GENESYS
AN ALCATEL·LUCENT COMPANY

Genesys Info Mart 8.0

Operations Guide

About Genesys

Alcatel-Lucent's Genesys solutions feature leading software that manages customer interactions over phone, Web, and mobile devices. The Genesys software suite handles customer conversations across multiple channels and resources—self-service, assisted-service, and proactive outreach—fulfilling customer requests and optimizing customer care goals while efficiently using resources. Genesys software directs more than 100 million customer interactions every day for 4000 companies and government agencies in 80 countries. These companies and agencies leverage their entire organization, from the contact center to the back office, while dynamically engaging their customers. Go to www.genesyslab.com for more information.

Each product has its own documentation for online viewing at the Genesys Technical Support website or on the Documentation Library DVD, which is available from Genesys upon request. For more information, contact your sales representative.

Notice

Although reasonable effort is made to ensure that the information in this document is complete and accurate at the time of release, Genesys Telecommunications Laboratories, Inc., cannot assume responsibility for any existing errors. Changes and/or corrections to the information contained in this document may be incorporated in future versions.

Your Responsibility for Your System's Security

You are responsible for the security of your system. Product administration to prevent unauthorized use is your responsibility. Your system administrator should read all documents provided with this product to fully understand the features available that reduce your risk of incurring charges for unlicensed use of Genesys products.

Trademarks

Genesys, the Genesys logo, and T-Server are registered trademarks of Genesys Telecommunications Laboratories, Inc. All other trademarks and trade names referred to in this document are the property of other companies. The Crystal monospace font is used by permission of Software Renovation Corporation, www.SoftwareRenovation.com.

Technical Support from VARs

If you have purchased support from a value-added reseller (VAR), please contact the VAR for technical support.

Technical Support from Genesys

If you have purchased support directly from Genesys, please contact Genesys Technical Support at the regional numbers provided on . For complete contact information and procedures, refer to the *Genesys Technical Support Guide*.

Ordering and Licensing Information

Complete information on ordering and licensing Genesys products can be found in the *Genesys Licensing Guide.*

Released by

Genesys Telecommunications Laboratories, Inc. www.genesyslab.com

Document Version: 80gim_op_02-2011_v8.0.101.00

Table of Contents

Preface

Welcome to the *Genesys Info Mart 8.0 Operations Guide.* This guide describes the procedures that you must follow to schedule and monitor the Genesys Info Mart jobs that extract, transform, and load (ETL) data. It is intended for system administrators and is valid only for Genesys Info Mart release 8.0.

Note: For versions of this document created for other releases of this product, visit the Genesys Technical Support website, or request the Documentation Library DVD, which you can order by e-mail from Genesys Order Management at `orderman@genesyslab.com`.

This preface contains the following sections:

For information about related resources and about the conventions that are used in this document, see the supplementary material starting on page 83.

About Genesys Info Mart

Genesys Info Mart produces a data mart that you can use for contact center historical reporting.

Genesys Info Mart includes a server component, administration graphical user interface (GUI), and database. The Genesys Info Mart server runs a set of predefined jobs that execute extract, transform, and load (ETL) processes to:

- Extract data that has been gathered by Interaction Concentrator from data sources such as Configuration Server, T-Server, Interaction Server, and Outbound Contact Server. Genesys Info Mart stores this low-level interaction data, which is consolidated from Interaction Concentrator databases (Interaction Databases [IDBs]), in the Info Mart database.

- Transform the low-level interaction data and load it into a dimensional model (or star schemas) in the Info Mart database.

Genesys Info Mart can also be configured to host an aggregation engine that aggregates or re-aggregates the data, and populates Aggregate tables in the Info Mart database.

You query the Fact and Dimension tables in the dimensional model, using Structured Query Language (SQL), to obtain results that enable you to examine the data in detail, identify patterns, and predict trends for your organization.

Intended Audience

This guide is primarily intended for database administrators and system administrators. The guide assumes that you have a basic understanding of the following topics:

- Relational database management systems (RDBMSs)
- Data extraction
- Data warehousing
- Data integration
- Network design and operation
- Your network and database configurations
- SQL (Structured Query Language)

Computer-telephony integration (CTI) concepts, processes, terminology, and applications

Making Comments on This Document

If you especially like or dislike anything about this document, feel free to e-mail your comments to `Techpubs.webadmin@genesyslab.com`.

You can comment on what you regard as specific errors or omissions, and on the accuracy, organization, subject matter, or completeness of this document. Please limit your comments to the scope of this document only and to the way in which the information is presented. Contact your Genesys Account Representative or Genesys Technical Support if you have suggestions about the product itself.

When you send us comments, you grant Genesys a nonexclusive right to use or distribute your comments in any way it believes appropriate, without incurring any obligation to you.

Contacting Genesys Technical Support

If you have purchased support directly from Genesys, contact Genesys Technical Support at the following regional numbers:

Region	Telephone	E-Mail
North America and Latin America	+888-369-5555 (toll-free) +506-674-6767	support@genesyslab.com
Europe, Middle East, and Africa	+44-(0)-1276-45-7002	support@genesyslab.co.uk
Asia Pacific	+61-7-3368-6868	support@genesyslab.com.au
Malaysia	1-800-814-472 (toll-free) +61-7-3368-6868 (International)	support@genesyslab.com.au
India	000 800 100 7136 (toll-free) +61-7-3368-6868 (International)	support@genesyslab.com.au
Japan	+81-3-6361-8950	support@genesyslab.co.jp
Before contacting technical support, refer to the *Genesys Technical Support Guide* for complete contact information and procedures.		

Document Change History

This section lists topics that are new or that have changed significantly since the first release of this document.

New in Document Version v8.0.101.00

The document has been updated to support Genesys Info Mart release 8.0.1. The following topics have been added or significantly changed since the previous version of this document:

- All information about database partitioning has been modified to include Microsoft SQL Server as a platform on which Genesys Info Mart supports partitioning of the Info Mart database.
- A new section in Chapter 2, "Job_MigrateGIM" on page 43, describes the new migration job. Information about the new job has also been added to Table 1 on page 18.

- Genesys Info Mart does not support deployments in which voice and multimedia interactions share the same virtual queues. In the subsection "Delayed and Early-Arriving Virtual Queue Data" on page 33, a former recommendation that you configure separate virtual queues for voice and multimedia interactions has been changed to a requirement.

- Information about when Genesys Info Mart performs a configuration check ("Configuration Checker" on page 45) has been updated to include other contact center objects, such as DNs, that can be configured to override Genesys Info Mart `Application` settings.

- The troubleshooting chapter has been updated to include information about a unique constraint violation error that the transformation job might encounter, related to the Genesys Info Mart support for 3[rd] Party Media. For more information, see "SQL Error" on page 73.

- A new section, "Changing Calendar Dimensions" on page 61, has been added to the chapter about working with Genesys Info Mart jobs to describe how you can safely change the values of the options that control custom calendar dimensions.

1 Introducing Genesys Info Mart 8.0

This chapter introduces you to Genesys Info Mart 8.0. It includes a brief overview and a list of new features available in this release:

About Genesys Info Mart 8.0

Genesys Info Mart 8.0 extracts data from one or more Genesys Interaction Concentrator databases to produce the Info Mart Database, your data store for contact center historical reporting.

In the release 8.0 model, Genesys Info Mart consolidates data from multiple Interaction Databases (IDBs) in Global Interaction Database (GIDB), which is part of the Info Mart schema. Data extracted from one or more IDBs passes through Merge tables (within the Info Mart database), where the merge operation combines related interactions together. The data is then moved into the GIDB tables where it is further processed before it is finally stored in Info Mart's fact and Dimension tables (*dimensional model*).

illustrates the Genesys Info Mart 8.0 architecture and the primary data flow between the Genesys Info Mart components and the other Genesys components. (The diagram does not depict high availability architecture for any components.)

Terminology Definition: ICON Details

The term, *ICON details* refers to the type of reporting data that Genesys Info Mart extracts from one or more IDBs that are populated by one or more Interaction Concentrator (ICON) applications.

Depending on the way it is configured, Genesys Info Mart stores the following types of data:

- ICON Configuration details, which include:
  - Configuration objects (such as a `DN, Person, Skill, or Place`)
  - Configuration object relationships (that is, associations between configuration objects, such as a `Person` assignment to a `Group`)

  The ICON data source for Configuration details is Configuration Server.

- ICON Voice details, which include:
  - Voice interaction data
  - Voice attached data and UserEvent-based key-value pair (KVP) data. (The combination of the two is also referred to as *user data*.)
  - Voice agent login data
  - Voice agent state and agent state reason details, including the ability to associate after-call-work with voice interactions
  - Voice DND mode details
  - Virtual queue data

  The ICON data source for Voice details is T-Server.

- ICON Multimedia details, which include:
  - Multimedia interaction data
  - Multimedia attached data
  - Multimedia agent login data
  - Multimedia agent state and agent state reason details
  - Virtual queue data

  The ICON data source for Multimedia details is Interaction Server.

- Outbound Contact details, which include:
  - History and results of campaigns, chains, and contact attempts
  - Associations between Outbound Contact objects (such as campaigns) and contact center objects (such as agent groups or place groups)
  - Precalculated Outbound Contact metrics

  The ICON data source for Outbound Contact details is Outbound Contact Server (OCS).

Note: For information about the meaning of other terms, such as *Voice details,* see the "Terminology Conventions" section in the *Genesys Info Mart 8.0 Deployment Guide*.
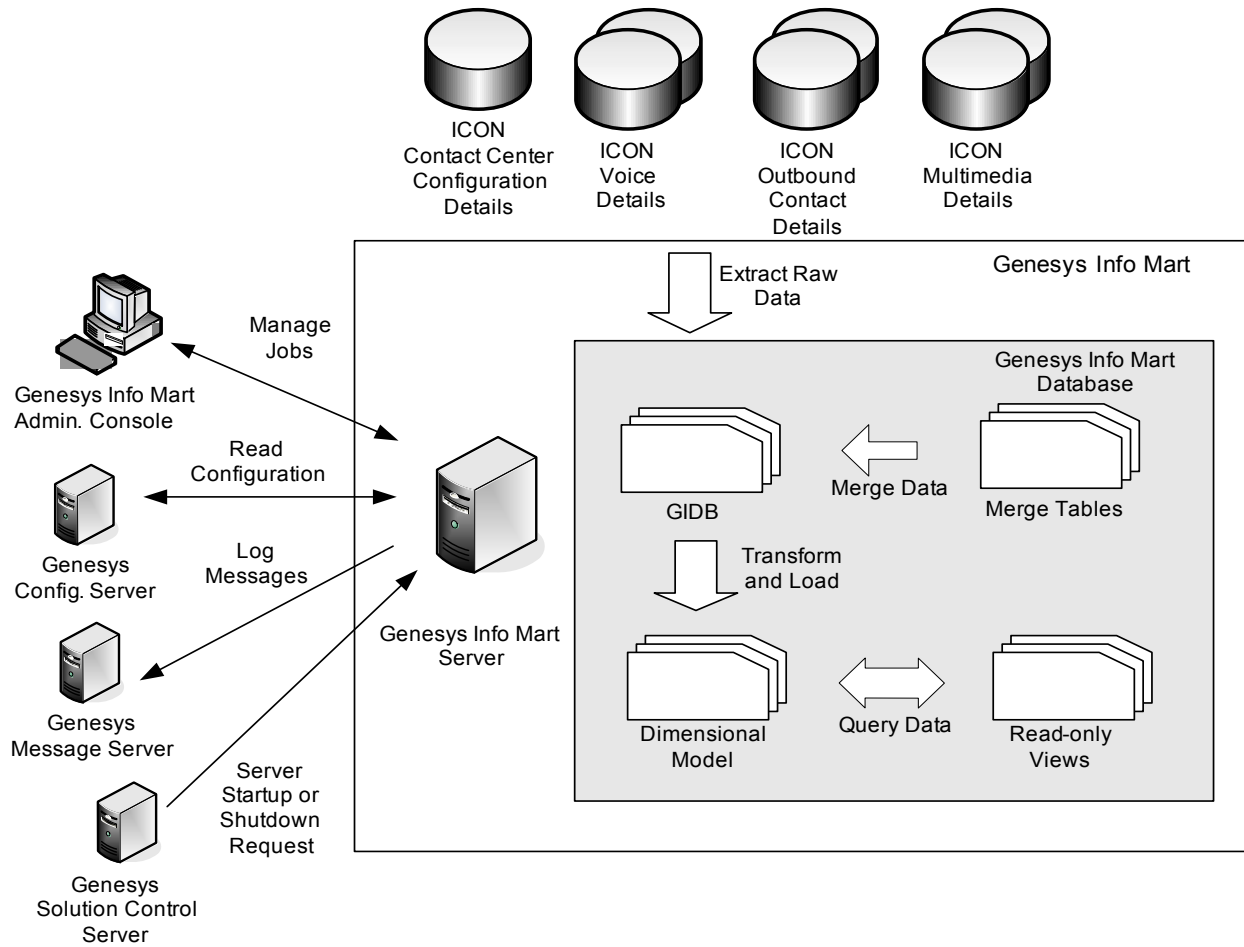
Figure 1: Genesys Info Mart Architecture and Data Flow Diagram

Components and Functions

The Genesys Info Mart operational environment consists of the following components:

- Genesys Info Mart Server—A server component that is configured with the Genesys Info Mart application

- Genesys Info Mart Administration Console—A graphical user interface (GUI) for managing jobs

- Info Mart Database—A database that is organized into a multi-level data model and contains data processed by Genesys Info Mart Server.

Genesys Info Mart Server

The Genesys Info Mart Server, a Java-based component, is the main executable process in Genesys Info Mart. Its main function is to run various functional jobs, including but not limited to jobs that extract and transform data based on the schedule that you configure in the Genesys Info Mart application. The Genesys Info Mart Server also processes requests from applications outside of the Genesys Info Mart Server, such as Genesys Info Mart Administration Console, to launch a specific job.

The Genesys Info Mart Server interfaces with:

- Solution Control Server (through Local Control Agent [LCA]), to control when the Genesys Info Mart Server starts and stops.

- Configuration Server, to read Genesys Info Mart application configuration options, as well as other configuration objects and options that affect Genesys Info Mart functionality.

- The Genesys Info Mart Administration Console, to start and stop jobs, and to provide the status of jobs.

- Message Server, to log messages to the Centralized Log Database.

Jobs in Genesys Info Mart

Jobs, which run under the Genesys Info Mart Server, perform extract, transform, and load (ETL) processes and other functions. The jobs do the following:

Initialize
- Set up your database for Info Mart use.

Extract
- Extract contact center configuration history details from the Interaction Database (IDB) component of Interaction Concentrator.

- Extract voice interaction, user data (including call-based attached data and UserEvent-based key-value pair (KVP) data), virtual queue, and agent activity details from one or more IDBs.

- Extract Outbound Contact solution details from one or more IDBs.

- Extract eServices/Multimedia solution interaction, attached data, virtual queue, and agent activity details from one or more IDBs.

- In an HA configuration, evaluate all redundant IDBs by comparing ICON-provided session information for each redundant set of IDBs (that store configuration, voice, OCS, or multimedia details), prior to extracting the data in a particular extract cycle.

Merge
- As part of the extract process, run a merge procedure to establish associations between related voice interactions that were extracted from the same or multiple IDBs—for example, to resolve intersite call linkages between related calls in a multi-site environment.

Transform • Transform the extracted data so that it becomes suitable for end-user reports and is available in a set of tables referred to as *dimensional model.* (Transform and load are performed by a single job.)

Maintain • Purge old data from the Info Mart database.

• Maintain calendar data by prepopulating calendar dimensions for use in reports.

• Create partitions on partitioned databases.

Migrate • Run the necessary scripts to update your database schema and perform any other steps necessary to move from an earlier 8.x release to the current one.

For more information about scheduling jobs using the Genesys Info Mart Server, see "Scheduling Jobs with Genesys Info Mart Server" on .

Genesys Info Mart Administration Console

The Genesys Info Mart Administration Console provides a graphical user interface that you can use to monitor job status and, when necessary, start or stop jobs outside of the normal schedule.

The Genesys Info Mart Administration Console provides the following functionality:

• Displays the current job execution status.

• Displays a history of job execution, including start time, stop time, duration, and final status.

• Filters the job execution history that is displayed based on time and/or status.

• Executes a single job on an ad hoc basis, either immediately or at a specified future time and date.

• Cancels a scheduled job.

• Shuts down a running job.

The Genesys Info Mart Administration Console interfaces with the following components:

• The Genesys Info Mart Server; to start, schedule, and stop jobs.

• The RDBMS, to query job status, job history, and job schedules from the Info Mart database. The Genesys Info Mart Administration Console connects to the Info Mart database by using a DB Server and a dedicated Database Access Point (DAP).

For more information, see "Managing Jobs with Genesys Info Mart Administration Console" on .

Info Mart Database

The Info Mart database contains all the data populated by the jobs that ETL data, including the low-level interaction data consolidated from any number of Interaction Databases (IDBs), as well as the data processed into a dimensional model that is suitable for end-user reports.

Genesys Info Mart Database Model

The Genesys Info Mart data resides in a database schema that includes:

- Control tables—Store information that controls ETL execution, such as the status of running jobs, job schedules, execution history, audit logs, and similar bookkeeping information.

- Merge tables—Used for merge of voice interactions. After the merge, voice interactions are moved into Global Interaction Database (GIDB) tables.

- GIDB tables—Store the low-level interaction data consolidated from any number of IDBs.

  There are separate sets of interaction-related GIDB tables for voice and multimedia interactions, to accommodate different requirements for transformation logic, indexes, data retention periods, and so on.

  The low-level reporting data in GIDB supports the possibility of custom reporting or drilling down from the dimensional model.

- Staging tables—Store information that is required in the transformation process. Staging tables contain data that has not yet been transformed or cannot be transformed due to incompleteness or inconsistency of source data and other auxiliary data necessary for the ETL process.

- Temporary tables—Store data which is used only within one job.

- Fact and Dimension tables—Contain transformed data that downstream reporting applications can query and combine in meaningful reports.

In deployments that use Genesys Interactive Insights (GI2) or Reporting and Analytics Aggregates (RAA) deployed as a separate package, the Info Mart database also includes Aggregate tables and views used by GI2. For more information, see the *Reporting and Analytics Aggregates 8.0 Reference Manual*.

