Genesys Info Mart 7.6

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 following regional numbers. For complete contact information and procedures, refer to the ***Genesys Technical Support Guide***.

Region	Telephone	E-Mail
North and Latin America	+888-369-5555 or +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
Japan	+81-3-6361-8950	support@genesyslab.co.jp

Prior to contacting technical support, please refer to the *Genesys Technical Support Guide* for complete contact information and procedures.

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: 76gim_op_03-2011_v7.6.006.00

Table of Contents

Preface

Welcome to the *Genesys Info Mart 7.6 Operations Guide*. Genesys Info Mart produces a data mart containing several star schemas that you can use for contact center historical reporting.

Genesys Info Mart includes a software platform and a set of predefined jobs. You configure these jobs to extract and transform data from several Genesys relational databases. The transformed data is loaded into dimension and fact database tables in Genesys Info Mart. You can query the data in these tables, using SQL, to display detailed data, reveal patterns, and predict trends.

This guide describes the procedures that you must follow to schedule and monitor the Genesys Info Mart Extraction, Transformation, and Loading (ETL) jobs. It is intended for system administrators and is valid only for Genesys Info Mart release 7.6.

Note: For versions of this document created for other releases of this product, please 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 chapter includes the following sections:

In brief, you will find the following information in this guide:

- Descriptions of Genesys Info Mart ETL jobs
- Instructions for scheduling and monitoring ETL jobs
- Tips for troubleshooting ETL jobs

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

Recommended Reading

Genesys Info Mart is designed to run under Genesys products. Genesys strongly recommends that you read the following documentation before you run any Genesys Info Mart ETL jobs:

- *Framework 7.6 Configuration Manager Help*
- *Framework 7.6 Stat Server User's Guide*
- *Interaction Concentrator Deployment Guide*
- *Interaction Concentrator User's Guide*
- *Genesys Info Mart 7.6 Deployment Guide*
- *Genesys Voice Platform 7.6 Voice Application Reporter Deployment and Reference Manual*

Chapter Summaries

In addition to this preface, this guide contains the following chapters:

- Chapter 1, "Introducing Genesys Info Mart 7.6," on page 17, introduces basic functionality and new features for this release.
- Chapter 2, "Understanding Genesys Info Mart ETL Jobs," on page 27, discusses the ETL Jobs and what they do.
- Chapter 3, "Working with ETL Jobs," on page 65, provides information on how to use the ETL jobs shipped with Genesys Info Mart.

- Chapter 4, “Managing Data Sources,” on page 105, provides information on how to migrate, move and purge data source databases.
- Chapter 5, “Troubleshooting ETL Jobs,” on page 123, provides information on troubleshooting ETL jobs.
- The appendix, “Using Stat Server in Legacy Environments” on page 135, provides information related to extracting and purging data from Stat Server in a legacy Reporting environment.

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:

76gim_op_03-2011_v76.001.00

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

Type Styles

Italic

In this document, italic is used for emphasis, for documents’ titles, for definitions of (or first references to) unfamiliar terms, and for mathematical variables.

Examples:

- Please consult the *Genesys Migration Guide* for more information.
- *A customary and usual practice* is one that is widely accepted and used within a particular industry or profession.
- Do *not* use this value for this option.
- The formula, $x + 1 = 7$ where x stands for . . .

Monospace Font

A monospace font, which looks like `teletype or typewriter text`, is used for all programming identifiers and GUI (graphical user interface) 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; and code samples.

Examples:

- Select the `Show variables on screen` check box.
- Click the `Summation` button.
- In the `Properties` dialog box, enter the value for the host server in your environment.
- In the `Operand` text box, enter your formula.
- Click `OK` to exit the `Properties` dialog box.
- The following table presents the complete set of error messages T-Server® distributes 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.

Monospace is also used for any text that users must manually enter during a configuration or installation procedure, or on a command line:

Example:

- Enter `exit` on the command line.

Screen Captures Used in This Document

Screen captures from the product GUI, as used in this document, may sometimes contain a minor spelling, capitalization, or grammatical error. 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.

Square Brackets

Square brackets indicate that a particular parameter or value is optional within a logical argument, a command, or some programming syntax. That is, the parameter's or value's presence is not required to resolve the argument, command, or block of code. The user decides whether to include this optional information. Here is a sample:

```
smcp_server -host [/flags]
```

Angle Brackets

Angle brackets indicate a placeholder for a value that the user must specify. This might be a DN or port number specific to your enterprise. Here is a sample:

```
smcp_server -host <confighost>
```

Related Resources

In addition to "Recommended Reading" on page 8, consult the following resources as necessary:

- The *Genesys Info Mart 7.6 User's Guide.*
- The *Genesys Info Mart 7.6 Reference Manual* for your database management system.
- The *Genesys Master Glossary*, which ships on the Genesys Documentation Library DVD, provides a list of Genesys and CTI terms and acronyms.
- The *Genesys Migration Guide*, also on the Genesys Documentation Library DVD, contains a documented migration strategy for each software release. Refer to the applicable portion or contact Genesys Technical Support for additional information.
- The *Genesys Interoperability Guide*, also on the Genesys Documentation Library DVD, provides a list of the versions of Genesys software that interoperate with the 7.6 release of Genesys Info Mart.
- The Release Notes and Product Advisories for this product, which are available on the Genesys Technical Support website.

