max_age = 1000000000
- autovacuum_max_workers = 10
- autovacuum_naptime = 20s
- autovacuum_vacuum_cost_delay = 10ms
- autovacuum_vacuum_cost_limit = 1000
- autovacuum_vacuum_scale_factor = 0.3
- autovacuum_vacuum_threshold = 100000
- vacuum_freeze_min_age = 10000000

- `vacuum_freeze_table_age = 800000000`

A Genesys Info Mart job, `Job_UpdateStats`, invokes PostgreSQL functionality to supplement the autovacuum process. Ensure that you also configure the Genesys Info Mart application to run `Job_UpdateStats` regularly. For more information, see `run-update-stats` and `update-stats-schedule`.

4. Consult your database administrator to further fine-tune these and other parameters if you find the performance of the ETL jobs in your environment unacceptable. See also [Performance tuning](#) for a summary of Genesys Info Mart configuration options that control the degree of parallelism for ETL processing, as well as [Additional Considerations](#).

Next Steps

Configure the required DAPs. For more information, see [Configuring Required DAPs](#).

Note about connections for extraction

For reasons described under [Database Connections](#), the extraction job in particular uses a large number of connections to the Info Mart database and IDBs. Genesys strongly recommends the connection-related settings specified above. However, if you want to refine your estimates of resource usage, the following calculation approximates the theoretical maximum number of processes and sessions (for Oracle) or connections (for Microsoft SQL Server and PostgreSQL) the extraction job might use.

```
min(80,extract-data-thread-pool-size) * Number of IDBs + 10 + extract-data-thread-pool-size
```

where:

- `extract-data-thread-pool-size` is the value of the `extract-data-thread-pool-size` configuration option.
- Number of IDBs equates to the number of extraction DAPs in the Genesys Info Mart application's connections.
- The formula adds:
 - 10 connections used for merging extracted data, as well as for other ETL needs.
 - The **`extract-data-thread-pool-size`** value, to account for the connections to the Info Mart database that are used to write data extracted from IDBs into the Info Mart database.

If the IDBs and Info Mart database are hosted by the same RDBMS instance, double the number of IDB connections.

Additional Considerations

In addition to the previously listed database tuning requirements, you might also need to do the

following:

- Periodically update statistics on the Info Mart fact tables. Failure to update them periodically can have a negative impact on the performance of end-user queries. Enable the automatic gathering of statistics on Info Mart tables if your RDBMS supports this feature.
- Have database administrators actively manage Genesys Info Mart after it is deployed.