In addition to the above tables that are used by Genesys Info Mart jobs, the following views are intended to simplify data retrieval for reports:

- Predefined views—Read-only views of certain configuration dimensions and facts contained in GIDB tables. These views are created in the Info Mart database schema.

- Genesys Info Mart Tenant Views—Tenant-specific read-only views of the dimensions and facts in the dimensional model and predefined views of the Info Mart schema. These views are created in a Genesys Info Mart Tenant

Views schema. Genesys Info Mart supplies a separate database schema for each tenant (including the Environment tenant), so that each tenant can access only its own data. In addition, these views shield business users from evolutionary changes to the underlying database schema and prevent users from accidentally changing the contents of the underlying database. Customers should use these views to query Genesys Info Mart data.

Genesys Info Mart provides SQL scripts that you execute to create database views that your reporting application can query. These views are created in the Genesys Info Mart Views and Genesys Info Mart Tenant Views schemas.

For detailed descriptions of the components, refer to the *Genesys Info Mart 8.0 Deployment Guide.* For a description of the database schema, refer to the *Genesys Info Mart 8.0 Reference Manual* for your RDBMS type.

Chapter

2 Understanding Genesys Info Mart Jobs

This chapter describes the jobs that ship with Genesys Info Mart. When executed, these jobs:

- Extract data from your source databases.
- Cleanse and transform the data.
- Load the data into the Info Mart Fact and Dimension tables.
- Optionally, calculate and load aggregated data into the Aggregate tables.
- Purge old data from the Info Mart database.
- Maintain calendar Dimension tables.
- Add and delete partitions for partitioned tables.
- Migrate your existing version 8.x Info Mart database so it is ready to use by the current 8.x release of Genesys Info Mart.

This chapter contains the following sections:

Job Summary Table

Table 1 on page 18 lists the jobs provided with Genesys Info Mart. Review the table to familiarize yourself with the job functions. For information about how

the jobs transform data, job interdependencies, and a sample schedule, see Chapter 3, "Working with Jobs," on page 45.

Table 1: Genesys Info Mart Job Summary

Name	Function	Frequency	Notes
Job_InitializeGIM	Populates many of the Dimension tables in the Info Mart database with fixed information. Adds partitions to partitioned tables for partitioned databases.	Once	This job automatically executes once during the first run of Genesys Info Mart after the initial deployment.
Job_ExtractICON	Extracts new and changed records from one or more Interaction Databases (IDBs), and stores those records in Global Interaction Database (GIDB) tables. The job first stores voice call records in the Merge tables of the Genesys Info Mart schema, and then merges the interaction data to reconcile links between interactions that contain related calls. After the Merge operation has completed, the job stores the merged interaction data in GIDB. Prior to extracting data in high availability (HA) deployments, the job evaluates which IDB to use for data extraction in a particular extract cycle.	Intraday, as scheduled	
Job_TransformGIM	Transforms and loads previously extracted data into the Fact and Dimension tables of the Genesys Info Mart database.	Intraday, depending on `Job_ExtractICON`	
Job_AggregateGIM	Aggregates or re-aggregates the facts based on data that was added or changed during the last transform job. Stores the data in historical Aggregate tables.	Continuous	This job is available in deployments with either Genesys Interactive Insights (GI2) reports or Reporting and Analytics Aggregates (RAA) package.

Table 1: Genesys Info Mart Job Summary (Continued)

Name	Function	Frequency	Notes
Job_MaintainGIM	Purges old data from the Info Mart database, using retention periods that are configured in the Genesys Info Mart application.	Daily	
	Maintains calendar dimensions by populating ahead the DATE_TIME and any custom calendar tables.	As configured (default is 183 days)	
	On partitioned databases, adds and removes partitions for partitioned tables.	As needed	
Job_MigrateGIM	Runs all scripts and makes any other updates necessary to prepare your database for the new release of Genesys Info Mart.	Once, as a part of the process of moving from an earlier 8.x release to the current one.	Must be run from the Genesys Info Mart Administration Console.

Job_InitializeGIM

This job performs the following functions:

- Populates the following dimensions with fixed information:

  - ATTEMPT_DISPOSITION
  - CALL_RESULT
  - CAMPAIGN_GROUP_STATE
  - CONTACT_INFO_TYPE
  - DATE_TIME
  - DIALING_MODE
  - INTERACTION_RESOURCE_STATE

  - INTERACTION_TYPE
  - MEDIA_TYPE
  - RECORD_STATUS
  - RECORD_TYPE
  - RESOURCE_STATE
  - TECHNICAL_DESCRIPTOR

Genesys Info Mart Server automatically launches Job_InitializeGIM to initialize the Info Mart database during the first run of Genesys Info Mart after the initial deployment.

Job_ExtractICON

`Job_ExtractICON` extracts data from one or more IDBs in discrete chunks and stores it either directly in the Global Interaction Database (GIDB) tables or, for voice interaction data, initially in the Merge tables within GIDB. As part of the extract process for Voice details, `Job_ExtractICON` merges related data in the Merge tables, and then moves the data to the GIDB tables.

In particular, `Job_ExtractICON`:

- Populates the `START_DATE_TIME_KEY` field in the GIDB and Merge tables.

- Merges call data in the Merge tables.

- Creates audit log records in the `CTL_AUDIT_LOG` table for each chunk.

In high availability (HA) deployments, the extract job also analyzes the Interaction Concentrator (ICON)–provided session information in the redundant IDBs that store the same type of data (configuration, voice, OCS, or multimedia details). This analysis occurs prior to data extraction in a particular extract cycle.

After the successful completion of the extract job, the Genesys Info Mart Server launches the job that transforms all the extracted ICON data.

Note: By default, all time dimension data is calculated in Coordinated Universal Time (UTC) format.

Extract Processing

There are some differences in the extract-processing details for different types of data and data flows.

This subsection contains information about the following topics:

Extracting Configuration Details

ICON Configuration details include data about configuration objects and configuration object relationships.

> **Note:** The `extract-data-cfg-facts-chunk-size` option enables you to configure a large extract window to minimize the possibility that any configuration data might be skipped.

The configuration data extracted includes the following:

- Configuration objects, which are extracted from IDB tables with the prefix `GC_` (for example, `GC_AGENT` and `GC_PLACE`)
- Configuration object relationships, which are extracted from IDB tables with the prefix `GCX_` (for example, `GCX_AGENT_PLACE`)

The paragraphs below explain the different approaches to data extraction for these two types of Configuration details.

Extracting Configuration Object Data

There is a single record for each configuration object in IDB. After the record is initially inserted, it might subsequently be updated. Each insertion or update is an independent snapshot of the configuration object.

The ETL extracts Configuration object data as soon as it becomes available in IDB, and merges the new data with existing records in the equivalent Global Interaction Database (GIDB) tables. If a record for a configuration object already exists, the record is updated if the newly arriving data has a later `LASTCHANGE` timestamp.

The ETL might create a particular GIDB table record based on data extracted from one IDB, and then might subsequently update the record based on data extracted from another IDB. Therefore, even if your deployment provides HA of Configuration details, it is important to ensure that available IDBs maintain synchronization with the Configuration Database.

The ETL extracts all available configuration object data in one cycle, and there is no expiration threshold for the extract window.

Extracting Configuration Object Relationship Data

Each configuration object relationship record covers the interval of an historical association between two or more configuration objects, which are uniquely identified by database identifiers that are assigned by Configuration Server (DBIDs).

The ETL extracts object relationship records from an available IDB based on comparing the extract high-water mark against the IDB timestamp of when the record was written to IDB. New relationship fact records are inserted into GIDB, and updates are simply merged into existing historical records. The

ETL might create a particular GIDB table record based on data extracted from one IDB, and then might subsequently update the record based on data extracted from another IDB.

For updates, the ETL updates only information about termination of the association. The ETL does not compare and update relationship attributes (for example, the skill level in a GCX_SKILL_LEVEL record for a particular association between an agent and a skill).

Extracting Voice Interaction Data

ICON Voice details include data about:

- Voice interactions, including virtual queue usage and attached data.
- Agent states and agent state reasons, including DND mode details.
- Agent login sessions.

Agent-related data for voice activity is extracted in the same way as for multimedia activity. For more information, see "Extracting Agent State, Agent State Reason, and DND Usage Data" on page 24 and "Extracting Agent Login Session Data" on page 25.

For voice interaction data, the ETL processes only *completed* voice interactions. The ETL identifies eligible data by comparing the extract high-water mark against the termination timestamp in the record. For reliable call data, the call must be visible to ICON for the entire call duration, from the time of call creation until call completion.

The extract job initially extracts eligible call-related data to Merge tables, and it does not move the data to GIDB until it has run the merge procedure to establish the associations between related voice interactions.

Data from the G_IR, G_CALL, and G_IS_LINK tables is always extracted from the same IDB.

In an HA deployment, when there are no disruptions, the ETL continues to extract each table independently from a reliable IDB that it used in the previous extract cycle. When disruptions occur, the ETL will switch to extracting from another IDB, if possible.

Extracting Multimedia Interaction Data

ICON Multimedia details include data about:

- Multimedia interactions, including virtual queue usage and attached data.
- Agent states and agent state reasons.
- Agent login sessions.

Agent-related data for multimedia activity is extracted in the same way as for voice activity. For more information, see "Extracting Agent State, Agent State

Reason, and DND Usage Data" on page 24 and "Extracting Agent Login Session Data" on page 25.

Unlike voice interactions, multimedia interactions might be very long-lived. The ETL does not wait for multimedia interactions to end before it starts to extract them. The ETL extracts interaction data from an available IDB, based on comparing the extract high-water mark against the record-creation timestamp.

Special Considerations When Restarting Multimedia ICON

Genesys Info Mart requires that the ICON `calls-in-the-past` and `om-force-adata` options be set to `true`. This means that, if multimedia interactions begin while ICON is down or has no connection to Interaction Server, ICON reconstructs operational data and stores a user data snapshot for multimedia interactions that are already in progress beginning with the next party that ICON sees being added to the interaction. The ETL extracts the reconstructed data in the usual way, based on a comparison of the extract high-water mark against the record-creation timestamp. However, be aware that, in these situations, information about previous parties and first values of user data keys might be missing or inaccurate.

The consequences may include:

- Multimedia interactions that end while the Multimedia ICON is down but Interaction Server is running are forever reported as active. Since ICON did not see these interactions end, Info Mart cannot report them as ended.

- ICON records interaction-related data (`IR, CALL, PARTY`) for a never-before-seen multimedia interaction beginning when ICON sees a party being added to the interaction.

- End times are only shown for a party if the party was created and ended within the same ICON session.

- If ICON sees an interaction end, but the interaction is neither currently known, nor known from a previous ICON session, then it does not record anything about this interaction.

Avoiding Data-Quality Issues   As noted above, some data quality issues regarding multimedia interactions can occur when the interaction data was recorded during different ICON sessions, and the high-availability architecture is *not* in use. Some information can be lost between the sessions. However, sometimes it is necessary to stop an ICON to apply upgrades, or so it can become aware of configuration changes when it restarts. It is possible to avoid nearly all of these data quality issues across a planned ICON outage by using the following procedure:

1. Stop the Interaction Server using the Solution Control Interface (SCI) prior to stopping the Multimedia ICON.

2. Stop the Multimedia ICON.

3. Restart the Multimedia ICON.

4. Restart Interaction Server using the Solution Control Interface (SCI).

By stopping Interaction Server first, interactions are moved to a 'home' state, returning to Interaction Queues or Workbins, and ICON will be aware of this.

It is best to follow this procedure during a period of agent inactivity (for example at the end of one shift, and before the next shift starts) to minimize disruption to the agents. When Interaction Server is stopped, all agents are logged out of Interaction Server, and any interactions that were actively on their desktops are returned to Interaction Queues or Workbins, so it is best to follow the procedure when agents are not in the middle of handling interactions.

Following this procedure eliminates nearly all data quality issues that can occur when a multimedia interaction spans more than one ICON session. However it is possible that some Virtual Queue activity that occurs during the time frame of the procedure will be lost:

- Interactions that were in Virtual Queues when Interaction Server was stopped are returned to Interaction Queues. The events sent by URS indicating the interactions were cleared from Virtual Queues may not reach Interaction Server before it stops its connection to URS. So the data may not show that the interaction left the Virtual Queue.

To avoid data-quality issues if a scheduled restart of ICON cannot be performed without affecting multimedia interaction-handling, or if ICON stops unexpectedly, Genesys recommends you use a high-availability architecture for multimedia.

Note: When Interaction Server is restarted, some interactions may be placed in Virtual Queues before ICON has successfully reestablished its connection to Interaction Server. However, ICON will have established its connection to Interaction Server by the time agents begin logging in, so no agent activity is lost, not even from agents who receive interactions from those Virtual Queues.

Extracting Agent State, Agent State Reason, and DND Usage Data

Agent state–related data is extracted in the same way for agents handling voice and multimedia interactions.

Note: For a detailed discussion of Agent States, Agent State Reasons, and DND, see the "Populating Agent Activity' section of the "Populating Genesys Info Mart Data" chapter in the *Genesys Info Mart 8.0 User's Guide*.

Agent state–related records are inserted once and never updated. These records do not contain artifacts that are created by ICON. Therefore, in an HA

deployment, they do not need to be synchronized between HA ICON instances.

The ETL extracts interaction data from any available IDB, based on comparing the extract high-water mark against one of the following:

- For agent state and DND history records, the timestamp when ICON detected a change in the agent's state.

- For agent state reason code records, the timestamp when ICON detected creation or end of the hardware or software reason code on an agent's state.

Extracting Agent Login Session Data

The configuration objects that are the subjects of agent login session records are uniquely identified by DBIDs. Agent login session data is processed similarly to configuration object relationship data (see).

Extracting Outbound Contact Details

The configuration objects that are the subjects of Outbound Contact records are uniquely identified by identifiers that are reported by Outbound Contact Server (OCS).

In some tables, records are inserted once and never updated (for example, GO_CAMPAIGNHISTORY table records). In other tables, records might be updated (for example, GO_CAMPAIGN table records); in these cases, the updated data is merged with the existing record in the GIDB table.

The ETL extracts Outbound Contact data from an available IDB, based on comparing the extract high-water mark against the timestamp of the OCS event that triggered the creation or update of the record.

Outbound Contact records do not contain artifacts that are created by ICON. Therefore, in an HA deployment, they do not need to be synchronized between HA ICON instances.

Extract Algorithm

Job_ExtractICON extracts IDB data in chunks. The size of the chunk is a configurable time interval. This enables Genesys Info Mart to manage the volume of data committed in one transaction, because each data chunk is extracted and committed in a separate transaction.

Extract Windows

Job_ExtractICON uses different algorithms to extract data from different types of data sources. There are separate extract windows for the data from Configuration Server, T-Server, Interaction Server, and Outbound Contact

Server, respectively. `Job_ExtractICON` extracts data concurrently, using a configurable number of threads.

`Job_ExtractICON` uses the timestamp of the server that produced the data as a high-water mark. For example, if the previous run of `Job_ExtractICON` extracted data with timestamps up to 3:00 PM, the next run will extract data with timestamps starting from 3:00 PM.

In order to be aligned in time, related data is extracted inside the same extract window. For all types of data except configuration object data, the timespan of the extract window depends on the configured chunk size, in combination with a configured stuck-data timeout. The configured stuck-data timeout (or threshold) determines the size of the window. If there is available data to extract, the extract window advances forward by the amount of the chunk size. Separate configuration options enable you to configure a larger chunk size (and, therefore, a larger extract window) for configuration relationship data than for Voice details, Multimedia details, or Outbound Contact details.

Extract Window Example

For example, say the configured chunk size is 15 minutes; the configured stuck-data threshold is 8 hours, 1 minute; in the previous extract cycle, `Job_ExtractICON` extracted data with timestamps up to 2:45 PM; and data is available beyond 3:00 PM. Then the extract window for the next extract cycle is 6:59 AM to 3:00 PM.

Single Extract Cycle for Configuration Object Data

For configuration object data, all available data is extracted in one extract cycle, so that any configuration-related information that may be required is available for transformation of other data.

Configuration Options

For information about the configuration options that control the size of the extract windows and the operation of `Job_ExtractICON`, see "Controlling the Volume of Data Extraction" on .

About Extract Roles

The extract algorithm used by `Job_ExtractICON` depends on the extract role that you configured in the DAP. These roles include:

* `ICON_CFG`

* `ICON_CORE`

* `ICON_OCS`

* `ICON_MM`

ICON_CFG

`ICON_CFG`—`Job_ExtractICON`:

* Extracts all new and changed data from IDB tables that store the contact center configuration history, and stores the data in the GIDB tables of the Genesys Info Mart database. Extracts all available configuration data in one extract cycle in order for the transformation of other data to proceed.

- Extracts object relationship records from IDB tables. New relationship fact records are inserted into GIDB, and updates are simply merged into existing historical records.

ICON_CORE

ICON_CORE—Job_ExtractICON extracts:

- Completed virtual queue details.
- Completed voice interaction details, such as calls and user data (including call-based attached data and UserEvent-based key-value pair (KVP) data).
- Both active and completed voice agent login session details.
- Both active and completed voice agent states.
- Completed voice agent state reason codes.
- Both active and completed voice do-not-disturb (DND) modes.

The job stores the information in the GIDB tables of the Genesys Info Mart database. The job also merges voice data in the Merge tables before transformation.

The Merge Operation   When Job_ExtractICON extracts voice details, it initially stores interaction data in the Merge tables of the Info Mart database. Job_ExtractICON then performs a Merge operation against the Merge tables to resolve linkages (IS-Links) between related calls. If Genesys Info Mart has information about both sides of an IS-Link, the merge operation combines both call records into a single, end-to-end interaction, and then Job_ExtractICON moves the successfully merged voice interactions to the GIDB tables.

Partially Merged Interactions   If Genesys Info Mart does not have information about both sides of an IS-Link, the merge operation is able to perform only a partial merge, before Job_ExtractICON moves the interaction data to GIDB.

The usual reasons for missing IS-Link information are:

- One of the sites in a multi-site call is not monitored by ICON—Data from any unmonitored site(s) will never arrive. To streamline processing, identify unmonitored sites when you prepare the Info Mart database, so that the merge operation does not wait for the stuck link timeout to expire before the procedure finalizes the partial merge of interactions with links to the unmonitored site. For more information about configuring unmonitored (remote) locations, see the procedure about configuring the Info Mart database for merge in the *Genesys Info Mart 8.0 Deployment Guide*.
- Link information from monitored sites is delayed or missing (for example, because of intermittent connectivity issues, or because of extended outage of an ICON application during an extract cycle)—The merge operation waits for the missing data until the stuck link timeout expires. If the data arrives before the timeout expires (delayed data), the interaction is

processed normally. If the data does not arrive before the timeout expires (missing data), the merge operation finalizes the partial merge, and `Job_ExtractICON` moves the partially merged interaction data to GIDB.