Information on 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*

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.

Making Comments on This Document

If you especially like or dislike anything about this document, please 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 information in this document only and to the way in which the information is presented. Speak to 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.

Document Change History

This section lists topics that are new in the current release *of this document*, or that have changed significantly from the preceding release.

New in Document Version v7.6.006.00

The document has been updated to support Genesys Info Mart release 7.6.012. The following topics have been added or significantly changed since the previous 7.6 release of this document:

- "New in Release 7.6.012" on page 26—A new subsection in the "New in This Release" section describes the feature enhancements introduced in Genesys Info Mart release 7.6.012.

  "Job_LoadRecent" on page 43—Has been updated with situations requiring special handling in order to prevent a single execution of `Job_LoadRecent` from aggregating facts over a large time span, and in so doing, taking many hours to complete if your intraday and/or historical fact tables already have data in them.
- "Calculating Aggregates Within a Time Span" on page 53—This section has been updated with new procedures you should follow to prevent a single execution of `Job_LoadRecent` from aggregating all eligible fact data and how to use input parameters to specify the time span for which you want `Job_AggregateGIM` to calculate or recalculate the aggregate tables.
- "Purging Rules"on page 59—The bulleted item which starts with "GI2 out-of-box aggregates" and discusses how these aggregates are evaluated for purging based on the standard tenant date, has been updated to include outbound aggregates.
- "Moving Data Sources" on page 107—New information has been added to this section indicating that it is possible to move the data source and change the database connection information for the ICON data sources with the roles listed on page 107.

New in Document Version v7.6.005.00

The document has been updated to support Genesys Info Mart release 7.6.011. The following topics have been added or significantly changed since the previous 7.6 release of this document:

- "New in Release 7.6.009" on page 24—A new subsection in the "New in This Release" section describes the feature enhancements introduced in Genesys Info Mart release 7.6.009.
- "New in Release 7.6.011" on page 25—A new subsection in the "New in This Release" section describes the feature enhancements introduced in Genesys Info Mart release 7.6.011.

- "Restarting a Multimedia ICON" on page 38—A new subsection in the "Job_ExtractICON" section describes the special considerations and procedures that are needed in order to avoid data quality issues when stopping and restarting a Multimedia Interaction Concentrator. This small section points users to the newly added section,"Restarting a Multimedia ICON" on page 117 which contains the detailed information about restarting a Multimedia Interaction Concentrator.
- "Restarting a Multimedia ICON" on page 117—A new section describes the special considerations, the Genesys recommended Interaction Concentrator options, and the procedures that you need to follow to properly stop and restart a Multimedia Interaction Concentrator in order to avoid data quality issues within your Genesys Info Mart.

New in Document Version v7.6.004.00

The document has been updated to support Genesys Info Mart releases 7.6.007 and 7.6.008. The following topics have been added or significantly changed since the previous 7.6 release of this document:

- "New in Release 7.6.007" on page 23—A new subsection in the "New in This Release" section describes the feature enhancements introduced in Genesys Info Mart release 7.6.007.
- "New in Release 7.6.008" on page 24—A new subsection in the "New in This Release" section describes the feature enhancements introduced in Genesys Info Mart release 7.6.008.
- High availability (HA) of Outbound Contact details—The following sections have been updated with information describing how Genesys Info Mart now provides high availability data extraction for Outbound Contact details:
  - Throughout the document sections have been updated to indicate that HA data extraction of Outbound Contact details via switchover is now supported and this functionality requires Interaction Concentrator 8.0. See the *Genesys Info Mart 7.6.x Release Notes* for the minimum release of Interaction Concentrator 8.0 that is required to support this functionality.
  - Information has been added to the "Job_ExtractICON" section of the "ETL Job Summary Table" on page 28 describing HA of Outbound Contact details.
  - Information has been added to "ICON_OCS" on page 36 indicating that HA of Outbound Contact details is now supported. A new subsection "HA Data Extraction" on page 37 describes in detail how Genesys Info Mart processes HA of Outbound Contact details.
  - Table 9, "ADMIN_EXTRACT_HISTORY View Columns," on page 99 has been updated to include three new view columns which track data source sequences:
    - LAST_INSERT_SEQUENCE