The stuck link timeout is controlled by the, `max-call-duration` configuration option. This option also affects a number of other functions, such as establishing a limit for waiting for ACW and for looking up multimedia party activity before transforming data.

As a stuck threshold for merge operation, `max-call-duration` specifies the amount of time Info Mart waits for IS-Link information from another site, before it considers the IS-Link to be stuck.

For more information about the `max-call-duration` parameter, see the section about partially-merged calls in the chapter on Data Processing in the *Genesys Info Mart 8.0 Deployment Guide*.

When the extract job moves partially-merged interactions into GIDB, the interaction records have unresolved linkages, which are referred to as *dangling links*. For information about how `Job_TransformGIM` processes partially merged interactions, see "Handling Partially Merged Interactions" on page 35.

Notes: Extracting voice details from IDB is optional, provided that you are extracting multimedia (`ICON_MM`) details from an IDB. Genesys Info Mart requires that there is at least one DAP that is configured for the role of either `ICON_CORE` or `ICON_MM`.

If you do not wish to extract voice details, do not configure a DAP for the `ICON_CORE` role.

If you plan to extract both voice and multimedia details, you must configure different ICON applications to store ICON voice and ICON multimedia details into separate IDBs. This means that you cannot configure both the `ICON_CORE` and `ICON_MM` roles for the same DAP. You can configure multiple DAPs for the `ICON_CORE` role, to extract voice details from more than one IDB. For more information, see the section in the *Genesys Info Mart 8.0 Deployment Guide* about configuring Genesys Info Mart and data sources to extract both voice and multimedia details.

ICON_OCS

ICON_OCS—Job_ExtractICON extracts all new and changed data from IDB tables that store OCS data, and stores the data in the GIDB.

Note: Extracting Outbound Contact details from IDB is optional, provided that you are extracting voice (ICON_CORE) details from an IDB. Genesys Info Mart requires that there is at least one DAP that is configured for the role of either ICON_CORE or ICON_MM.

- If you do not wish to extract Outbound Contact details, do not configure a DAP for the ICON_OCS role.
- If you do configure a DAP for ICON_OCS, the Genesys Info Mart Server will run Job_ExtractICON for the DAP with this role.

If you plan to extract both voice and Outbound Contact details, you must configure different ICON applications to store ICON Voice and ICON Outbound Contact details into separate IDBs. This means that you cannot configure both the ICON_CORE and ICON_OCS roles for the same DAP. You can configure multiple DAPs for the ICON_OCS role, to extract Outbound Contact details from more than one IDB. For more information, see the information in the *Genesys Info Mart 8.0 Deployment Guide* about supported topologies and about configuring Genesys Info Mart to extract Voice and Outbound Contact details.

ICON_MM

ICON_MM—Job_ExtractICON extracts:

- All new and changed data from IDB tables that store multimedia interactions. Both active and completed multimedia interactions are extracted along with user data (including interaction-based attached data and eServices/Multimedia-specific attributes). These interactions do not need to be merged.
- Both active and completed virtual queue details.
- Both active and completed multimedia agent login session details.
- Both active and completed multimedia agent states.
- Completed multimedia agent state reason codes.
- Both active and completed multimedia do-not-disturb (DND) modes.

The job stores the extracted multimedia data in the GIDB tables of Genesys Info Mart database.

Notes: Extracting Multimedia details from IDB is optional, provided that you are extracting voice (`ICON_CORE`) details from an IDB. Genesys Info Mart requires that there is at least one DAP that is configured for the role of either `ICON_CORE` or `ICON_MM`.

- If you do not wish to extract Multimedia details, do not configure a DAP for the `ICON_MM` role.
- If you do configure a DAP for `ICON_MM`, the Genesys Info Mart Server will run `Job_ExtractICON` for the DAP with this role.

If you plan to extract both voice and multimedia details, you must configure different ICON applications to store ICON Voice and ICON Multimedia details into separate IDBs. This means that you cannot configure both the `ICON_CORE` and `ICON_MM` roles for the same DAP. You can configure multiple DAPs for the `ICON_MM` role, to extract Multimedia details from more than one IDB. For more information, see the information in the *Genesys Info Mart 8.0 Deployment Guide* about supported topologies and about configuring Genesys Info Mart to extract Voice and Multimedia details.

Determining Data Availability

At the start of each extract job, Genesys Info Mart:

1. Looks into the connections lists of all the ICON `Applications` to which Genesys Info Mart has a configured connection, and compiles a list of all the data sources and IDBs that have been configured for the deployment. (If the same data source is present in the connections list of more than one ICON application, HA mode is automatically in effect.)

2. For each data source in ICON connections, checks if the data source application is enabled (the `State Enabled` check box on the `General` tab is selected). Genesys Info Mart uses the `State Enabled` option as a flag to identify whether it should consider the data source as one that is intended to be available.

3. Checks configuration of the ICON DAPs, and compiles a list of all the IDBs that have been configured for the deployment.

4. Opens connections to all available IDBs, reads session information from the `G_DSS_*_PROVIDER` tables in the IDBs, and compiles a list of all available data sources.

5. Compares the list of configured data sources against the list of available data sources. Any discrepancy between the lists enables Genesys Info Mart to infer that a failure has occurred.

During extract processing, Genesys Info Mart uses information in the `G_DSS_*_PROVIDER` tables in IDB to monitor activity on each data source session.

If Genesys Info Mart identifies that data from a configured data source (or HA pair) is not available in any IDB during an extract cycle, Genesys Info Mart logs an error and does not proceed with the extract job.

Controlling Extract Job Functioning

Genesys Info Mart uses a number of configurable settings that control `Job_ExtractICON`. This subsection describes application and database settings that affect the operation of `Job_ExtractICON`, as well as operational suggestions for improving performance.

Controlling the Volume of Data Extraction

Except in the case of configuration object data, you can configure Genesys Info Mart to control how much data should be extracted from an IDB at any one time. For example, you can configure the number of concurrent extractions and the chunk size.

Configuration Options to Control Data Extraction

The following configuration options in section [`gim-etl`] control `Job_ExtractICON`. For more information about these options, see the chapter about configuration options in the *Genesys Info Mart 8.0 Deployment Guide.*

- `extract-data-cfg-facts-chunk-size`— Specifies the size of the time interval, in seconds, for configuration relationship data that is committed in one transaction. This option also affects `Job_TransformGIM`—see "Transaction Size" on .

- `extract-data-chunk-size`— Specifies the size of the time interval, in seconds, for data committed in one transaction. This option also affects `Job_TransformGIM`—see "Transaction Size" on .

- `extract-data-stuck-threshold`— Specifies the time, in seconds, that `Job_ExtractICON` waits for stuck data.

- `extract-data-thread-pool-size`—Specifies the maximum number of threads that can be used to extract data concurrently.

Streamlining Merge Operation Functioning

The following database-related settings affect the merge operation:

- The `max-call-duration` parameter specifies the time, in seconds, that the merge stored operation waits for stuck IS-Links before finalizing a partial merge. For a complete description of this option, see the Configuration Options chapter in the *Genesys Info Mart 8.0 Deployment Guide*.

- The `merge-chunk-size` option (in the `gim-etl` section) enables you to manage situations in which the Merge area could grow to such a size that the transform job has too much data to transform, thus impacting processing speed. The option makes the Merge process create multiple chunks of manageable size.

  For a complete description of this option and how to determine the right chunk size for your environment, see the *Genesys Info Mart 8.0 Deployment Guide*.

- The `REMOTELOCATION` column in the `GSYS_DNPREMOTELOCATION` table stores the names of unmonitored switches, so that the lack of information from these switches will not delay merge processing. The `GSYS_DNPREMOTELOCATION` table is one of the Merge tables in the Info Mart database.

Data Extraction in HA Deployments

`Job_ExtractICON` ensures Genesys Info Mart support for HA of all types of reporting data.

The extract algorithm for redundant data includes selection of the IDB from which to extract the data in a given extract cycle, if more than one IDB stores the same type of data for the corresponding timespan.

For detailed descriptions of the IDB selection criteria and switchover algorithm, refer to the "High Availability" chapter of the *Genesys Info Mart 8.0 Deployment Guide.*

Job_TransformGIM

This job transforms data extracted from all IDBs. The Genesys Info Mart Server launches this job during each ETL cycle after it has extracted data from all IDBs. The job `Job_TransformGIM` transforms GIDB data and then loads it into the Info Mart database as the last step of the transformation process. For more information about the scope of the data transformation, see Chapter 3, "Working with Jobs," on .

Transaction Size

`Job_TransformGIM` transforms extracted GIDB data, chunk by chunk. For purposes of transformation, a chunk of data is considered to be all the data in a primary table (for example, `GIDB_G_IR`) plus all the related data in the secondary table (for example, `GIDB_G_PARTY_HISTORY`) marked by the same audit keys.

Because the audit keys are assigned to GIDB data by the upstream extract job (`Job_ExtractICON`), this means that the transaction size of `Job_ExtractICON` determines the transaction size of `Job_TransformGIM`. Therefore, when you set the configuration options to control the volume of data extracted (see

"Controlling the Volume of Data Extraction" on), also consider transaction size for transform transactions.

Horizontal Transform

The transform logic implements a dependency between the data in the primary and secondary tables. Data from the primary table for time *T* can be transformed only when the data in secondary table(s) for same time *T* is available. Before transforming a chunk of data, `Job_TransformGIM` checks that all related data is extracted at least up to the maximum timestamp in the chunk. If it is not all extracted, the transformation is deferred. Under normal conditions, the next extraction brings in new data, enabling the transformation to proceed. Occasionally, the extract cycle will never pick up the missing data (for example, an ICON was stopped and uninstalled). In this case, the data will be missing for longer than the wait time specified by the `extract-data-stuck-threshold` configuration option. When `Job_ExtractICON` moves the high-water mark forward, this enables `Job_TransformGIM` to proceed with the transformation of the available data.

Each time `Job_TransformGIM` checks for dependent data but does not find it, it generates a log message.

Delayed Data

If data from one data source is delayed, transformation of data from all data sources of the same type is delayed while the transform job waits for the missing data.

To prevent excessive or unnecessary delays, a configuration option in the `[gim-etl]` section, `delayed-data-threshold`, enables you to specify a timeout after which Genesys Info Mart logs a detailed message (55-20110) if expected data does not arrive. Genesys recommends that you set an alarm on this message, so that you can investigate the reasons for data delays in a timely manner and take appropriate action, such as:

* Temporarily disabling an unavailable data source
* Interrupting the job or temporarily stopping the ETL schedule, so that the next run of the extract job does not advance the high-water mark

For more information about the `delayed-data-threshold` option, see the chapter about data processing in the *Genesys Info Mart 8.0 Deployment Guide*. For more information about actions you can take in response to log message 55-20110, see "Recovering from Data Source Unavailability" on .

Delayed and Early-Arriving Virtual Queue Data

Because the T-Server or Interaction Server that handles virtual queue data might not be the same T-Server or Interaction Server that handles the associated interaction, the virtual queue and interaction data streams are not necessarily synchronized.

Delayed and early-arriving virtual queue data requires additional processing. For example, delayed virtual queue data requires fact records in the dimensional model to be updated, while early-arriving virtual queue data must be stored and maintained in GIDB until the associated interaction data arrives. There is reasonably intensive processing to update Fact tables in the dimensional model and to maintain the Staging tables (to delete processed records). The additional processing has a performance impact.

Genesys Info Mart requires that you use separate virtual queues for voice and multimedia interactions in your contact center deployment. For example, if a `Virtual Queue` has been configured on a voice switch, the DNs that use that virtual queue must also be configured on a voice switch; conversely, if the DNs that use a particular virtual queue are configured on a multimedia switch, also configure the `Virtual Queue` on a multimedia switch.

Error Handling

The Genesys Info Mart server handles errors differently, depending on the type of data and the reasons for the error.

Handling Configuration Data Errors

During transform, Genesys Info Mart might encounter references to configuration objects about which it has no other data.

- When `Job_TransformGIM` transforms Configuration details, missing configuration objects are an error that indicates a problem in the source data. For example, if an agent was created before ICON started, then the agent was added to a group and subsequently deleted, there will be a record in the `GCX_GROUP_AGENT` table but no corresponding record in the `GC_AGENT` table.

  When `Job_TransformGIM` encounters this kind of missing configuration data, it discards the record and records the information in the `STG_IDB_FK_VIOLATION` table.

- When `Job_TransformGIM` transforms interaction data, it does not delay transform because of missing configuration objects.

  When `Job_TransformGIM` encounters unresolved references to configuration objects in the extracted data, it creates the missing configuration objects using available attributes (for example, the object's ID), and continues processing data.

  In many cases, you can resolve the error when you force ICON to perform a resynchronization of the configuration data between the Configuration Database and the IDB from which Genesys Info Mart extracts the configuration history. When the resynchronized configuration data is extracted, the incomplete Genesys Info Mart data about the resources created by `Job_TransformGIM` is updated.

Handling Partially Merged Interactions

As described in "Partially Merged Interactions" on page 27, there are two situations that result in partially merged voice interactions:

- One of the sites is registered as unmonitored—The transform job ignores dangling links that point to unmonitored sites, and transforms the interaction as if it were an inbound or outbound call to a remote site. For example, if the dangling link provides information about the target for an internal transfer, the interaction will be transformed as an inbound call rather than as a transfer.

- Information is missing—The transform job treats dangling links as a data inconsistency, which it handles as described in "Handling Data Inconsistencies" on page 35. The interaction-level error policy option that controls behavior when the transform job encounters this data inconsistency is `error-policy-islink-dangling`.

Partially merged calls have implications for data quality. For more information, see the section about data issues in a partially monitored environment in the *Genesys Info Mart 8.0 User's Guide*.

Handling Data Inconsistencies

`Job_TransformGIM` does not make a special effort to validate data consistency. However, if it encounters data inconsistencies during transform, it will report them, and then proceed in accordance with configurable error-handling policies.

Genesys Info Mart has a two-level approach to handling data inconsistencies during transform:

- Interaction-level policies
- Job-level policies (see page 37)

Interaction-Level Error Policies

At the interaction level, error policy options enable you to configure the behavior of the job when it encounters various types of data inconsistency during processing. Table 2 summarizes the interaction-level error policy options.

If the option that specifies error-handling policy at the interaction level is set to interrupt processing with an exception, the exception is handled as specified by the error-policy option for the job (see "Job-Level Error Policy" on page 37). By default, the interaction-level error policy options are set to generate an exception. This setting favors data consistency.

If the option that specifies error-handling policy at the interaction level is set to resume processing without generating an exception, `Job_TransformGIM` continues processing, and handles the error as described in Table 2.

Table 2: Interaction-Level Error Policy Options

Error Policy Option	Reason for the Error	Result If Job Resumes
error-policy-ipurpose-numberformat	The `IPurpose` attached data key-value-pair (KVP) is present but the value of `IPurpose` is not a number. This error usually arises because of incorrect configuration.	The transform job processes the data as if the `IPurpose` KVP was not attached.
error-policy-islink-dangling	A partially merged voice interaction is missing information for one side of the interaction. For the reasons that this error might arise, see "Partially Merged Interactions" on page 27.	The transform job processes the data as if the missing link information is 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	For a voice interaction, there are multiple (>1) source links with the same `LINKID`.	The transform job chooses one of the source records randomly, and ignores the other source records.
error-policy-islink-multiple-targets	For a voice interaction, there are multiple (>1) target links with the same `LINKID`.	The transform job chooses one of the target records randomly, and ignores the other target records.
error-policy-islink-multiple-vertices	For a voice interaction, there are more than two bidirectional links with the same `LINKID`. This data inconsistency occasionally occurs with older T-Servers.	The transform job chooses one of the bidirectional records randomly, and ignores the other bidirectional records.
error-policy-islink-source-party-missing	For a voice interaction, the source call for the link for a dial-out attempt does not have a remote dialed party.	The transform job builds the order for Interaction Resource Facts randomly as it processes the interaction.
error-policy-party-created-duplicated	`G_PARTY_HISTORY` contains multiple records with `ChangeType=1(party_created)` for some party.	The transform job treats the first record it reads as the party created record, and ignores the other party created records.
error-policy-party-created-missing	`G_PARTY_HISTORY` does not contain a record with `ChangeType=1(party_created)` for some party.	The transform job constructs a party created record based on assumptions from the first party history record it reads.
error-policy-party-parent-missing	The party refers to a parent, but party records with the referenced `PARTYID` do not exist.	The transform job ignores the missing data, and continues processing.

Job-Level Error Policy

An error policy option, `error-policy-irf-exception,` controls the behavior of `Job_TransformGIM` when the job encounters an exception. The job-level error might be generated as a result of a configured interaction-level error policy (see "Interaction-Level Error Policies" on).

The job-level error policy option can have the following values:

- `log_db_resume`—`Job_TransformGIM` discards the problematic interaction thread, writes a corresponding record into the `STG_TRANSFORM_DISCARDS` table, and resumes processing. (This is the default value.)

- `resume`—`Job_TransformGIM` discards the problematic interaction thread and resumes processing without writing a corresponding record into the database.

- `exception`—`Job_TransformGIM` fails.

Job_AggregateGIM

In an environment where either GI2 reports or RAA package are deployed, `Job_AggregateGIM` calculates or recalculates the historical Aggregate tables in the Info Mart database based on:

- Data that changed since the last load of the historical Fact tables.

- New settings for configuration options that control aggregation.

In release 8.0, `Job_AggregateGIM` runs continually within the time window you specify.

`Job_TransformGIM` notifies the aggregation engine when new or updated data is available for aggregation and in this way keeps `Job_AggregateGIM` continually aggregating new data or data that has been updated, during the active window you set for aggregation.

For more information about the aggregation engine, refer to the:

- Reporting and Analytics Aggregates documentation set, if you are only deploying RAA package that ships with Genesys Info Mart.

Scheduling

`Job_AggregateGIM` is an optional job. You can run `Job_AggregateGIM` from the Info Mart Administration Console in integrated mode if you plan to use GI2 reports.You can also aggregate directly from GI2 in autonomous aggregation mode. For more information about these aggregation modes, see the Reporting and Analytics Aggregates documentation set.

Aggregation Scheduling Configuration Options

Configuration options enable you to specify:

- Whether `Job_AggregateGIM` will run the aggregation engine within the Genesys Info Mart server process, under the control of the scheduler.

- The start times and duration of the daily intervals within which `Job_AggregateGIM` will run. Within these intervals, `Job_AggregateGIM` will run continuously.

For more information about the scheduling options, see the configuration options reference chapter in the *Genesys Info Mart 8.0 Deployment Guide*.

If you need to calculate or recalculate the aggregates for a certain time span, this capability is available outside of Genesys Info Mart Administration Console. Refer to the Reporting and Analytics Aggregates documentation set for instructions.

Job_MaintainGIM

This job performs the following tasks:

Purge Info Mart Database
- Purges the Info Mart database in accordance with configurable data retention policies. `Job_MaintainGIM` purges:
  - Completed and active fact data from GIDB.
  - Completed and active fact data from the dimensional model, including certain Staging tables.
  - Discarded operational data from discard tables.
  - Outdated information from the `AUDIT_LOG` and History tables.
  - For partitioned tables, purges old data that's eligible for deletion by dropping appropriate partitions.

For more information about the retention policy options, see the descriptions of the `days-to-keep-gim-facts` and `days-to-keep-gidb-facts` options in the configuration options reference chapter in the *Genesys Info Mart 8.0 Deployment Guide*.

For more information about data retention policies and about the purging algorithms that `Job_MaintainGIM` uses, see "Purging the Info Mart Database" on .

Note: `Job_MaintainGIM` does *not* purge old aggregate data or dimension data.

Maintain Calendars
- Populates the calendar table(s) for future reports. The calendar tables are the default `DATE_TIME` Dimension table and any custom calendar tables that have been created to support your reporting.

`Job_InitializeGIM` initially populates the calendar table(s) as far ahead as you specify (in the `date-time-max-days-ahead` option), so that calendar dimensions are available for your reports. `Job_MaintainGIM` continues to populate the calendar tables when the next batch of calendars is required (as specified by the `date-time-min-days-ahead` option).

For more information about creating custom calendars, see the *Genesys Info Mart 8.0 Deployment Guide*.

Maintain Partitions • If you are using partitions on a partitioned database, adds partitions as necessary to process incoming data.

Purging the Info Mart Database

Job_MaintainGIM uses different algorithms to purge different categories of data from various areas of the Info Mart schema. Separate configuration options enable you to configure different retention policies for the different categories of data.

Retention Periods

The Genesys Info Mart configuration specifies the number of days that completed and active facts are retained in the Info Mart database, as well as the number of days that discarded operational data is retained. Records for facts that started earlier than the applicable retention periods are purged from the Info Mart database. Similarly, operational data that was discarded by a job that ran earlier than the specified retention period is purged.

For the configuration options that determine the retention periods, the values that you choose must allow the Info Mart database to retain the data long enough for you to complete deployment-specific tasks (calculating aggregates, archiving data, or uploading data to a data warehouse). For example, if you want to retain data for one year, consider setting the value to 380 days. This allows for the extra day in a leap year, plus two extra weeks for aggregate calculation, archiving, or uploading data to a data warehouse.

Retention Periods for Facts

The subjects of fact records are entities such as interactions, agent login sessions, or Outbound Contact campaigns. The purging algorithms divide dimensional model and GIDB Fact tables into two categories:

Short-Living and • Tables that contain data about *short-living entities*—Entities whose
Long-Living Facts lifespans are well defined and of short durations, typically less than a day (for example, agent login sessions). Job_MaintainGIM purges records relating to short-living entities based on the start time of facts in the applicable tables. The days-to-keep-gim-facts and days-to-keep-gidb-facts options in the [gim-etl] section of the Genesys Info Mart configuration directly specify the retention period for short-living entities.

• Tables that contain data about *long-living entities*—Entities whose lifespans are indefinite and potentially of long durations (for example, e-mail interactions). Job_MaintainGIM purges long-living entities based on the start time of the earliest active fact in the interval defined by the

`days-to-keep-gim-facts` and `days-to-keep-gidb-facts` options (the retention period for short-living facts) and the `days-to-keep-active-facts` option (the retention period for active long-living facts).

- If there are active facts within the interval, the start time of the earliest active fact determines the purge eligibility threshold (in other words, the retention period) for both active and completed long-living facts. All completed long-living facts that started earlier than the purge threshold in a group of related tables are purged, but the purge eligibility threshold itself is based on the start timestamp in a root table.

- If there are no active facts within the interval, the `days-to-keep-gim-facts` and `days-to-keep-gidb-facts` options determine the retention period for completed long-living facts.

For a list of the tables in GIDB and in the dimensional model that contain records for short- and long-living entities (as defined for purging purposes), including the root tables for each group of long-living facts, see the appendix in the *Genesys Info Mart 8.0 Deployment Guide* about Info Mart tables purged by the maintenance job.

Figure 2 illustrates the purging strategy for long-living entities.
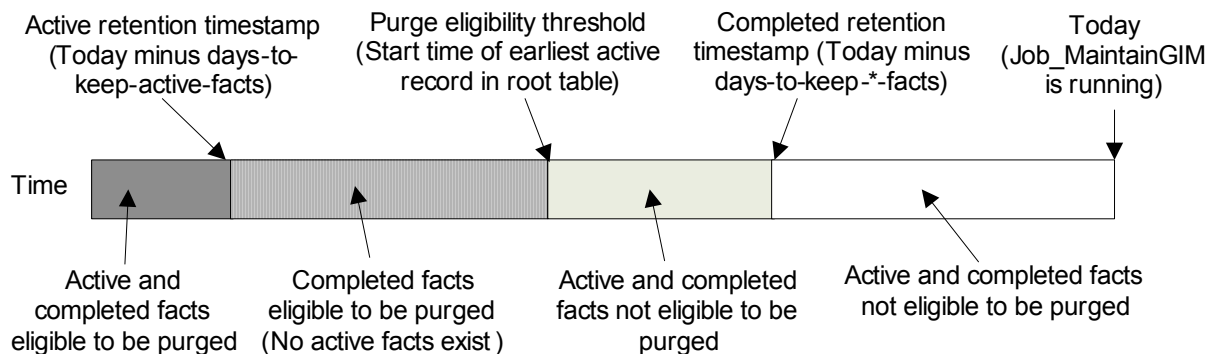
Figure 2: Purge Threshold for Long-Living Entities

Voice and Multimedia Interaction Data

Genesys Info Mart does not extract data for active voice interactions. By definition, therefore, all facts relating to voice interactions in both GIDB and the dimensional model are completed. Facts relating to multimedia interactions in both GIDB and the Dimensional Model might be either active or completed.

Voice and multimedia interaction data is stored in separate, media-specific tables in GIDB. In the algorithm that the maintenance job uses to purge GIDB, voice interactions are considered to be short-living entities, and multimedia interactions are considered to be long-living entities.

Voice and multimedia interaction data is stored in common tables in the dimensional model. In the algorithm that the maintenance job uses to purge the

dimensional model, both voice and multimedia interactions are considered to be long-living entities.

> **Note:** Active multimedia interactions might delay the purging of voice data from the dimensional model. For example, if `days-to-keep-gim-facts=400, days-to-keep-active-facts=600,` and the earliest active interaction is an e-mail that has been left in a workbin for 550 days, then voice interactions will be retained for 550 days.

Retention Period for Discarded Operational Data

The retention period for records in the discard, audit, and history tables is calculated from the time that the ETL process discarded the data. The discard tables are Staging tables that store operational data that the transform job was unable to process—for example, voice interaction data with unresolved IS-Links, or Configuration facts (configuration relationship 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 tables is useful for troubleshooting. Purging of discarded data is controlled by the `days-to-keep-discards-and-job-history` option in the `[gim-etl]` section.

Purging Rules

The following rules apply to purging:

- Facts are evaluated for purging based on the start time of the fact.
- Both active and completed facts are eligible for purging, but they have different retention periods, controlled by separate configuration options. Similarly, the retention periods for completed facts in the dimensional model and in GIDB are different.
- For both short-living and long-living entities, records are purged independently from the applicable tables, for all records that have a start time earlier than the purge threshold.
  - For short-living entities, the purge threshold is the configured data retention period for completed facts, as specified by the `days-to-keep-gim-facts` and `days-to-keep-gidb-facts` options.
  - For long-living entities, the purge threshold is determined by the start time of the earliest active fact in the interval between the completed and active retention periods, as specified by the `days-to-keep-gim-facts` and `days-to-keep-gidb-facts` and `days-to-keep-active-facts` options, respectively.

  For more information, see "Retention Periods for Facts" on page 39 and "Purging the Genesys Info Mart Database" on page 62.

- Discarded operational data is evaluated for purging based on the date portion of the timestamp of the ETL processing event.
- Dimension table rows are not eligible to be purged in either the dimensional model or GIDB.

Purging Mechanism

The actual SQL commands that `Job_MaintainGIM` issues depend on whether the tables are partitioned.

- When `Job_MaintainGIM` deletes rows in nonpartitioned tables, the job issues SQL `DELETE` operations against the tables. Running this job daily results in a small percentage of the table being deleted, which minimizes the time that it takes the RDBMS server to find the rows, delete them, and make index adjustments.
- When `Job_MaintainGIM` purges partitioned tables, the job issues the appropriate SQL commands against the tables, as required by the RDBMS implementation, to drop partitions.

Scheduling

To use `Job_MaintainGIM` to purge data, configure the time of day you want Genesys Info Mart Server to launch this job. This must be done once a day.

For more information about enabling or disabling a purging schedule, see "Setting Scheduling Options for Genesys Info Mart Server", particularly Step 10 on page 52.

Optimizing the Transaction Size

You can configure Genesys Info Mart to control the size of database transactions for purging. For this job, the purge transaction size is specified as the number of deleted records per table that will be committed in a single transaction. The `purge-transaction-size` option in the section `[gim-etl]` controls the transaction size.

You can also specify the maximum number of concurrent purging transactions. The `purge-thread-pool-size` option controls this aspect of `Job_MaintainGIM` operation.

Finding optimal settings for these options for your environment should help you balance purging execution time against the RDBMS resources that are required to purge the Info Mart database.

Database Partitioning

You can use partitioning on Oracle 10g or 11g databases (range partitioning only) and on Microsoft SQL Server databases. Fact tables and associated

indexes in GIDB and the dimensional model are partitioned. Configuration object tables, configuration relationship Fact tables, and Dimension tables are not partitioned. To create partitioned schema, use `make_gim_partitioned.sql` instead of `make_gim.sql`.

During initialization, `Job_InitializeGIM` creates the first set of partitions, and `Job_MaintainGIM` subsequently creates additional partitions as required to be populated during ETL cycles.

Three configuration options are used to specify the size of the partitions in GIDB and the dimensional model (`partitioning-interval-size-gidb` and `partitioning-interval-size-gim`), and to control how far ahead the Genesys Info Mart jobs (`Job_InitializeGIM` in the first instance, then `Job_MaintainGIM` on an ongoing basis) will create partitions, in preparation for future ETL cycles (`partitioning-ahead-range`). For additional information on these configuration options and their settings, see the *Genesys Info Mart 8.0 Deployment Guide*.

Job_MigrateGIM

When you need to migrate from an existing Genesys Info Mart 8.x deployment to the most recent 8.x release of Genesys Info Mart, you will run `Job_MigrateGIM` as a part of the transition process. `Job_MigrateGIM` automatically runs all of the scripts necessary to prepare your existing Info Mart database for use with the current release of Genesys Info Mart 8.x.

- When you restart an upgraded Genesys Info Mart Server application with an unmigrated Info Mart database, Genesys Info Mart automatically detects the out-of-date database schema version and puts the Genesys Info Mart Server into migration state. In this state, you cannot run any jobs other than `Job_MigrateGIM`.

- Ensure that no queries or other activities are performed against the Info Mart database while `Job_MigrateGIM` runs. Be sure to take this into account when you plan migration.

Note: `Job_MigrateGIM` migrates only existing 8.x deployments to the most recent 8.x release of Genesys Info Mart. There is no migration path from 7.x to 8.x.

For complete migration preparations and procedures, see the Genesys Info Mart 8.x section of the *Genesys Migration Guide*.

3 Working with Jobs

This chapter describes how to execute the Genesys Info Mart jobs, either automatically or manually, using Genesys Info Mart Server and Genesys Info Mart Administration Console.

Note: For descriptions of the jobs, see "Understanding Genesys Info Mart Jobs" on page 17.

This chapter provides information on the following:

Note: Before you can execute any job, you must complete the tasks in the *Genesys Info Mart 8.0 Deployment Guide*.

Configuration Checker
The Genesys Info Mart Server checks the integrity of the Genesys Info Mart deployment and prevents any new job from starting if it encounters errors in the items it checks.

Genesys Info Mart automatically performs this check:

- On startup.
- Whenever there are changes to the configuration of the Genesys Info Mart Application (option settings or connections).
- Whenever there are changes to the configuration of any ICON applications that are in the Genesys Info Mart application connections (in other words, ICON Applications that are on the Connections tab of the Genesys Info Mart Application object).

- Whenever there are changes to the configuration of any contact center configuration objects that can override settings in the Genesys Info Mart `Application` (option settings on `Switch`, `Business Attribute`, `DN`, or `Script` objects).

Note: For more information about the items that Genesys Info Mart checks, see the section about deployment verification in the "Data Processing" chapter of the *Genesys Info Mart 8.0 Deployment Guide*. For more information about how Genesys Info Mart handles other configuration errors that it might encounter, see "Configuration Errors" on page 72.

Scheduling Jobs with Genesys Info Mart Server

The Genesys Info Mart Server launches jobs based on the schedule options that you configure in the Genesys Info Mart application, with any necessary adjustments to accommodate logical rules that guide scheduling. (For convenience, this function is called the Scheduler.)

Scheduling the ETL Cycle

The basic unit of scheduling is the ETL cycle. The ETL cycle performs the following functions:

- Extracts data from each data source.
- Transforms the extracted data, and loads the transformed data into the Dimension tables and the Fact tables.

Additional jobs run on a scheduled basis to perform the following functions:

- An optional aggregation process, which runs in parallel with the ETL cycle, runs continuously within configured time intervals to populate Aggregate tables, in an environment where either Genesys Interactive Insights (GI2) reports or Reporting and Analytics Aggregates (RAA) package are deployed.
- `Job_MaintainGIM,` which runs outside the ETL cycle, purges data from the Info Mart database, in accordance with configurable data retention policies. The job also maintains the default and custom calendars.

ETL Cycle Configuration Options

You configure options that control the ETL cycle, including:

- Time of day the first ETL cycle should begin
- Time of day the final ETL cycle should begin

- Frequency of the ETL cycle

You configure additional options to specify:

- Whether calculation of aggregates occurs in parallel with the ETL cycle.
- The times of day when the purging of old Info Mart data should start and end.
- The time zone in which the schedule will be defined.

You can also set configuration options to:

- Temporarily stop Genesys Info Mart Server from launching scheduled jobs.
- Stop Genesys Info Mart Server from launching the job that calculates the Aggregate tables.
- Stop Genesys Info Mart from launching the job that purges old data from the Info Mart database.

Job Sequencing Rules

`Job_ExtractICON`, `Job_TransformGIM`, and `Job_MaintainGIM`, have a logical dependence on each other. That is, until you perform an extraction, there is nothing to transform, and until you have performed extraction and transformation, there is no data to maintain.

When you have Scheduler run jobs, it does not start a particular job until the logically prior job(s) are completed. You can start a job manually at any time, but if the logically prior job is not completed, the "later" job has no data to process.

Scheduler manages jobs automatically according to the following rules, whether you use the Genesys Info Mart Administration Console to run a job immediately or schedule it to run at some future time:

- Scheduler does not start a job if there is another instance of that job already running.
- During a scheduled ETL cycle, Scheduler does launch the transform job after the extract job completes. You can launch the transform job manually from the Genesys Info Mart Administration Console while an extract job is running, but the transform job will not do anything if no data has been extracted.
- Scheduler ensures that the extract and transform jobs do not run at the same time as the maintenance job, which purges Info Mart data. You can launch the maintenance job manually from the Genesys Info Mart Administration Console while an extract or transform job is running, but the maintenance job is delayed until extraction and transformation are completed.

  If a scheduled ETL cycle is set to begin before a maintenance job is finished, it stops the maintenance job and starts the ETL cycle.

- In the case of the aggregation job, Genesys Info Mart Server does not allow an instance of the job to run outside the regularly scheduled intervals within which `Job_AggregateGIM` has been configured to run. For example, if `Job_AggregateGIM` has been configured to run every day between 01:00 AM and 06:00 AM, you will not be able to launch `Job_AggregateGIM` manually from the Genesys Info Mart Administration Console at any time outside that time period (for example, at 08:00 AM).

  For information about running `Job_AggregateGIM` manually, see "Managing Jobs with Genesys Info Mart Administration Console" on .

Using the Solution Control Interface

Genesys Info Mart Server can be started, stopped, and monitored by using the Genesys Solution Control Interface (SCI).

Note: For a more detailed discussion of running Genesys Info Mart via SCI, see the "Starting and Stopping" chapter of the *Genesys Info Mart 8.0 Deployment Guide*.

Starting Genesys Info Mart Server

When you start Genesys Info Mart Server by using SCI, the Genesys Info Mart application status reflects the status of the Genesys Info Mart Server itself, and not the status of any jobs. In other words, the `STARTED` status that is reported by Solution Control Server (SCS) indicates that Genesys Info Mart Server is operational, but it does not indicate whether jobs are currently running or whether a job has failed.

Viewing Job Status

To learn the job status, you can:

- Use the Genesys Info Mart Administration Console.
- View the job status in the `ADMIN_ETL_JOB_HISTORY` and `ADMIN_ETL_JOB_STATUS` database views.
- Check the logs (for example, in the Centralized Log Database) for the job status messages.

Stopping Genesys Info Mart Server

When you stop the Genesys Info Mart Server using SCI, the server shuts down all currently running jobs and terminates gracefully. The `STOPPED` status reported by SCS indicates that Genesys Info Mart Server has stopped, but it does not indicate the status of jobs.