    - LAST_UPDATE_SEQUENCE
    - LAST_MERGE_SEQUENCE
  - A new section, “Moving Data Sources” on page 107 has been added to Chapter 4, describing how to replace the IDB from which you extract Configuration details, and how to replace one IDB in an HA pair from which you extract Outbound Contact details.
  - A note has been added to“High Availability Recommendations” section on page 132 indicating that the Genesys recommendations described in this section are only valid for HA of Voice and Configuration details. They are not valid for HA of Outbound Contact details.
- Automatic Retry of Failed ETL Jobs—The following sections have been updated with information describing how Genesys Info Mart supports automated retries of failed ETL jobs when they are started by Genesys Info Mart Server scheduler:
  - A new section,“Automatically Retrying Failed ETL Jobs with Genesys Info Mart Server” on page 72 has been added to Chapter 3, describing the automated retry of failed ETL jobs and how the retry is handled differently when the failure is due to unresolved references to configuration objects.
  - Information about the configuration options, `job-retry-count` and `job-retry-wait`, which can be configured to work with this new functionality, has been added to “Configuring Automatic Retry of ETL Jobs” on page 72.
  - Information has been added to the “Handling Configuration Errors” section on page 42 describing how Genesys Info Mart uses the automatic retry of `Job_transformGIM`.
  - Information about the dependency between the automatic retry of `Job_TransformGIM` and the `ignore-missing-config-objs` configuration options has been added to “Handling Configuration Errors” on page 42.

New in Document Version v7.6.003.00

The document has been updated to support Genesys Info Mart releases 7.6.005 and 7.6.006. The following topics have been added or significantly changed since the previous 7.6 release of this document:

- “New in This Release” on page 18—New subsections in the “New in This Release” section describe the operations-related feature enhancements introduced in releases 7.6.005 and 7.6.006, respectively.
- Throughout the document, references to preceding Genesys Info Mart releases no longer use a maintenance release numbering scheme; the preceding maintenance releases are referred to as Genesys Info Mart release 7.6.004 and Genesys Info Mart 7.6.005, respectively.

- Information on pages 42 and 125 about handling missing configuration objects has been modified to include the functionality provided by the `ignore-missing-config-objs` configuration option.
- The descriptions of `Job_LoadRecent` on page 43 and page 85 have been modified to include the capability to optionally control the frequency with which `Job_LoadRecent` performs intraday aggregation. The new configuration option, `intraday-aggregates-frequency`, has been included in the discussion of Genesys Info Mart Server scheduling options on page 70. It has also been included in the list of relevant scheduling options in the sample intraday loading schedules on pages 88 and 90.
- The descriptions of `Job_LoadRecent` on page 43 and aggregation intervals for `Job_AggregateGIM` on page 52 have been modified to include the capability to optionally limit how far back in time the job will re-aggregate when it encounters newly loaded, late-arriving facts. A note on page 56, about optional parameters for manually running `Job_AggregateGIM`, has been expanded to clarify that the new configuration option, `max-late-arriving-fact-time-limit`, can be overridden by a time range that you specify manually.
- Various statements about the dependency of `Job_TransformGIM` on `Job_LoadRecent` have been modified to clarify that, if intraday loading is configured, `Job_TransformGIM` will not start until the load steps of `Job_LoadRecent` have completed. Figure 6 on page 89 and Figure 7 on page 90 have been modified to show that `Job_TransformGIM` might now overlap with portions of `Job_LoadRecent`.
- The description of `Job_LoadRecent` results for voice and Multimedia interactions (page 46) has been updated.
- The note on page 56, about optional parameters for manually running `Job_AggregateGIM`, has been modified to reflect changed behavior—you no longer have to specify the `-tenantKey` parameter if you want to specify a time range.
- Information about purging the Info Mart database, starting on page 112, has been updated to include new purge functionality provided in Interaction Concentrator release 7.6.1.
- The section “Database Connection Errors” on page 124 has been updated to include information about errors that might arise in large-scale deployments that have implemented the parallel-processing options to improve performance.

New in Document Version v7.6.002.00

The following topics have been added or significantly changed since the previous 7.6 release of this document:

- “New in This Release” on page 18

  - Ability to extract UserEvent-based key-value pair (KVP) data that is sent within a configurable timeout after the associated voice interaction ends
  - Option to include the last five minutes of extracted voice agent activity data when transforming data in a simple contact center environment
- Table 2, “IDB Tables for Gathering Statistics,” on page 38
  - Addition of `G_CUSTOM_DATA_S` table for the `ICON_CORE` role
- “Planning for Latency” on page 40 in the “Job_ExtractICON” section
  - Introduction of `user-event-data-timeout` configuration option
- “Unrecognized Call Flow Errors” on page 125, “Resources to Consult for Additional Information” on page 128, and Table 14, “Sample STG_ICON_DATA_DISCARDS Table Entries,” on page 126
  - New section to introduce the `STG_ICON_DATA_DISCARDS` table
  - Detailed description of the `STG_ICON_DATA_DISCARDS` table and table entries