Note: When Genesys Info Mart is installed on the Windows platform, it is installed as a Windows service. The startup type for the Windows service is Automatic. If the machine is restarted, the Windows service automatically launches Genesys Info Mart. For more information on Windows service, see the *Genesys Info Mart 8.0 Deployment Guide.*

Connecting to Configuration Server

Genesys Info Mart Server keeps an active connection to Configuration Server. This enables it to receive notification of any configuration changes that affect its operation. The Genesys Info Mart Server adjusts to any dynamic changes to the configuration options in the corresponding Genesys Info Mart Server `Application` object. Configuration changes that affect the operation of a currently running job take effect the next time the job starts.

If Genesys Info Mart Server cannot connect to Configuration Server on startup, it reads the values of configuration options previously stored in a local file and re-attempts to make a connection every 30 seconds.

Note: When the Genesys Info Mart Server cannot connect to Configuration Server because the Genesys Info Mart application in SCI is already connected, the Genesys Info Mart Server exits immediately. This situation may occur when the Genesys Info Mart Server is currently running and an attempt is made to start it from another location.

If the Genesys Info Mart Server is able to connect to Configuration Server, it retrieves the backup Configuration Server information from the `Server Info` tab. If Genesys Info Mart Server then later loses its connection to Configuration Server, Genesys Info Mart Server repeats the following process until a connection is established. Genesys Info Mart Server:

1. Attempts to connect to Configuration Server.

2. Attempts to connect to the backup Configuration Server, if configured.

3. Waits 30 seconds, the time specified for the reconnection timeout.

Setting Scheduling Options for Genesys Info Mart Server

The information in this subsection supplements information in the *Genesys Info Mart 8.0 Deployment Guide* about configuring the Genesys Info Mart `Application`.

Use the following steps to configure the Genesys Info Mart Server `schedule` options. Genesys Info Mart Server uses these options to launch the ETL, aggregation, and maintenance jobs. Each configuration option is related to one or more of the jobs.

Job schedules are defined in 24-hour time spans in the format `HH:mm`, where `HH` is the number of hours (00–23), and `mm` represents the number of minutes (00–59). The 24-hour schedule can span two calendar days. For example, if the `etl-start-time` is defined as `18:00` and the `etl-end-time` is defined as `06:00`, the start time is 6:00 PM one day and the end time is 6:00 AM the following day.

For information on the functions of the jobs, see "Job Summary Table" on page 17. For information on job interdependencies, see "Job Sequencing Rules" on page 47.

To set a job schedule:

1. In Configuration Manager, navigate to the `Options` tab of the Genesys Info Mart application.

2. Navigate to the `[schedule]` section.

Time Zone 3. (Optional) Enter a value for `timezone` to specify a local time zone in which you want to define the schedule. 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. The default time zone is `GMT`.

For more information about supported time zones, see the Java documentation about calendar time zones on the Java developer web site (http://java.sun.com).

ETL Schedule 4. Enter a value for `etl-start-time` to specify the time of day the first ETL cycle begins.

5. Enter a value for `etl-end-time` to specify the time of day the final ETL cycle begins. This should be a time of day when no other ETL cycles will begin.

6. Enter a value for `etl-frequency` to specify the number of minutes between the start times of adjacent ETL cycles.

If the time it takes to complete a cycle is shorter than this value, the next cycle is delayed until the time is met. If the time it takes to complete a cycle is greater than this value, the next cycle starts immediately.

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

Note: Various `extract-*` options in the `[gim-etl]` section control aspects of extract and transform job functioning that significantly affect ETL cycle performance. When setting the ETL scheduling options, consider the values of these related options as well. For more information, see "Controlling the Volume of Data Extraction" on page 31. See also the option descriptions in the configuration options reference chapter in the *Genesys Info Mart 8.0 Deployment Guide*.

---

7. To start or resume the ETL schedule, set `run-scheduler` to `TRUE`. You can set this option to `FALSE` to temporarily stop Genesys Info Mart Server from launching jobs.

Aggregation Schedule

8. If you plan to use the Info Mart historical Aggregate tables:

    a. Set `run-aggregates` to `TRUE`. This option specifies whether the Genesys Info Mart Server launches the job, `Job_AggregateGIM`. This job calculates the Aggregate tables based on newly added or changed Fact table data.

    b. Set `aggregate-schedule` to the time you want the aggregation job to start (as long as it is not currently running, such as following the initial deployment of the aggregates). The value must be expressed as a valid CRON expression: a string, in which in this case uses only two fields, minute and hour, separated by white space,

        For example, to set the aggregation start time to 2:30 AM, enter the value `30 2`. More examples follow in "Sample Schedule" on page 52.

        Note: For a detailed description of this option, see the Configuration Options section of the *Genesys Info Mart 8.0 Deployment Guide*.
        ---

    c. Set `aggregate-duration` to length of the period, in `HH:mm` format, during which the aggregation job will run after each launch. Within the time intervals defined by the `aggregate-schedule` and `aggregate-duration` options, the aggregation job runs continuously.

        For example, to have aggregation run for the twelve hours following the start time, set the value for the duration option to `12:00`. More examples follow in "Sample Schedule" on page 52.

    d. If you want the aggregation job to run continuously for an indefinite period, without any breaks (for example, for daily maintenance activities), set `aggregate-duration=24:00`.

    e. Set `aggregation-engine-class-name` as described in the *Genesys Info Mart 8.0 Deployment Guide*. This option specifies the class name of the aggregation package.

    Note: You must deploy either Genesys Interactive Insights (GI2) or Reporting and Analytics Aggregates (RAA), in order for Aggregate tables to be created and populated in the Info Mart database.
    ---

9. If you do *not* plan to use the Info Mart Aggregate tables, keep the default value of the `run-aggregates` option (`FALSE`). This setting ensures that the Genesys Info Mart Server does not launch `Job_AggregateGIM` in deployments where the Genesys-provided aggregation engine is not deployed.

Maintenance Schedule

10. If you plan to purge eligible fact data from the Info Mart database (Global Interaction Database [GIDB], the dimensional model, Staging tables, and Discard tables):

   a. Set `maintain-start-time` to the time of day that you want Genesys Info Mart to launch `Job_MaintainGIM`. The time of day must be outside the range that is specified by `etl-start-time` and `etl-end-time`. `Job_MaintainGIM` will not start until the `maintain-start-time` has been reached.

Sample Schedule

Figure 3 on depicts a sample Genesys Info Mart Server job schedule that runs repeated ETL cycles throughout the day. The deployment includes aggregation, and the aggregation job has been scheduled to run without interruption.

The schedule options are set as follows:

• `etl-start-time` = 06:30	• `maintain-start-time` = 03:00
• `etl-end-time` = 00:30	• `aggregate-schedule` = 0 0
• `etl-frequency` = 30 minutes	• `aggregate-duration` = 24:00
	• `timezone` = <business user local time zone>

The schedule has been defined in the local time zone of the business user that will query the Info Mart data. Therefore, from the perspective of the business user, the Genesys Info Mart server is in the same time zone, and the final extract, which begins at 00:30, contains all the reporting data from the previous day.

Note: Midnight in local time is considered to be the end of the reporting day for a particular time zone.
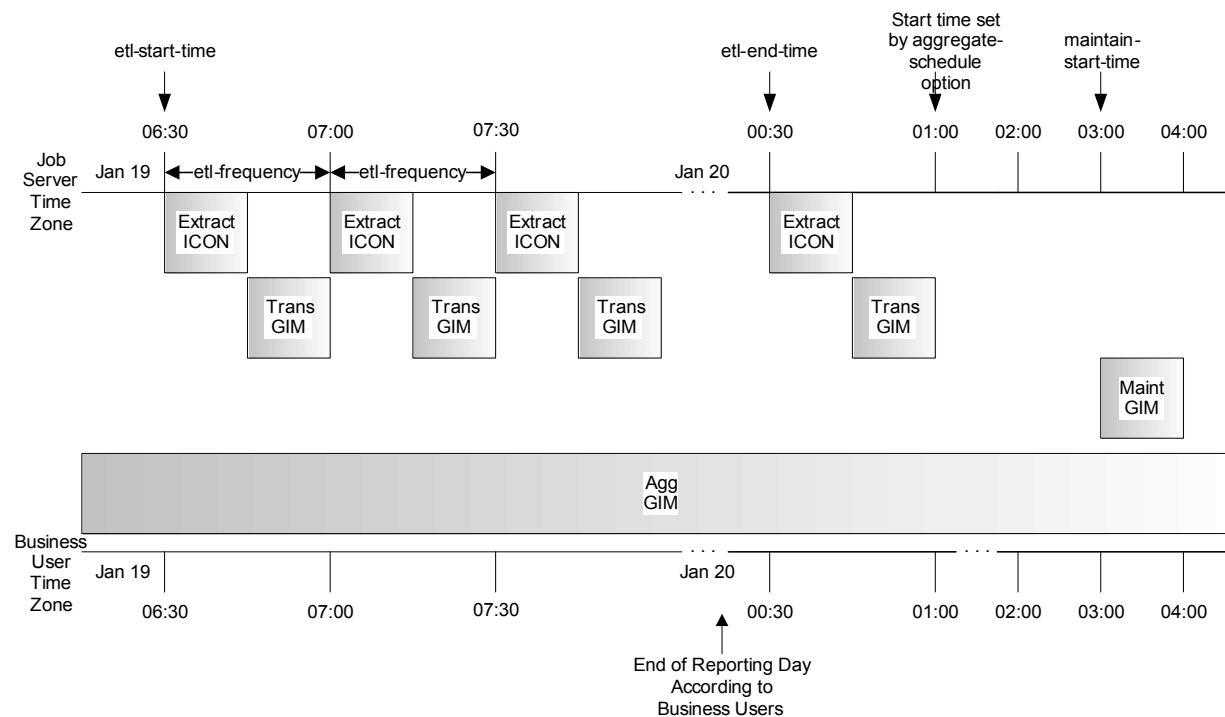
Figure 3: Sample Genesys Info Mart Server Schedule

Managing Jobs with Genesys Info Mart Administration Console

You can use the Genesys Info Mart Administration Console to monitor job status and, when necessary, manually start or stop a job (outside of the normal schedule).

Although the Genesys Info Mart Server automatically launches scheduled jobs, you can use the Genesys Info Mart Administration Console to:

* Execute or schedule one or more jobs to recover from job failures.
* Execute or schedule Job_ExtractICON following the update of configuration data in IDB with the Interaction Concentrator (ICON) on-demand resynchronization feature.
* Display the current job execution status.
* Display a history of job execution, such as start time, stop time, duration, and final status.
* Filter the jobs based on time and status.
* Execute a single job as needed (either now or at some future specified time and date).
* Remove a scheduled job.

•   Selectively shut down a running job.

Accessing Genesys Info Mart Administration Console

Genesys Info Mart Administration Console is a graphical user interface that, by using the Wizard Framework, functions as an extension to Configuration Manager.

To access the Administration Console:

1. Open Configuration Manager.

2. Select the Genesys Info Mart application object.

3. Right-click and select `Wizard > Configure` options. This will invoke the Genesys Info Mart Administration Console for the Genesys Info Mart application object.

Monitoring Jobs

You can monitor any job in your environment. You can monitor the overall status of jobs or the statistics that are associated with a job.

To monitor your jobs, view the `Status` tab window in the Genesys Info Mart Administration Console. Figure 4 shows the `Status` tab window.
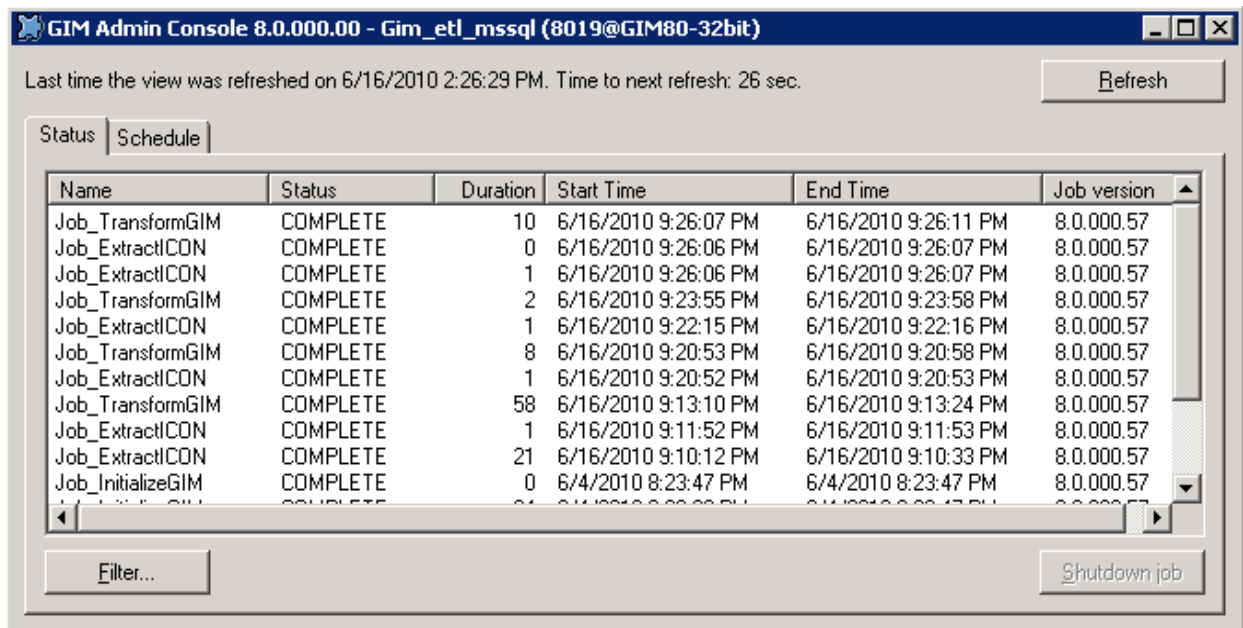
Figure 4: Genesys Info Mart Administration Console Status Tab

For each job, the Genesys Info Mart Administration Console `Status` tab displays the following information:

•   The current state of the job on each refresh cycle. The status for each job is listed under the `Status` column. Values for status include:

- ◆ `Running`
- ◆ `Complete`
- ◆ `Shutdown`
- ◆ `Waiting`
- ◆ `Failed`

- The time of the last update and the time until the next update. To determine when a job completed, see the `End Time` column.

Note: The Genesys Info Mart Administration Console displays timestamps in the time zone of the host on which the Genesys Info Mart Administration Console is installed.

The Genesys Info Mart Administration Console view refreshes the status of the jobs:

- Automatically every minute

- When the `Refresh` button is clicked

- After a job is shut down with the `Shutdown Job` button

Customizing the Display

You can customize the way the Genesys Info Mart Administration Console displays information:

- You can resize columns in the Genesys Info Mart Administration Console window view by clicking and dragging the column heading.

- You can sort the jobs alphabetically or numerically by clicking the column heading.

- You can filter the view to determine which jobs will appear. See "Filtering Genesys Info Mart Administration Console View" on page 60 for more information.

Executing and Scheduling Jobs

Use Genesys Info Mart Administration Console to run a job once, either immediately or at a scheduled time.

Normally, the Genesys Info Mart Server launches scheduled jobs automatically. However, you can use the Genesys Info Mart Administration Console to:

- Run one or more jobs to recover from job failures.

- Run `Job_ExtractICON` following the update of configuration data in IDB with the ICON on-demand resynchronization feature.

- Manually run the job that aggregates Info Mart data, provided that Genesys Info Mart has not been configured to run the aggregation job on a schedule. In other words, you can start or stop `Job_AggregateGIM` from the Genesys Info Mart Administration Console only if the `run-aggregates` configuration option, in the `[schedule]` section, has been set to `false`.

- Run the maintenance job outside the scheduled time, provided that no instances of other jobs are running.

The following subsections describe how to use the Genesys Info Mart Administration Console to manage jobs:

Running Jobs One-by-One

Initially after deploying Genesys Info Mart, you may want to run the jobs one-by-one to test the best values for various configuration options, rather than immediately scheduling them to run routinely. Or, rarely, you may need to run the jobs one-by-one while troubleshooting. In these special situations, use the following procedure to run the jobs one-by-one:

ETL Cycle Jobs 1. Set the `run-scheduler` configuration option (in the `[schedule]` section on the Genesys Info Mart `Options` tab) to `FALSE`.

2. Follow the steps in "Running a Job Immediately" on page 57 to run the jobs that perform ETL processing.

   Run the jobs in the following order:
   - `Job_ExtractICON` (see the following Note)
   - `Job_TransformGIM`

   > **Note:** The extract job extracts data from all available IDBs, for all extract roles. When you have completed your trial runs or troubleshooting, restore normal running conditions by setting the `run-scheduler` configuration option to `TRUE` to have the Genesys Info Mart Server launch jobs based on a schedule you configure in the Genesys Info Mart application. For more information, see "Setting Scheduling Options for Genesys Info Mart Server" on page 49.

Aggregation Job 3. If your deployment includes aggregation and you want to run `Job_AggregateGIM` as a non-scheduled job, set the `run-aggregates` configuration option (in the `[schedule]` section on the Genesys Info Mart `Options` tab) to `FALSE`.

4. Follow the steps in "Running a Job Immediately" on page 57 to start `Job_AggregateGIM`. The job will run continuously until you manually stop it, or until you reset the `run-aggregates` configuration option from `FALSE` to `TRUE` and the configured daily schedule comes into effect.

Maintenance Job　**5.** Set the `run-maintain` configuration option (in the [`schedule`] section on the Genesys Info Mart `Options` tab) to `FALSE`.

6. On the `Status` tab of the Genesys Info Mart Administration Console, verify that no other jobs are running.

7. Follow the steps in "Running a Job Immediately" on page 57 to start `Job_MaintainGIM`.

After the initial run, you can set the `run-maintain` configuration option to `TRUE`, so that the Genesys Info Mart Server will launch the maintenance job based on the schedule you configured in the Genesys Info Mart application. For more information, see "Setting Scheduling Options for Genesys Info Mart Server" on page 49.

Running a Job Immediately

Use the following steps to run a job immediately:

1. In the Genesys Info Mart Administration Console, select the `Schedule` tab. The window displays any scheduled jobs that have not started yet.

2. Click the `Run Job` button at the bottom of the window. The `Run Job` window appears, as shown in Figure 5.
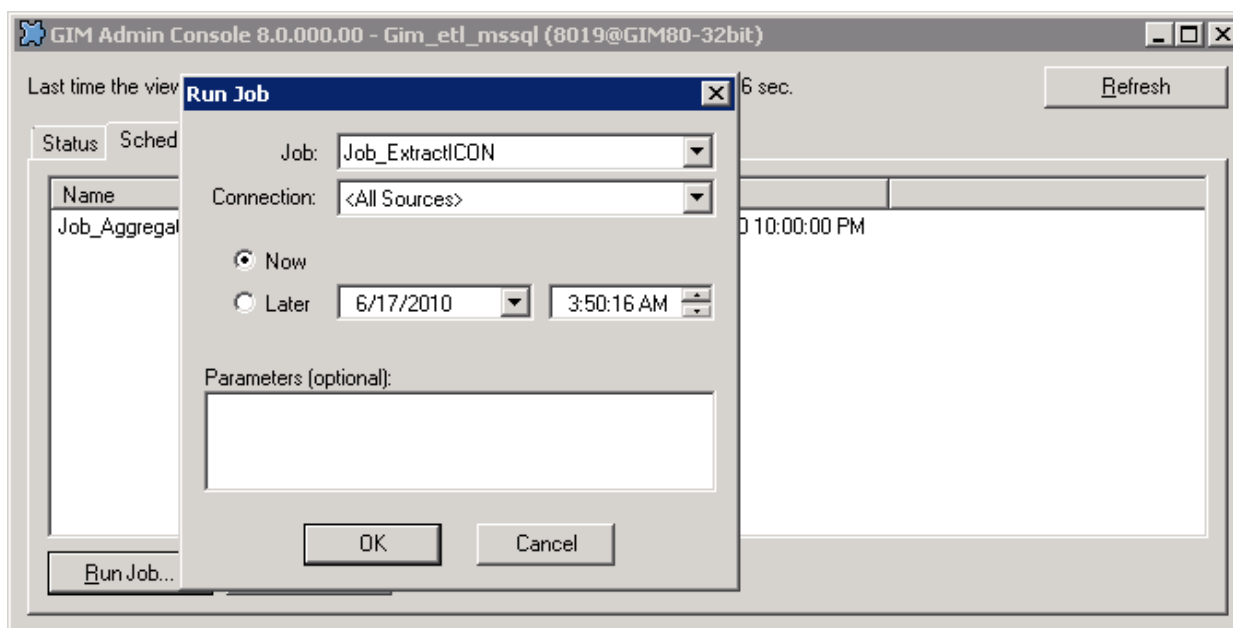
Figure 5: Administration Console Run Job Window

3. From the `Job` drop-down list, select the job that you wish to execute.

4. Select the option to run the job `Now`.

> **Note:** Genesys Info Mart 8.0 does not accept any parameters for any job.

5. Click the `OK` button.

6. Click the `Status` tab to see the job that you have just started.

> **Note:** Jobs that are scheduled to run immediately appear on the `Status` bar, but not on the `Schedule` tab. These jobs will be in a `Waiting` or `Running` state. If an error occurred during job startup, the status will be `Failed`.

Scheduling a Job to Run Later

Rather than executing a job immediately, you can use Genesys Info Mart Administration Console to schedule the job to run at a future time.

Use the steps below to schedule a job to run at a future time:

1. In the Genesys Info Mart Administration Console, select the `Schedule` tab. The window displays any scheduled jobs that have not started yet.

2. Click the `Run Job` button at the bottom of the window. The `Run Job` window appears, as shown in Figure 5 on .

3. From the `Job` drop-down list, select the job that you wish to execute.

4. Select the option to run the job `Later`. From the applicable drop-down lists, select the date and time that you want the job to run.

5. Click the `OK` button.

Jobs that are scheduled to run in the future appear on the `Schedule` tab until the scheduled time has been reached. At that point, the job is removed from the `Schedule` tab and appears on the `Status` tab, with a status of `Waiting` or `Running`.

Canceling a Scheduled Job

To cancel a previously scheduled job:

1. In the Genesys Info Mart Administration Console, select the `Schedule` tab. The window displays any scheduled jobs that have not started yet.

2. Select the job from the `Scheduled` list and click the `Cancel Scheduled Job` button.

A message appears to confirm that the job has been canceled.

Shutting Down a Job

The Genesys Info Mart Administration Console enables you to shut down a running job by using the following steps:

1. If the job that you want to shut down is the aggregation job, set the `run-aggregates` configuration option (in the `[schedule]` section on the Genesys Info Mart `Options` tab) to `FALSE`.

   Note: If `run-aggregates` is set to `TRUE`, the scheduler ensures that `Job_AggregateGIM` runs continuously during the interval specified by the between the scheduled start time and duration. If `Job_AggregateGIM` is running under the control of the scheduler and you try to stop the job during the interval in which it has been configured to run, the scheduler will automatically restart the job almost immediately.

2. In the Genesys Info Mart Administration Console, select the `Status` tab. The window displays the status of current and completed jobs (see Figure 4 on).

3. Select the job that you want to stop. You can stop jobs that have a status of `Running` or `Waiting`.

4. Click the `Shutdown Job` button at the bottom of the window.

   A confirmation window appears (see Figure 6).
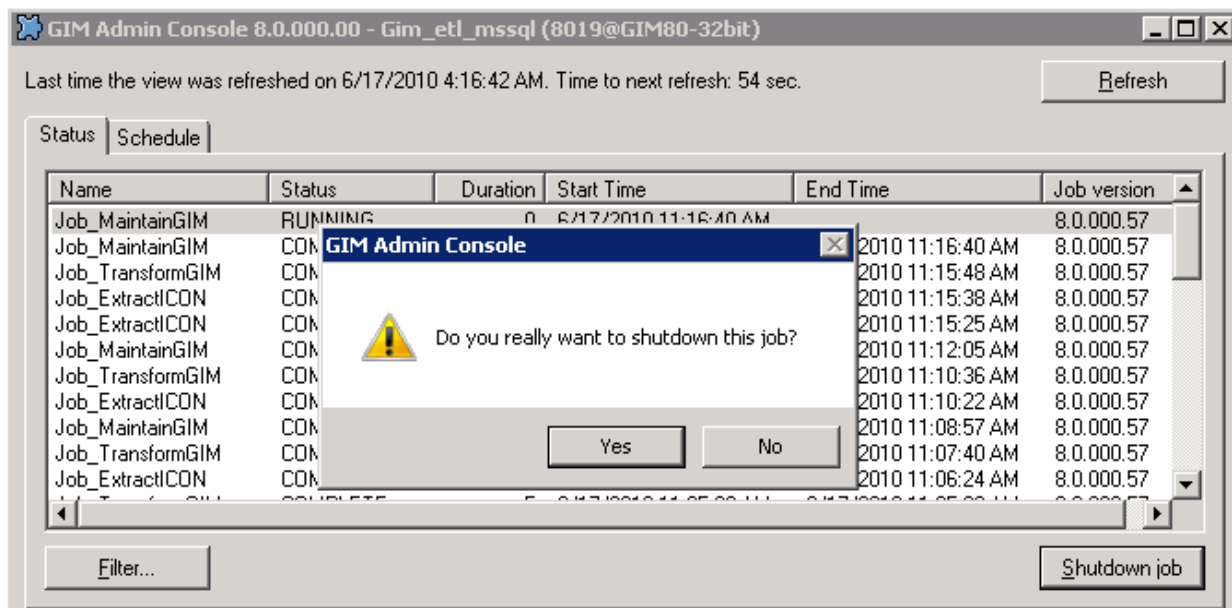
Figure 6: Shutdown Job Confirmation

5. Click `Yes` to confirm shutdown or `No` to cancel the shutdown.

After the `Shutdown Job` command is confirmed, it may take a few moments for the job to shut down.

Filtering Genesys Info Mart Administration Console View

You can filter the view of the Genesys Info Mart Administration Console to determine which jobs will appear on the tabs.

1. In the Genesys Info Mart Administration Console, select the `Status` tab. The window displays the status of current and completed jobs (see Figure 4 on).

2. Click the `Filter...` button. The `Filter` dialog box appears (see Figure 7).
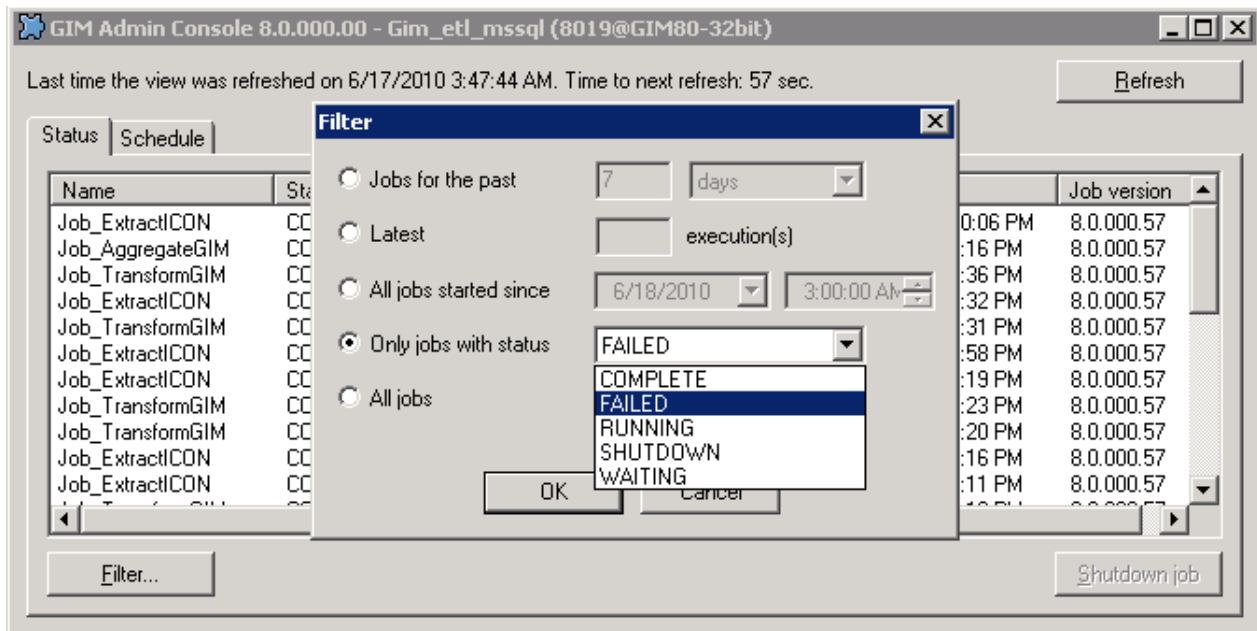
Figure 7: Filtering Options for Administration Console

3. Select the filtering options that you want. Filters can be set to show:
   - Jobs that have run in the past over a specified time interval.
   - A specified number of recent jobs that have run.
   - Jobs that have run since a specified date.
   - Jobs that have a specified status. Available status options are:
     - `Complete`
     - `Failed`
     - `Running`
     - `Waiting`
     - `Shutdown`
   - All jobs in your environment.

4. Click `OK` to set the filter. The jobs that are visible in the Genesys Info Mart Administration Console view change accordingly.

> **Note:** Once you have set a filter in the Genesys Info Mart Administration Console, the filter preference is saved. The next time that you use the `Filter...` option on the `Status` tab, the filter options will be preset to the last settings that you used.

Changing Calendar Dimensions

After the default or custom calendar dimension tables have been populated, you can safely change the values of the options in the `[date-time]` or custom `[date-time-*]` configuration section that control when `Job_MaintainGIM` populates the table with the next batch of calendar dimensions (`date-time-min-days-ahead`) and how far ahead the table will be populated (`date-time-max-days-ahead`).

However, if you want to change the values of any of the other `date-time` options during runtime, you must perform additional steps to avoid compromising the consistency of calendar data, which can adversely affect your reporting results. For example, if you change the time zone option (`date-time-tz`) without performing the additional steps, your reports might mix the results for different time zones within the same reporting interval.

The following procedure describes how you can safely modify the settings that control the content of existing calendar dimension tables.

Changing Calendar Settings During Runtime

1. If your deployment includes aggregation, stop aggregation. For information about stopping the aggregation process, see the *Reporting and Analytics Aggregates 8.0 User's Guide.*

2. In Configuration Manager, change the settings for the `date-time` options in the Genesys Info Mart `Application` object, as required. For information about the available options and valid values, see the configuration options reference chapter in the *Genesys Info Mart 8.0 Deployment Guide*.

3. Manually truncate the corresponding calendar table in the Info Mart database.

4. Run `Job_MaintainGIM`. For information about running the job manually, see "Running a Job Immediately" on .

5. If your deployment includes aggregation:
   a. Restart aggregation.
   b. Re-aggregate your data. For information about re-aggregating data over a certain time range, see the *Reporting and Analytics Aggregates 8.0 User's Guide.*

Purging the Genesys Info Mart Database

Purging old data from a database should be performed periodically to prevent database size from growing too large. This section illustrates strategies for purging old data from the Genesys Info Mart database.

In GIDB, voice interaction data and multimedia interaction data are stored in separate, media-specific tables.

- When purging *GIDB*, voice interactions are treated as short-living entities, and multimedia interactions are treated as long-living entities.

- When purging the *dimensional model*, voice interaction data and multimedia interaction data are stored in common tables. For the purposes of purging the dimensional model, both voice interactions and multimedia interactions are treated as long-living entities. Be aware that active multimedia interactions might delay the purging of voice data in the dimensional model.

  For example, if `days-to-keep-gim-facts=400` but, on the day that `Job_MaintainGIM` is running, the earliest active interaction is an e-mail that has been left in a workbin for 550 days, then voice interaction data are retained for 550 days.

- The transition from one day to the next is considered to occur at 00:00:00 (midnight) in GMT.

For more information about how `Job_MaintainGIM` applies data retention policies and defines short-living and long-living entities for purging purposes, see "Retention Periods" on page 39.

For information about purging old data from the Interaction Database (IDB), see "Interaction Concentrator Purge Procedures" on page 79.

Note: In partitioned databases, data is purged by partition. Therefore, partitioning might further delay the purging of some voice interactions because a partition will be dropped only when all data in that partition is eligible to be purged.

Purging Fact Data from the Dimensional Model

Figure 8 on page 65 shows a purging eligibility example for Fact tables in the dimensional model. When looking at the example, keep in mind the following:

- The `days-to-keep-gim-facts` is configured as 400 days.
- The `days-to-keep-active-facts` is configured as 600 days.
- Data has been extracted until January 21, 2012.

`Job_MaintainGIM` determines the maximum start dates for purge eligibility by the following actions:

- `Job_MaintainGIM` subtracts 400 days from January 21, 2012, which is December 17, 2010. This is the retention timestamp for short-living entities, and the upper bound of the purge eligibility interval for completed long-living interactions.

- `Job_MaintainGIM` subtracts 600 days from January 21, 2012, which is June 1, 2010. This is the retention timestamp for active long-living entities.

- From the root table in the group of related tables for the particular type of data, `Job_MaintainGIM` calculates the start time of the earliest active fact in the time range between the `days-to-keep-gim-facts` and the `days-to-keep-active-facts` timestamps. For example, for multimedia interactions, the root table is `STG_ACTIVE_IF`.

  - The start time of the earliest active fact between December 17, 2010 and June 1, 2010 sets the purge threshold for completed long-living entities in the group of interaction-related tables. In the example shown in Figure 8, this date is August 3, 2010.

    In all related tables, completed facts that started on August 3, 2010 or earlier are eligible to be purged.

  - Active facts that started earlier than June 1, 2010 are also eligible to be purged.

In Figure 8, `Job_MaintainGIM` purges facts relating to the following entities:

Long-Living Entities Purged

- `Active Multimedia Interaction 1`, because the start date is earlier than the retention timestamp for active facts.

- `Completed Multimedia Interaction 3`, because the start date is earlier than the calculated purge threshold for completed long-living entities.

- `Voice Interaction 1` (by definition, a completed fact), because the start date is earlier than the calculated purge threshold for completed long-living entities.

- `Outbound Contact Campaign 1`, because, for the group of Outbound Contact–related tables, there are no active facts within the time range between the `days-to-keep-gim-facts` and the `days-to-keep-active-facts` timestamps, and the start date is earlier than the `days-to-keep-gim-facts` timestamp.

Short-Living Entities Purged

- `Agent Login Session 1`, because the start date is earlier than the retention timestamp for short-living entities.

`Job_MaintainGIM` does not purge facts relating to the following entities:

Long-Living Entities Not Purged

- `Active Multimedia Interaction 2`, because the start date is not earlier than the retention timestamp for active facts. However, as the earliest active interaction within the interval defined by the retention policy options for active and completed facts, `Active Multimedia Interaction 2` sets the purge threshold for completed long-living entities.

- `Completed Multimedia Interaction 4`, because the start date is later than the calculated purge threshold for completed long-living entities.

- `Voice Interaction 2` or `Voice Interaction 3`, because the start date is later than the calculated purge threshold for completed long-living entities.

- `Outbound Contact Campaign 2`, because it is still active and the start date is not earlier than the `days-to-keep-gim-facts` and `days-to-keep-gidb-facts` timestamps.

Short-Living Entities Not Purged

- `Agent Login Session 2`, because the start date is later than the retention timestamp for short-living entities.