Chapter

1 Introducing Genesys Info Mart 7.6

This chapter introduces you to Genesys Info Mart 7.6. It includes a brief overview and a list of new features available in this release:

About Genesys Info Mart 7.6

Genesys Info Mart 7.6 extracts data from multiple Genesys data sources—Interaction Concentrator (ICON), Genesys Voice Platform Voice Application Reporter (GVP VAR), and Stat Server (for backward compatibility only)—to produce the Info Mart Database, your data store for contact center historical reporting. Genesys Info Mart provides the following components to produce and manage the Info Mart Database:

- Genesys Info Mart Server (configured with the Genesys Info Mart application)
- Genesys Info Mart Administration Console

For detailed descriptions of the components, refer to the *Genesys Info Mart 7.6 Deployment Guide.*

Note: Genesys Info Mart continues to provide data extraction of voice agent state and reason details from Stat Server database for backward compatibility with deployments of prior Genesys Info Mart releases only. This guide provides information for using Stat Server database as a data source only as a courtesy to the customers whose legacy reports continue to use Stat Server data. New Genesys Info Mart deployments must extract data related to voice agent activity from Interaction Concentrator 7.5.or later.

Genesys Info Mart Server

The Genesys Info Mart Server launches a series of Extraction, Transformation, and Load (ETL) jobs 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. For more information about scheduling ETL jobs using the Genesys Info Mart Server, see "Launching ETL Jobs with Genesys Info Mart Server" on page 66.

Genesys Info Mart Administration Console

The Genesys Info Mart Administration Console provides a graphical user interface that you can use to monitor ETL job status and, when necessary, start or stop ETL jobs outside of the normal schedule. For more information, see "Managing ETL Jobs with Genesys Info Mart Administration Console" on page 74.

New in This Release

This section describes new or changed functionality that was introduced in the initial 7.6 release of Genesys Info Mart or in subsequent maintenance releases.

Note: If you choose to enable some of the new functionality that requires you to upgrade Interaction Concentrator, and if your deployment includes an earlier release of Interaction Concentrator, do not create a new ICON `Application` object in the Configuration Layer when upgrading Interaction Concentrator. Instead, use the existing `Application` object in the Configuration Layer when you install the Interaction Concentrator upgrade. For more information, see the "Recommendations on ICON Deployment and Upgrade" section of the *Genesys Info Mart 7.6 Deployment Guide*.

The 7.6 release of Genesys Info Mart provides the following new or changed functionality:

- Extracts voice agent state and reason details from Interaction Database (IDB), instead of from Stat Server. New Info Mart fact tables store details about states, reasons, and do-not-disturb (DND) modes for voice and Multimedia.

> **Note:** For backward compatibility with deployments of earlier Genesys Info Mart releases only, Genesys Info Mart continues to provide data extraction of voice agent state and reason details from Stat Server database. See Appendix , "Using Stat Server in Legacy Environments" on page 135 for more information.

- Provides high availability (HA) data extraction for voice agent login session, state, state reason, and DND mode details, which are extracted from an IDB that is populated by Interaction Concentrator release 7.6 that has been configured appropriately. (T-Server release 7.6 is required.)
- Loads open media interaction and agent activity details from an IDB into the Info Mart database, in an environment with release 7.6 of Interaction Concentrator and Interaction Server. (*Open Media* refers to a custom media channel that is supported on top of Genesys Multimedia. The `Workitem` media type is an example of Open Media.).
- Loads active Multimedia virtual queue details into the Info Mart database, and links virtual queue details to their corresponding target Multimedia interaction segment details.
- Loads active Multimedia chat interactions into the Info Mart database. Previously, only completed chat interactions were loaded into the Info Mart database.
- Provides HA data extraction for contact center configuration history details, which are extracted from IDBs that are populated by Interaction Concentrator release 7.6.
- Extracts data from IDB following the use of the Interaction Concentrator feature for resynchronization of configuration data.
- Provides data-quality improvements in HA data extraction for voice interaction details by comparing voice interaction data between the IDBs that constitute the HA pair.
- Provides detailed reasons for interactions that are cleared from a virtual queue, such as:
  - Target is cleared by routing strategy.
  - Interaction is routed by another, parallel virtual queue.
  - Interaction is default-routed by strategy.
  - Multimedia interaction is pulled back from strategy due to timeout.

Notes:

- The support for all four clearance scenarios requires an environment with Universal Routing Server of release 7.6 and Interaction Concentrator of release 7.6 that has been configured appropriately.
- The fourth clearance scenario requires Interaction Server 7.6 to report when a Multimedia interaction is cleared from a virtual queue or pulled from a routing strategy because it was not routed within the timeout configured for routing in Interaction Server.

- Provides uninterrupted durations for `After Call Work` (ACW) (for voice only) and `Not Ready` states, when interactions are initiated or received while in these states, in an environment with Interaction Concentrator release 7.6 that has been configured appropriately.

Note: For voice, the newly introduced fact tables in release 7.6 contain the data for this feature; the data is not available in the legacy fact tables that are implemented in previous releases.

- Provides data to calculate the number of voice interactions that are initiated or received while the agent is in ACW (voice only) or `Not Ready` states, in an environment with Interaction Concentrator release 7.6 that has been configured appropriately.

Note: For voice, the newly introduced fact tables in release 7.6 contain the data for this feature; the data is not available in the legacy fact tables that are implemented in previous releases.

- Associates `After Call Work` with the ACD or routed call, instead of with a consultation call, for the case in which the consultation call outlasts the original inbound customer call, in an environment with Interaction Concentrator release 7.6 that has been configured appropriately.
- Provides data to measure agent-to-agent consult talk duration, even if the consultation included an Interactive Voice Response (IVR) application or voice treatment port before the target agent answered the consultation.

Note: The newly introduced fact tables in release 7.6 contain the data for this feature; the data is not available in the legacy fact tables that are implemented in previous releases.

- Provides a set of new agent and interaction summary tables that facilitate aggregation for agent state and inbound voice interaction reporting.
- Provides several new interval-based and disposition-based aggregates, for use with either Genesys Interactive Insights or your own custom reports.

- Provides configurable control of transaction sizes for loading, aggregating, and purging data in the Info Mart database. This functionality provides improved capability for customers to control the database resources that are required to run the ETL jobs.
- Provides a new ETL job, `Job_MigrateGIM`, to migrate the data from the Staging Area and Genesys Info Mart databases of release 7.5 to release 7.6.

New in Release 7.6.004

Starting with Genesys Info Mart release 7.6.004, Genesys Info Mart provides the following new or changed functionality:

- Provides the ability to extract UserEvent-based key-value pair (KVP) data that is sent within a configurable timeout after the associated voice interaction ends.
- Provides the ability, at your option, to include the last five minutes of extracted voice agent activity data when transforming data in a simple contact center environment. This functionality improves the accuracy of agent reports for a given business day in a contact center that operates less than 24 hours a day. (A simple contact center is the one where an agent only logs in to a single switch, DN, or queue at a time, and where reporting requirements do not include the factoring of Do-Not-Disturb [DND] mode into summarized resource states and resource state reasons.)

New in Release 7.6.005

Starting with Genesys Info Mart release 7.6.005, Genesys Info Mart provides the following new or changed functionality that is relevant to the topics discussed in this *Operations Guide*:

- Improves ETL performance by enabling you to specify the frequency with which the intraday aggregation portion of `Job_LoadRecent` will run. In high-volume deployments, a short ETL cycle improves performance by keeping data sizes reasonable. However, the short cycle can result in repeated re-aggregation of overlapping time ranges, and this degrades the performance of intraday aggregation. Running intraday aggregation less frequently than `Job_LoadRecent` enables you to achieve a better balance between the two processes.

  A new configuration option, `intraday-aggregates-frequency`, supports this optional functionality. For more information about the new configuration option, see the *Genesys Info Mart 7.6 Deployment Guide.*
- When you run `Job_AggregateGIM` manually, enables you to selectively specify the optional parameters. You can now specify the start and end date parameters without also having to specify the tenant key. For more information about the optional parameters, see page 56.

For information about additional functionality introduced in release 7.6.005, see the *Genesys Info Mart 7.6 Deployment Guide* for this release.

New in Release 7.6.006

Starting with Genesys Info Mart release 7.6.006, Genesys Info Mart provides the following new or changed functionality that is relevant to the topics discussed in this *Operations Guide*:

- Improves ETL performance for large-scale inbound voice contact centers, primarily by running several steps in parallel, rather than sequentially. Improvement in ETL performance was observed during testing in large-scale deployments using Oracle 10 and running the ETL on either Solaris 10 or Windows 2003.

  Several new configuration options in a new configuration section, [gim-tuning], enable these performance improvements. You should implement these improvements only if the servers and databases that the ETL uses have sufficient processing power and resources.

  For more information about the new configuration options and their effect on ETL processing, see the *Genesys Info Mart 7.6 Deployment Guide.*
- Provides the option to allow the ETL to automatically ignore unresolved references to configuration objects, rather than require the user to manually run the failed ETL job from the Genesys Info Mart Administration Console with the -ignoreMissingConfigObjs parameter.