Note: Figure 8 does not depict all Fact tables from which `Job_MaintainGIM` purges old data. For a list of the Fact tables in the dimensional model that `Job_MaintainGIM` purges, see the appendix about Info Mart tables purged by the maintenance job in the *Genesys Info Mart 8.0 Deployment Guide*.
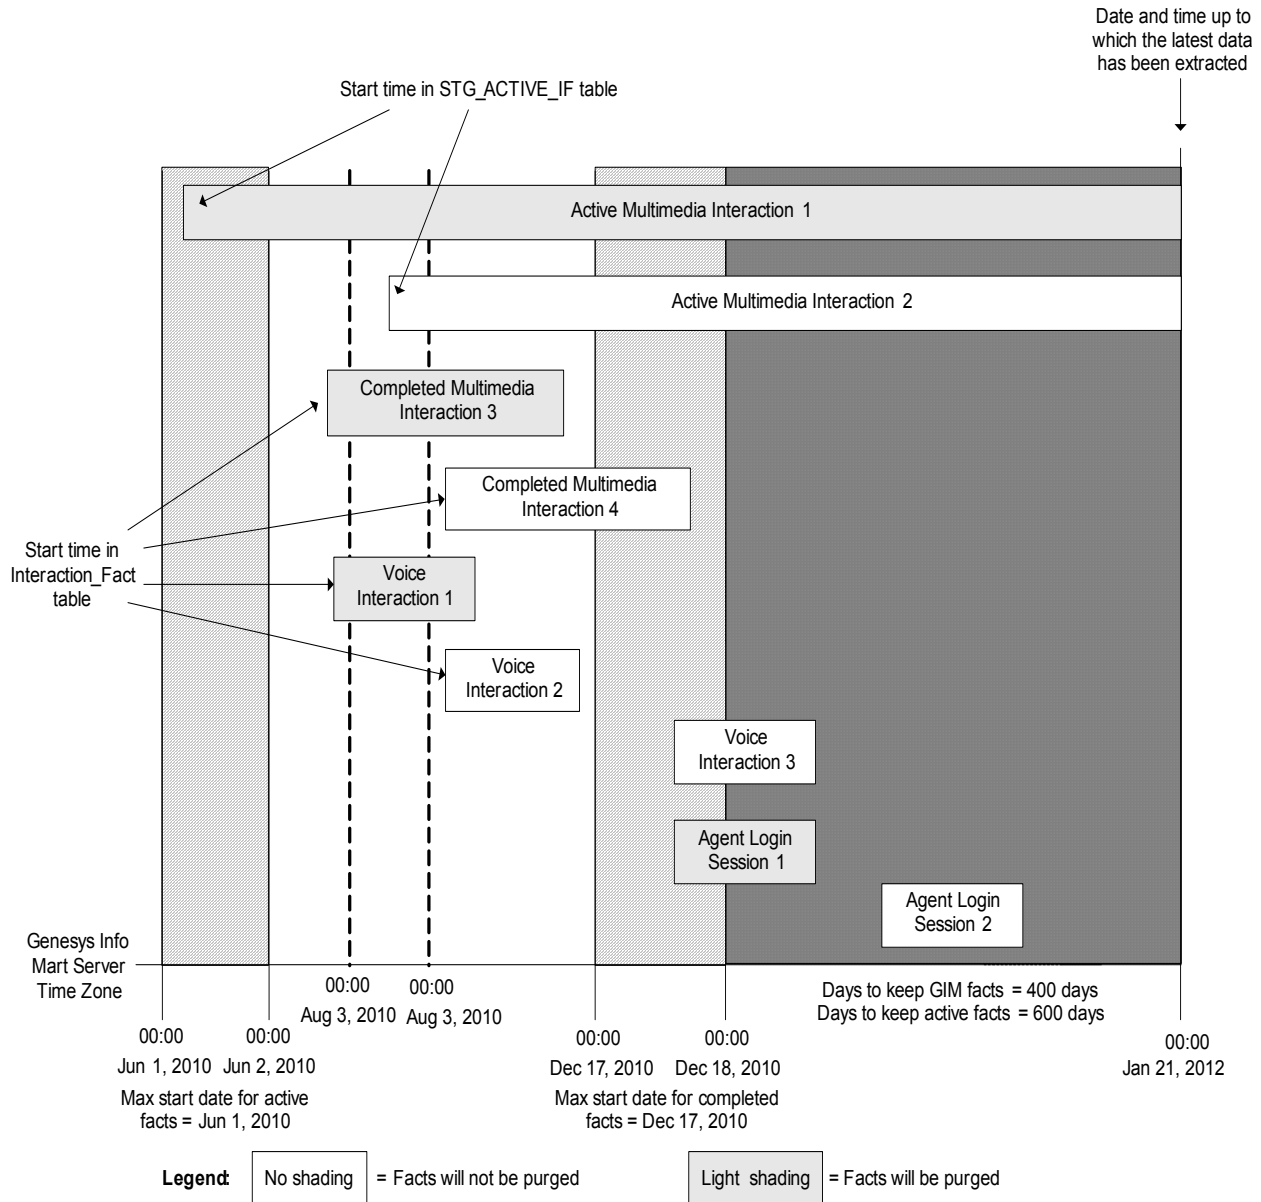
Figure 8: Purging Eligibility Example—Dimensional Model

Purging Fact Data from GIDB

Figure 9 on page 68 shows a purging eligibility example for Fact tables in GIDB. When looking at the example, keep in mind the following:

- The `days-to-keep-gidb-facts` is configured as 14 days.
- The `days-to-keep-active-facts` is configured as 600 days.
- Data has been extracted until January 21, 2012.

`Job_MaintainGIM` determines the maximum start dates for purge eligibility by the following actions:

- `Job_MaintainGIM` subtracts 14 days from January 21, 2012, which is January 7, 2012. This is the retention timestamp for short-living entities, and the upper bound of the purge eligibility interval for completed long-living interactions.

- `Job_MaintainGIM` subtracts 600 days from January 21, 2012, which is June 1, 2010. This is the retention timestamp for active long-living entities.

- From the root table in the group of related tables for the particular type of data, `Job_MaintainGIM` calculates the start time of the earliest active fact in the time range between the `days-to-keep-gidb-facts` and the `days-to-keep-active-facts` timestamps. For example, for multimedia interactions, the root table is `STG_ACTIVE_IF`.

    - The start time of the earliest active fact between January 7, 2012 and June 1, 2010 sets the purge threshold for completed long-living entities in the group of interaction-related tables. In the example shown in Figure 8, this date is August 3, 2010.

        In all related tables, completed facts that started on August 3, 2010 or earlier are eligible to be purged.

    - Active facts that started earlier than June 1, 2010 are also eligible to be purged.

In Figure 9, `Job_MaintainGIM` purges facts relating to the following entities:

Long-Living Entities Purged
- `Active Multimedia Interaction 1`, because the start date is earlier than the retention timestamp for active facts.

- `Completed Multimedia Interaction 3`, because the start date is earlier than the calculated purge threshold for completed long-living entities.

- `Outbound Contact Campaign 1`, because, for the group of Outbound Contact–related tables, there are no active facts within the time range between the `days-to-keep-gidb-facts` and the `days-to-keep-active-facts` timestamps, and the start date is earlier than the `days-to-keep-gidb-facts` timestamp.

Short-Living Entities Purged
- `Voice Interaction 1`, `Voice Interaction 2`, and `Voice Interaction 3` (by definition, completed facts), because the start date is earlier than the calculated purge threshold for completed short-living entities.

- `Agent Login Session 1`, because the start date is earlier than the retention timestamp for short-living entities.

`Job_MaintainGIM` does not purge facts relating to the following entities:

Long-Living Entities Not Purged
- `Active Multimedia Interaction 2`, because the start date is not earlier than the retention timestamp for active facts. However, as the earliest active interaction within the interval defined by the retention policy options for active and completed facts, `Active Multimedia Interaction 2` sets the purge threshold for completed long-living entities.

- `Completed Multimedia Interaction 4`, because the start date is later than the calculated purge threshold for completed long-living entities.

- `Outbound Contact Campaign 2`, because it is still active and the start date is not earlier than the `days-to-keep-gim-facts` and `days-to-keep-gidb-facts` timestamps.

Short-Living Entities Not Purged

- `Voice Interaction 4` because the start date is later than the calculated purge threshold for completed short-living entities.

- `Agent Login Session 2`, because the start date is later than the retention timestamp for short-living entities.

Note: Figure 9 does not depict all Fact tables from which `Job_MaintainGIM` purges old data. For a list of the Fact tables in GIDB that `Job_MaintainGIM` purges, see the appendix about Info Mart tables purged by the maintenance job in the *Genesys Info Mart 8.0 Deployment Guide*.
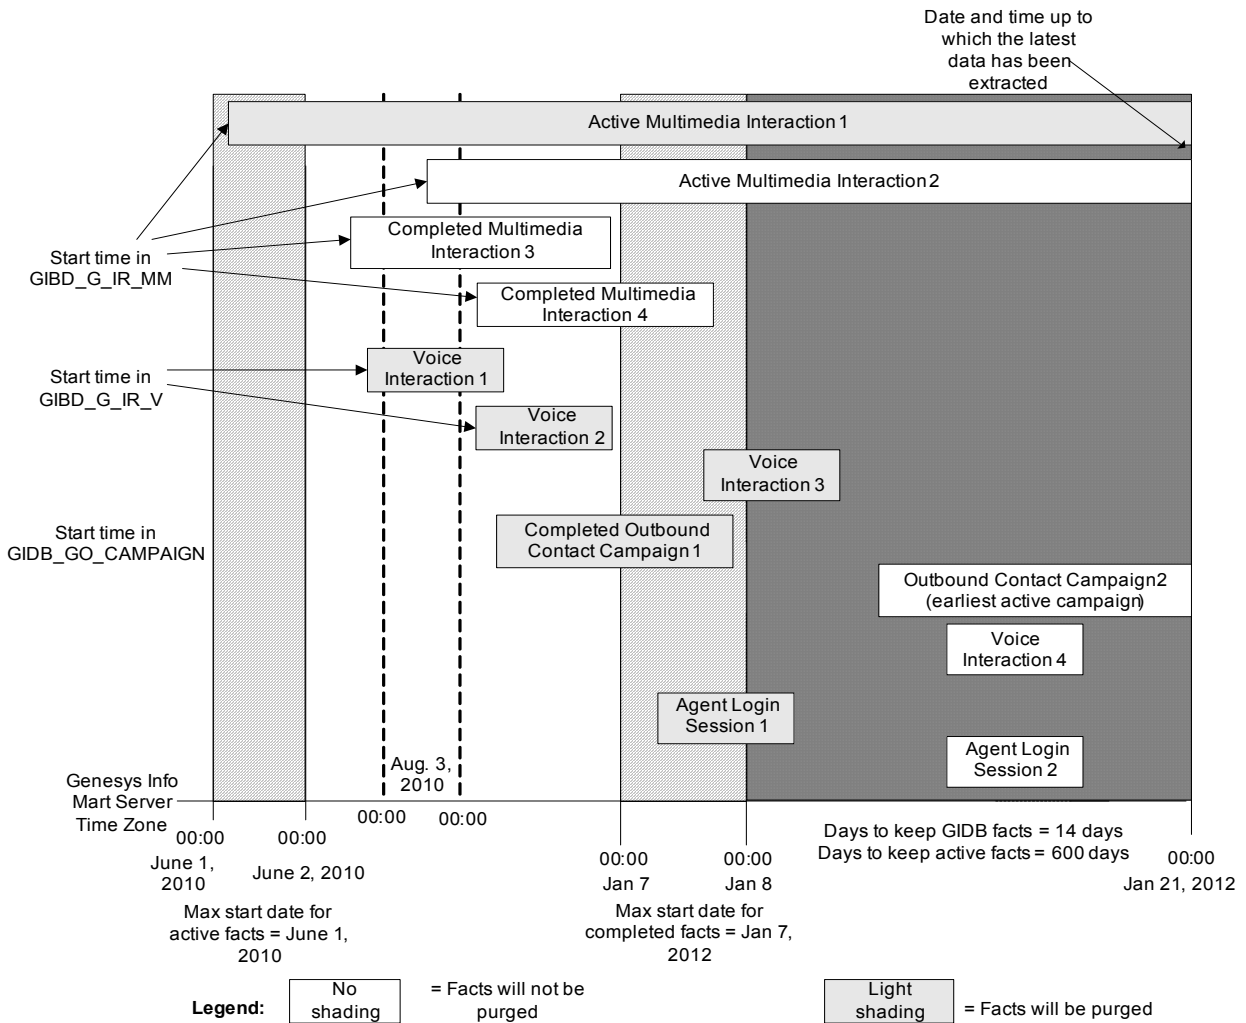
Figure 9: Purging Eligibility Example—GIDB

About Voice of Process

Voice of process functionality provides administrators with information related to the processing history of Genesys Info Mart jobs. In particular, voice of process functionality enables administrators to:

- Quickly check the state of Genesys Info Mart jobs.

- Track the data extraction progress of each extract job cycle.

- Track the data transformation progress of each transform cycle.

- Identify the job that inserted a particular row of data or made the most recent update to it.

The Genesys Info Mart database provides several service control tables, which in addition to existing administrative views, store the ETL processing history details.

Checking the State of ETL Processing

The `ADMIN_ETL_JOB_HISTORY` administrative view is updated each time a job is executed. By monitoring this view, administrators can quickly assess the current state of all jobs.

The `ADMIN_ETL_JOB_HISTORY` administrative view provides the following information related to the jobs:

• Name of the job

• Time of execution

• Time of completion

• Duration

• Status (success or failure)

Tracking the Progress of Data Extraction

A row is added to the `ADMIN_EXTRACT_HISTORY` administrative view when `Job_ExtractICON` successfully completes extracting a source data table. Administrators can closely track the progress of the data extract cycle by monitoring this view.

The `ADMIN_EXTRACT_HISTORY` provides the following information related to the data extract job including:

• Name of the source table

• Number of records extracted

• Start and end time of each extract job cycle

Tracking the Progress of Data Transformation

Administrators can closely track the progress of the data transform cycle by viewing the contents of the `CTL_TRANSFORM_HISTORY` table.

The `CTL_TRANSFORM_HISTORY` provides the following information related to the data transform job:

• Name of the destination table

• Number of records transformed

• Start and end time of each transform job cycle

Identifying the Job Inserting or Updating Data

In Info Mart's dimensional model, every table that receives inserts has a `CREATE_AUDIT_KEY` service field. Every table that can also receive updates has an additional `UPDATE_AUDIT_KEY` service field. Both these fields contain a reference to a row in the `CTL_AUDIT_LOG` table.

By linking fact data to the `CTL_AUDIT_LOG` table, administrators can determine:

*   The identifier of the job that inserted/updated the data. This information can be used to quickly locate relevant portions of the log file when troubleshooting data quality issues.

*   The total number of records processed by the job.

*   The range of partitions (when using partitioning) that received the new/updated data.

Additionally, fields `CREATE_AUDIT_KEY` and `UPDATE_AUDIT_KEY` can be used to identify newly arriving data for subsequent aggregation or other processing.

For more information about any of the Info Mart tables, views, and fields, see the *Genesys Info Mart 8.0 Reference Manual* for your RDBMS.

4 Troubleshooting Genesys Info Mart Jobs

This chapter describes errors that Genesys Info Mart jobs encounter. Some of the errors cause jobs to fail, while others result in incorrectly processed data.

The jobs that Genesys Info Mart Server launches, or that you execute or schedule, may encounter errors that cause them to fail. Some job failures are caused by an error that the job encounters directly. These are called *single-job failures*. Because of the job interdependencies that are described in "Job Sequencing Rules" on page 47, some job failures are caused by an error that some other job encounters. These are called *job-interdependency job failures*.

When you use Genesys Info Mart Server to launch jobs, you likely will not see job-interdependency job failures, because Genesys Info Mart Server will not launch a job until all interdependent jobs have completed successfully.

If the jobs are run automatically by scheduler, jobs are recovered automatically in cases of failures caused by environment situations. For example, if a database shutdown occurs during transformation, the transform job restarts from the point it left off when the database is restarted.

This chapter also provides recommendations for situations in which the extract, transform, and load (ETL) cycle has not run for an extended period of time. These recommendations help you process the backlog of source data in a way that leads to the best ETL performance.

This chapter contains the following sections:

Types of Errors

Several types of errors can cause a job to fail or to produce incomplete data:

Configuration Errors

If Genesys Info Mart Server reports a configuration error, refer to the log messages to identify the issue so that you can correct it.

For additional information on the configuration checking process, including information on when Genesys Info Mart performs the check, what it checks for, and how it responds when it finds errors, see the section about deployment verification in the "Data Processing" chapter of the *Genesys Info Mart 8.0 Deployment Guide*.

In addition to the items Configuration Checker reviews, other possible configuration errors include the following:

- The transform job (`Job_TransformGIM`) encountered errors in the configuration of either OCS record fields in Configuration Database or Outbound Contact–related user data fields in the Info Mart database. In either case, the job does not process the incorrectly configured fields and/or user data. `Job_TransformGIM` depends on this configuration for Outbound Contact data transformation. Although the transform job does not fail when these errors are encountered, the resulting data in the Info Mart database would be missing certain information about Outbound Contact–related interactions.

- `Job_TransformGIM` encountered missing configuration data during transformation of Configuration details or data from other domains. For more information about how `Job_TransformGIM` handles configuration data errors, see "Handling Configuration Data Errors" on page 34.

You should monitor the Genesys Info Mart local log for messages indicating configuration errors and take action to resolve them promptly to avoid incomplete data. Genesys recommends setting up an alarm condition using the Solution Control Interface as this error can occur at any time during report configuration when `Field` Annex options are changed. Refer to the *Framework 8.0 Solution Control Interface Help* for information about how to create alarm conditions using the Alarm Condition Wizard.

Database Connection Errors

- The extraction job could not connect to a source database from which it extracts data because the database is not running or the network connection between the Genesys Info Mart Server and the database is down.

- The Genesys Info Mart job could not connect to a target database because the database is not running or the network connection between the Genesys Info Mart Server and the database is down.

- The Genesys Info Mart job or Genesys Info Mart Server could not connect to a source or target database because of a `JDBC Driver class not found:` exception. Ensure that the `CLASSPATH` environment variable has been updated to include the JDBC-specific jar files needed for the appropriate database type. Restart the Genesys Info Mart Server after the classpath is updated.

- In a High Availability (HA) configuration, `Job_ExtractICON` could not connect to one of the HA Interaction Databases (IDBs) because the database is not running or the network connection between the Genesys Info Mart Server and the database is down.

At the start of each extract job, the Genesys Info Mart Server automatically checks connectivity for each database access point (DAP) to which the Genesys Info Mart `Application` has a configured connection. If it determines that a configured data source or IDB is not available, Genesys Info Mart logs an error and does not proceed with the extract.

You should always check the logs to obtain additional detailed information about any possible connection errors.

SQL Error

- The job encountered a Structured Query Language (SQL) error that caused the failure. For example, there may be insufficient resources, such as memory or physical storage, on the database.

- As a special case in deployments that use 3rd Party Media, the transformation job might encounter a unique constraint violation error if you were trying to add media types to the `MEDIA_TYPE` dimension table at the exact moment that the transformation job was dynamically adding an unknown media type to the `MEDIA_TYPE` dimension.

  If the transformation job fails for this reason, no action is required. During the next ETL cycle, the transformation job will take the appropriate action with regard to the interaction and the associated media type.

  For more information, see the section about preparing the Info Mart database for online interactions in the *Genesys Info Mart 8.0 Deployment Guide*.

- Genesys Info Mart is not able to access a Microsoft SQL database. In order for the connection to be made, Microsoft SQL JDBC driver version 1.1 or higher must be installed.

Genesys Info Mart Job Error

- The job encountered a critical error that caused the failure. For example, there may be insufficient operating system resources or a software defect.

- Genesys Info Mart determined that data from a configured data source is not available. Genesys Info Mart logs an error, and `Job_ExtractICON` does not proceed in this situation. For more information, see "Determining Data Availability" on page 30.

- `Job_TransformGIM` has been configured to fail when errors are encountered, and the job encountered a data inconsistency for which the applicable interaction-level error policy generated an exception. For more information about the configurable error policy options and the situations that might give rise to data inconsistency errors, see "Error Handling" on page 34.

Resources to Consult for Additional Information

Consult the following resources for information that will help you resolve problems:

- *Genesys Info Mart 8.0 Deployment Guide*—Contains information to help you tune performance parameters for your Info Mart database and correct errors in Genesys Info Mart configuration parameters. This guide also contains information that you can use to configure the ICON applications that populate the databases from which you want to extract data.

- *Genesys Info Mart 8.0 Operations Guide* (this document)—Contains information that you can use to correct errors in the job configuration. This guide also contains information about how to execute or schedule jobs and job interdependencies.

- Genesys Central Logger—Contains events that are logged by jobs. The logs indicate configuration errors, when a job begins, when a job ends, and whether it ends successfully or unsuccessfully. When a job fails, use one of the following methods to obtain detailed information about the failure.

  - Use a log file in the Genesys Info Mart Server's local directory to view log messages.
  - Use SCI to view log messages that are received by Message Server, provided that the Genesys Info Mart application has been configured with a connection to the Genesys Message Server.

- Genesys Info Mart local log—Contains detailed events that are logged to the local log file on the Genesys Info Mart Server host by Genesys Info Mart Server and some of the jobs. When a job fails, view these logs to obtain detailed information about the failure.

- Voice of Process—Provides administrators with information related to the processing history of jobs. You can use this information to quickly check the state of all jobs, to track the data extraction or data transformation progress of each data extract or data transform job cycle, to detect new and/or updated data for subsequent aggregation or other processing, and to diagnose ETL problems. See "About Voice of Process" on page 68 for more information.

- Publications for your database—Contain information for your specific RDBMS about database connections, SQL errors, configuration parameter settings that affect database performance, and the usage of operating system resources on the database server.

Error Recovery

In the following sections, Genesys provides several recommendations to consider when Genesys Info Mart jobs fail:

General Recommendations

When the Genesys Central Logger or Genesys Info Mart Administration Console indicates that a job failed, the cause of the failure dictates the recovery steps. Messages in the Genesys Info Mart local log or Genesys Central Logger will indicate the error that caused the failure. Correct the cause of the failure before attempting to restart any job.

If a job continues to fail and it takes a long time to resolve the issue, follow the suggestions provided in "Recovering from a Prolonged ETL Outage" on page 75 for the time period that the ETL is not running.

Recovering from a Prolonged ETL Outage

If certain circumstances (such as a failure of a particular job) prevent you from running ETL for an extended period of time, no special steps are required to process the backlog when normal processing resumes. However, carefully review settings for the options that control transaction size and limit data extraction. They might be set too large for the situation in which you are running normal ETL cycles to catch up a large backlog.

Recovering from Data Source Unavailability

The extract algorithm will fail the extract job when data from a particular data source is not available. Alternatively, you might be alerted by an alarm on log message 55-20110 that the transform job is experiencing delays. (See "Delayed Data" on for information about the conditions under which the log message is generated.)

If `Job_ExtractICON` fails or `Job_TransformGIM` is held up because of missing data:

1. Examine the Genesys Info Mart logs to determine the reason for the failure and identify the data source whose data is no longer available.

2. Evaluate whether it is worth continuing to extract data from remaining data sources under the risk of permanently losing data from the unavailable data source.

   This can be the case when, for example, Genesys Info Mart connection to a non-HA IDB experiences frequent connectivity problems for a certain period of time.

3. If you make the decision to proceed with the ETL process, temporarily disable the data source, by clearing the `State Enabled` check box on the `General` tab of the data source `Application` object. Disabling the data source, such as T-Server or Interaction Server, means that Genesys Info Mart will no longer expect to receive any data from it.

   • The configuration check will no longer identify the data source as one that is supposed to be available, so `Job_ExtractICON` will proceed to extract data from available data sources.

   • `Job_TransformGIM` will not delay transformation of data from other data sources in the same data domain.

Warning! Extracting data from a limited number of data sources impacts the data quality in the Info Mart database and, therefore, undermines the accuracy of contact center reports. To ensure that the accuracy of reports is not compromised on a permanent basis, reset the `State Enabled` check box on the data source `Application` object as soon as the data source availability issue is resolved.

High Availability Recommendations

In a high availability (HA) deployment, the IDB redundancy provides an additional layer of failure prevention for the extract job. Yet, you may experience a failure of `Job_ExtractICON` in a situation when it cannot access any IDBs in a redundant set. In this case, follow the guidelines in "Recovering from Data Source Unavailability" to decide whether you want the extract job

to ignore the missing data sources and, hence, compromise the completeness or accuracy of the reporting data.

Purging IDB

The size of Interaction Database (IDB) is one of the important factors that affects Interaction Concentrator (ICON) and Genesys Info Mart operational performance. You should periodically purge old data from the IDBs that Genesys Info Mart uses as sources of data.

This appendix provides recommendations and considerations relevant to Genesys Info Mart data source requirements, to prevent accidental purging of IDB data before Genesys Info Mart has an opportunity to extract it. In particular, this appendix provides guidance for selecting the smallest safe value to specify for retaining IDB data.

This appendix provides information about the following:

- Interaction Concentrator Purge Procedures, page 79
- IDB Data Retention, page 80

For detailed information about purging IDB, see the chapter about using special stored procedures in the *Interaction Concentrator 8.0 User's Guide*.

For more information about purging the Info Mart database, see "Purging the Genesys Info Mart Database" on page 62.

Interaction Concentrator Purge Procedures

Genesys Info Mart does not provide automated purging of old IDB data. However, various releases of Interaction Concentrator provide several stored procedures that purge old IDB data. Use your RDBMS utility, or write your own program or stored procedure, to invoke the Interaction Concentrator functionality to purge IDB data in a way that:

- Avoids deleting data that Genesys Info Mart has not yet extracted.
- Retains enough historical data to allow for error recovery and problem determination.

IDB Purge Frequency

Genesys Info Mart recommends that you run the Interaction Concentrator stored procedure(s) to purge old IDB data once a day, during off-peak hours, when contact center activity is low, and when Genesys Info Mart is not accessing IDB. This means you should run the stored procedures at the same time that the Genesys Info Mart runs its daily maintenance job, `Job_MaintainGIM`. Interaction Concentrator stored procedures may take some time to finish, so run them as early as possible to allow them to complete before Genesys Info Mart starts the next extract, transform, and load (ETL) cycle.

IDB Data Retention

The amount of historical data you are able to retain in IDB depends on the database server's hardware resources, such as memory and disk space, and disk subsystem performance. Because Genesys Info Mart initially copies almost all data from IDB into Global Interaction Database (GIDB) tables in the Info Mart schema, it is not necessary to retain IDB data for long periods. Nevertheless, for prudence, you should retain as much IDB data as possible, without impacting either ICON's operational performance or Genesys Info Mart's data extraction performance.

The following scenarios require special consideration when determining the number of days to retain IDB data:

*   A new Genesys Info Mart deployment—In a new Genesys Info Mart deployment, where ICON stores data prior to Genesys Info Mart's first ETL cycle, there is a backlog of IDB data waiting to be processed. By default, Genesys Info Mart limits the amount of data extracted during each ETL cycle (refer to the description of the `extract-data-chunk-size` configuration option on) while it works through a backlog of data. You might need to temporarily increase the IDB data retention period to take this backlog into account. You may also choose to suspend purging IDB data until Genesys Info Mart has had an opportunity to extract the backlog of data.

*   ETL failure—If some network, hardware, or software outage occurs that prevents Genesys Info Mart from maintaining its regular ETL schedule, consider temporarily increasing the IDB data retention period to account for the time that Genesys Info Mart has not been able to extract IDB data. You may also choose to suspend purging IDB data until Genesys Info Mart has had an opportunity to extract the backlog of data.

*   IDB data archiving—If your environment requires long-term storage of IDB data (longer than ICON's operational performance or Genesys Info Mart's data extraction performance permits), consider archiving IDB data. This allows the operational data store used by ICON and Genesys Info

Mart to be small enough to allow acceptable and predictable performance, while providing an alternate data store for long-term archiving of IDB data. Work with your Database Administrator to determine an appropriate archival strategy.

Determining the Retention Period

The sample procedure documented below gives an estimate for the data retention threshold for IDB data. This procedure takes into account the last time that Genesys Info Mart successfully extracted all of the data for an ETL cycle, and whether Genesys Info Mart is limiting the amount of data it extracts while it processes a backlog of data. This procedure also includes a safety buffer of seven days.

In an environment where Genesys Info Mart is maintaining a regular schedule of ETL cycles that run every sixty minutes or less, Genesys recommends that the number of days to retain IDB data is between 7 and 30 days. The procedure documented below will never return fewer than 7 days. If you are able to store more than 7 days of IDB data, you may choose to use a value larger than what is returned.

Note: If the Genesys Info Mart ETL cycles have not yet begun, this procedure returns a large value to prevent the accidental purging of data that has not yet been extracted.

Run an RDBMS-specific SQL statement against your Info Mart database to return the minimum value for the data retention threshold you will provide as an input parameter for the Interaction Concentrator purge procedure. To ensure an accurate calculation, issue the SQL statements prior to running each purge procedure. Also, make sure to log in using the Genesys Info Mart database Owner account before issuing the statements.

Example SQL Query

The following SQL query shows how to run a query against a Microsoft SQL Server–based Info Mart database. If you are using a different RDBMS, convert this query as appropriate.

Note: If you have performed only one ETL cycle, this query returns the result 999. However, if you have run more than one ETL cycle, it functions correctly.

```
select cast(max(x) as decimal) as MINIMUM_DAYS from (
/* number of days since 1st extract preceding last completed transformation */
/* (and following the prior completed transformation) or 999, if none */
   select coalesce((
```

```
                    select round(datediff(second, min(start_time),
                               CURRENT_TIMESTAMP)/86400.0, 0) + 7
                 from admin_etl_job_history
                 where job_name like '%Extract%'
                 and start_time <
                     (select max(start_time)
                         from admin_etl_job_history
                         where job_name = 'Job_TransformGIM' and status = 'COMPLETE')
                 and start_time >
                     (select max(start_time)
                         from admin_etl_job_history
                         where job_name = 'Job_TransformGIM' and status = 'COMPLETE'
                         and start_time <
                             (select max(start_time)
                                 from admin_etl_job_history
                                 where job_name = 'Job_TransformGIM'
                                  and status = 'COMPLETE'))
            ), 999) as x

        ) z;
```

Related Documentation Resources

The following resources provide additional information that is relevant to this software. Consult these additional resources, as necessary.

Framework

- The *Framework 8.0 Management Layer User's Guide* provides information about the concepts, terminology, and procedures that apply to this layer of the Genesys Framework.

- The *Framework 8.0 Configuration Options Reference Manual* provides information about configuration options for Framework components.

- The *Framework 8.0 Configuration Manager Help* provides information about using Configuration Manager in either an enterprise or a multi-tenant environment.

- The *Framework 8.0 Deployment Guide* provides information about configuring, installing, starting, and stopping Framework components.

- The *Framework 8.0 Combined Log Events Help* describes log events that Genesys server applications generate and that Solution Control Interface displays. The *Framework 8.0 Combined Log Events Help* includes descriptions of Genesys Info Mart log events.

Interaction Concentrator

- The *Interaction Concentrator 8.0 Deployment Guide* provides information about architecture, configuration requirements, and installation steps for Interaction Concentrator, and it describes how to make data from the Genesys Outbound Contact solution available in Interaction Database (IDB).

- The *Interaction Concentrator 8.0 User's Guide* provides basic information about IDB architecture and detailed information about Interaction Concentrator features and functionality, including attached data processing, available stored procedures, and integration with other Genesys products.

- The *Interaction Concentrator 8.0 Physical Data Model* for your relational database management system (RDBMS) provides information about the IDB schemas.

Genesys Info Mart

- The *Genesys Info Mart 8.0 Deployment Guide* provides information about architecture, configuration requirements, and installation steps for Genesys Info Mart and the Genesys Info Mart Administration Console.

- The *Genesys Info Mart 8.0 Operations Guide* provides information about the Genesys Info Mart jobs. The jobs extract, transform, and load (ETL) data, maintain the Info Mart database, and migrate the database schema as required. The guide also explains how to use the Genesys Info Mart Administration Console to monitor and administer the jobs and how to purge the Info Mart database.

- The *Genesys Info Mart 8.0 User's Guide* provides information about how to use data that is stored by Genesys Info Mart for contact center historical reporting.

- The *Genesys Info Mart 8.0 Reference Manual* for your RDBMS provides information about the Info Mart database schema.

- The *Genesys Info Mart 8.0 Database Size Estimator* helps you estimate the size of your Info Mart database when you are planning your deployment. The estimator is a Microsoft Office Excel 2007 spreadsheet that is available from the Genesys Technical Support website.

- Release Notes and Product Advisories for this product, which are available on the Genesys Technical Support website at
  http://genesyslab.com/support.

Reporting and Analytics Aggregates

- The *Reporting and Analytics Aggregates 8.0 Deployment Guide* describes how to deploy the Reporting and Analytics Aggregates (RAA) package provided with Genesys Info Mart.

- The *Reporting and Analytics Aggregates 8.0 Reference Manual* describes the aggregate tables that are available to Genesys Info Mart customers with deployment of RAA.

- The *Reporting and Analytics Aggregates 8.0 User's Guide* describes the aggregation process, provides the aggregation hierarchy, and explains how to enable aggregation of user data.

Genesys Interactive Insights

- The *Genesys Interactive Insights 8.0 Deployment Guide* describes how to install Genesys Interactive Insights (GI2) and set up the environment required in order to run the GI2 reports.

- The *Genesys Interactive Insights 8.0 Universe Guide* describes, in detail, the reports and measures that are provided in the GI2 release.

- The *Genesys Interactive Insights 8.0 User's Guide* summarizes how to operate GI2 reports and provides basic instructions for customizing your own reports.

Genesys

- The *Genesys Technical Publications Glossary,* which ships on the Genesys Documentation Library DVD, provides a comprehensive list of the Genesys and computer-telephony integration (CTI) terminology and acronyms that are used in this document.

Information about supported hardware and third-party software is available on the Genesys Technical Support website in the following documents:

- *Genesys Supported Operating Environment Reference Manual*

- *Genesys Supported Media Interfaces Reference Manual*

Consult the following additional resources as necessary:

- The *Genesys Migration Guide*, which ships on the Genesys Documentation Library DVD, provides documented migration strategies for Genesys product releases. Contact Genesys Technical Support for more information.

- The *Genesys Hardware Sizing Guide* provides information about Genesys hardware sizing guidelines for the Genesys 8.x releases.

- The *Genesys Interoperability Guide* provides information on the compatibility of Genesys products with various Configuration Layer Environments; Interoperability of Reporting Templates and Solutions; and *Gplus* Adapters Interoperability.

- The *Genesys Licensing Guide* introduces you to the concepts, terminology, and procedures that are relevant to the Genesys licensing system.

- The *Genesys Database Sizing Estimator 8.0 Worksheets* provides a range of expected database sizes for various Genesys products.

For additional system-wide planning tools and information, see the release-specific listings of System Level Documents on the Genesys Technical Support website. These documents are accessible from the `system level documents by release` tab in the Knowledge Base `Browse Documents` Section.

Genesys product documentation is available on the:

- Genesys Technical Support website at `http://genesyslab.com/support`.

- Genesys Documentation Library DVD, which you can order by e-mail from Genesys Order Management at `orderman@genesyslab.com`.

Document Conventions

This document uses certain stylistic and typographical conventions—introduced here—that serve as shorthands for particular kinds of information.

Document Version Number

A version number appears at the bottom of the inside front cover of this document. Version numbers change as new information is added to this document. Here is a sample version number:

80gim_dep_09-2010_v8.0.001.00

You will need this number when you are talking with Genesys Technical Support about this product.

Screen Captures Used in This Document

Screen captures from the product graphical user interface (GUI), as used in this document, may sometimes contain minor spelling, capitalization, or grammatical errors. The text accompanying and explaining the screen captures corrects such errors *except* when such a correction would prevent you from installing, configuring, or successfully using the product. For example, if the name of an option contains a usage error, the name would be presented exactly as it appears in the product GUI; the error would not be corrected in any accompanying text.

Type Styles

Table 3 describes and illustrates the type conventions that are used in this document.

Table 3: Type Styles

Type Style	Used For	Examples
Italic	• Document titles • Emphasis • Definitions of (or first references to) unfamiliar terms • Mathematical variables Also used to indicate placeholder text within code samples or commands, in the special case where angle brackets are a required part of the syntax (see the note about angle brackets on page 88).	Please consult the *Genesys Migration Guide* for more information. Do *not* use this value for this option. A *customary and usual* practice is one that is widely accepted and used within a particular industry or profession. The formula, $x + 1 = 7$ where x stands for . . .
Monospace font (Looks like `teletype` or `typewriter text`)	All programming identifiers and GUI elements. This convention includes: • The *names* of directories, files, folders, configuration objects, paths, scripts, dialog boxes, options, fields, text and list boxes, operational modes, all buttons (including radio buttons), check boxes, commands, tabs, CTI events, and error messages. • The values of options. • Logical arguments and command syntax. • Code samples. Also used for any text that users must manually enter during a configuration or installation procedure, or on a command line.	Select the `Show variables on screen` check box. In the `Operand` text box, enter your formula. Click `OK` to exit the `Properties` dialog box. T-Server distributes the error messages in `EventError` events. If you select `true` for the `inbound-bsns-calls` option, all established inbound calls on a local agent are considered business calls. Enter `exit` on the command line.
Square brackets ([ ])	A particular parameter or value that is optional within a logical argument, a command, or some programming syntax. That is, the presence of the parameter or value is not required to resolve the argument, command, or block of code. The user decides whether to include this optional information.	`smcp_server -host [/flags]`
Angle brackets (< >)	A placeholder for a value that the user must specify. This might be a DN or a port number specific to your enterprise. **Note:** In some cases, angle brackets are required characters in code syntax (for example, in XML schemas). In these cases, italic text is used for placeholder values.	`smcp_server -host <confighost>`

Index

Symbols

Numerics

A

B

C

D