  For links to more information about implementing this functionality, see the "Document Change History" on page 12. For more information about the new configuration option, ignore-missing-config-objs, see the *Genesys Info Mart 7.6 Deployment Guide.*
- Modifies existing functionality to enable the purging of fact data from Info Mart that is only three days old. Previously, fact data had to be at least 30 days old before it was eligible to be purged.
- Provides the option to limit how far back in time the aggregation processing in Job_LoadRecent and Job_AggregateGIM will go when considering newly loaded, late-arriving facts. Previously, it was not possible to set a limit, and all late-arriving facts would be re-aggregated. Extensive re-aggregation can result in excessive run times for Job_LoadRecent and Job_AggregateGIM, particularly when contact center agents forget to log off, and therefore stay in the same state for many days, or when Interaction Concentrator (ICON) restarts and creates login sessions for agents that have been logged in for many days.

  For links to more information about implementing this functionality, see the "Document Change History" on page 12. For more information about the new configuration option, max-late-arriving-fact-time-limit, see the *Genesys Info Mart 7.6 Deployment Guide.*

- Provides the option not to store Interaction Segment Fact data for voice media. This feature is for deployments where reports can be created from Interaction Resource Fact data for voice media. Disabling the storage of Interaction Segment Fact data can improve ETL performance.

  A new configuration option, `populate-voice-ixn-seg-facts`, supports this optional functionality. For more information about the new configuration option, see the *Genesys Info Mart 7.6 Deployment Guide.*
- Modifies a number of aggregation log event messages to include tenant information.

  For example—information message 55-32002 has been changed from `"Aggregation - aggregation complete for table [name]"` to `"Aggregation - aggregation complete for table [name] and tenant [name]."`

For information about additional functionality introduced in release 7.6.006, see the *Genesys Info Mart 7.6 Deployment Guide* for this release.

New in Release 7.6.007

Starting with Genesys Info Mart release 7.6.007, Genesys Info Mart provides the following new or changed functionality that is relevant to the topics discussed in this *Operations Guide*:

- Provides support to improve ETL and Genesys Info Mart Administration Console performance by providing an optional mechanism for purging the historical information about steps that the ETL has performed related to job execution, source data extraction, target table loading, table purging, and data aggregation from the Staging Area database.

  A new configuration option, `days-to-keep-stg-history`, enables this functionality. For more information about `days-to-keep-stg-history` and its effect on ETL processing, see the *Genesys Info Mart 7.6 Deployment Guide.*
- Provides support to improve the performance of the interval-based aggregation queries used to populate data for the `Interaction-Agent Interval` and `Agent-State Interval` aggregates.

  A new configuration option, `interval-aggregates-fact-time-window`, enables this functionality. For more information about the `interval-aggregates-fact-time-window` configuration option, see the *Genesys Info Mart 7.6 Deployment Guide.*
- Modifies the time range of data that is aggregated in a single database transaction by `Job_LoadRecent` and `Job_AggregateGIM` aggregation queries. The value can now be specified in `HOURS` in the `aggregate-time-range-units` configuration option on the `[gim-etl]` section of the Genesys Info Mart `Application` object's `Options` tab.

  For more information about the `aggregate-time-range-units` configuration option, see the *Genesys Info Mart 7.6 Deployment Guide.*

New in Release 7.6.008

Starting with Genesys Info Mart release 7.6.008, Genesys Info Mart provides the following new or changed functionality that is relevant to the topics discussed in this *Operations Guide*:

- Provides high availability (HA) data extraction of Outbound Contact details, which are extracted from an IDB that is populated by Interaction Concentrator release 8.0 and has been configured appropriately. See the *Genesys Info Mart 7.6.x Release Notes* for the minimum release of Interaction Concentrator 8.0 that is required to support this functionality.
- Enables automatic retry of any failed job or Genesys Info Mart Server exception using a configured number of retries with a configurable delay between retries.

  Two new configuration options, `job-retry-count` and `job-retry-wait` enable this functionality. For more information about the `job-retry-count and job-retry-wait` configuration options, see the *Genesys Info Mart 7.6 Deployment Guide.*
- Provides support for the automated re-run of `Job_ExtractICON` for `role=ICON_CFG` and `Job_TransformGIM` when `Job_TransformGIM` fails after encountering an unresolved configuration object.
- Provides support for extracting voice interaction data from topologies where not all T-Servers or IVR Servers involved in the call flow are monitored by ICON.

  This feature enables Info Mart to provide reporting data in the following types of environments:
  - Network routing or network parking are used, but you want Genesys Info Mart to store data for only the premise-portions of the interactions.
  - There are multiple sites or multiple tenants, but you want Genesys Info Mart to store data for only some of the sites or tenants.

  A new configuration option, `extract-partially-merged-interactions`, enables this functionality. For more information about the `extract-partially-merged-interactions` configuration option, see the *Genesys Info Mart 7.6 Deployment Guide.*

New in Release 7.6.009

Starting with Genesys Info Mart release 7.6.009, Genesys Info Mart provides the following new or changed functionality that is relevant to the topics discussed in this *Operations Guide*:

- Provides the option to maintain database table statistics for fact tables in the Info Mart database on your own, or to continue letting Genesys Info Mart maintain them as needed by the ETL.

Two new Genesys Info Mart application options have been added to support this new capability, `run-intraday-fact-table-stats`, and `run-historical-fact-table-stats`. For more information about the `run-intraday-fact-table-stats`, and `run-historical-fact-table-stats` configuration options, see the *Genesys Info Mart 7.6 Deployment Guide*.

- Provides the option to specify how long Genesys Info Mart stores information about login sessions that have been extracted from one side of an High Availability (HA) pair of Interaction Concentrator Databases (IDBs), but not the other; this allows the duplicate session to be ignored when it does eventually appear in the second side of the HA pair.

  A new Genesys Info Mart application option, `days-to-keep-stg-ha-login-sessions`, has been added to support this new functionality. For more information about the `days-to-keep-stg-ha-login-sessions` configuration option, see the *Genesys Info Mart 7.6 Deployment Guide*.

New in Release 7.6.011

Starting with Genesys Info Mart release 7.6.011, Genesys Info Mart provides the following new or changed functionality that is relevant to the topics discussed in this *Operations Guide*:

- Supports the Interaction Concentrator 7.6.1 and 8.0 capability to continue storing information about multimedia interactions that were active when the Multimedia ICON is stopped and subsequently restarted.

  Two Interaction Concentrator application options support this capability, `calls-in-the-past` and `om-force-adata`. For more information about the `calls-in-the-past` and `om-force-adata` configuration options, see the *Genesys Info Mart 7.6 Deployment Guide.*

  In this guide, see "Restarting a Multimedia ICON" on page 117 for considerations, and the procedure you should follow when you need to restart a Multimedia ICON in order to minimize data loss or data quality issues within Genesys Info Mart.

- For voice interactions, supports the Interaction Concentrator 8.0 capability to associate call-based Key-Value Pair (KVP) data with the Routing Point or Agent party that attached or updated the KVP data when they are no longer an active call party. For more information about this Interaction Concentrator 8.0 capability, see the *Interaction Concentrator 8.0 User's Guide*.

New in Release 7.6.012

Starting with Genesys Info Mart release 7.6.012, Genesys Info Mart provides the following new or changed functionality that is relevant to the topics discussed in this *Operations Guide*:

- Provides new disposition-based aggregates, `AG2_OUT_V_IXN_AGENT_*` and `AG2_OUT_V_IXN_AGENT_GRP_*`, from which you can build your own custom reports to measure agent and agent group handling of outbound and internal voice interactions based on key business attributes, such as customer segment, service type, and service subtype.

  A new configuration option, `populate-ixn-agent-out-aggregates`, enables this functionality. For more information about the `populate-ixn-agent-out-aggregates` configuration option, see the *Genesys Info Mart 7.6 Deployment Guide*.

Chapter

2 Understanding Genesys Info Mart ETL Jobs

This chapter describes the ETL 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.
- Calculate and load aggregated data into the Info Mart aggregate tables.
- Purge old data from the Info Mart fact tables.
- Migrate existing data to the latest release of the Staging Area and Info Mart database schemas.

This chapter contains the following sections:

ETL Job Summary Table

Table 1 lists the jobs that ship 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 "About the Intraday ETL Cycle" on page 83.

Table 1: ETL Job Summary

Name	Function	Frequency			Notes
		Once	Intraday	Daily	
Job_InitializeGIM	Populates many of the dimension tables in the Info Mart database and Staging Area database with fixed information. Updates the `ENTERPRISE_DATE` dimension with configured fiscal periods.	✔			Run this job once manually during the initial deployment. If the configuration of `enterprise-fiscal-periods` changes, run this job again.
Job_ExtractICON	Extracts new and changed records from a single extraction role from one or more IDBs, and stores those records in the Staging Area. This job performs intra-IDB and multi-IDB interaction merges to reconcile links between interactions that contain related calls.		✔		The Genesys Info Mart Server launches this job once for each extraction role (`ICON_CFG`, `ICON_CORE`, `ICON_OCS`, `ICON_MM`—see page 34) on each IDB, as long as multi-IDB merge and HA are not configured. **Note:** To support new features introduced in Genesys Info Mart 7.6, Genesys Info Mart requires Interaction Concentrator release 7.6 or 8.0.

Table 1: ETL Job Summary (Continued)

Name	Function	Frequency			Notes
		Once	Intraday	Daily	
Job_ExtractICON (continued)	This job extracts data from redundant data sources to protect against a loss of source data in two different ways: • For high availability (HA) of Configuration and Voice details, deduplicates redundant data from paired HA IDBs. • For HA of Outbound Contact details, detects gaps in source data and performs a time based switch over from one IDB to the other in the paired HA IDBs.		✔		In particular, to support high availability (HA) of Outbound Contact details, Genesys Info Mart (release 7.6.008 or later) requires Interaction Concentrator 8.0. See the *Genesys Info Mart 7.6.x Release Notes* for the minimum release of Interaction Concentrator 8.0 that is required to support this functionality. If multi-IDB merge is configured for the Staging Area database access point (DAP), then one instance of `Job_ExtractICON` is launched to extract voice data from all DAPs with the role of `ICON_CORE` in order to merge multi-IDB calls. If multi-IDB merge has not been configured, but HA for ICON voice details is configured, `Job_ExtractICON` will run once for each HA pair of DAPs. If HA for ICON configuration details is configured, one instance of `Job_ExtractICON` is launched to extract configuration data from the HA pair of DAPs with the role of `ICON_CFG` in order to deduplicate redundant configuration data. If HA for ICON Outbound Contact details is configured, one instance of `Job_ExtractICON` is launched to extract Outbound Contact data from the HA pair of DAPs with the role of `ICON_OCS`. Rather than deduplicate, the job switches between IDBs when gaps in the source data are detected.

Table 1: ETL Job Summary (Continued)

Name	Function	Frequency			Notes
		Once	Intraday	Daily	
Job_ExtractICON (continued)			✔		If you need to run this job manually from the Genesys Info Mart Administration Console, you must run it once and specify `ALL SOURCES`. **Note:** From the Genesys Info Mart Administration Console, make sure that you select `ALL SOURCES`.
Job_ExtractGVP	Extracts new and changed records from a single Genesys Voice Platform Voice Application Reporter (GVP VAR) database and stores those records in the Staging Area database.		✔		The Genesys Info Mart Server launches this job once for each GVP VAR database. If you need to run this job manually from the Genesys Info Mart Administration Console, you must run it once for each GVP VAR database, or specify `ALL SOURCES`. **Note:** From the Genesys Info Mart Administration Console make sure you select the correct DAP, or select `ALL SOURCES`.
Job_ExtractSS	Extracts new and changed records from a single Stat Server database (in legacy reporting environments only) and stores those records in the Staging Area database.		✔		The Genesys Info Mart Server launches this job once for each Stat Server database. If you need to run this job manually from the Genesys Info Mart Administration Console, you must run it once for each Stat Server database, or specify `ALL SOURCES`. **Note:** From the Genesys Info Mart Administration Console make sure you select the correct DAP, or select `ALL SOURCES`.

Table 1: ETL Job Summary (Continued)

Name	Function	Frequency			Notes
		Once	Intraday	Daily	
Job_TransformGIM	Cleanses and transforms previously extracted IDB, GVP VAR, and Stat Server (in legacy environments only) data within the Staging Area.		✔		
Job_LoadRecent	Loads previously transformed data from the Staging Area database into the dimension and intraday fact tables in the Info Mart database and re-aggregates the most recently loaded intraday data.		✔	✔	If you plan to use intraday loading, this job will be run as part of the intraday ETL cycle. If you do not plan to use intraday loading, this job will be run once per day, after the last intraday ETL cycle.
Job_LoadGIM	Moves data from the intraday fact tables to the historical fact tables in the Info Mart database.			✔	
Job_AggregateGIM	Aggregates or re-aggregates the historical aggregate tables based on data that has changed since the last load of the historical fact tables.			✔	
Job_MaintainGIM	Purges old data from the historical fact and aggregate tables in the Info Mart database, or marks the data for purging, using rules that are configured in the Genesys Info Mart application.			✔	

Table 1: ETL Job Summary (Continued)

Name	Function	Frequency			Notes
		Once	Intraday	Daily	
Job_MigrateGIM	Migrates the data from the Staging Area and Info Mart databases of release 7.5 to release 7.6, either in one run or according to the schedule that is configured in the Genesys Info Mart application.			✔	For new deployments, this job is not required. If you are migrating from a previous release of Genesys Info Mart, you must run this job manually from the Genesys Info Mart Administration Console to perform the data migration from the critical Staging Area and Info Mart database tables; the data migration from critical tables is required in order to run the ETL jobs with this release of Genesys Info Mart. Once the critical data migration is complete, you can schedule this job to complete the migration of historical data from non-critical tables outside of the intraday ETL cycle. You can configure the amount of data that each execution of the job migrates.

Job_InitializeGIM

This job performs the following functions:

- Populates the following dimensions with fixed information:
  - CURRENCY
  - RESOURCE_STATE
  - MEDIA_TYPE
  - INTERACTION_TYPE
  - TECHNICAL_DESCRIPTOR
  - ENTERPRISE_DATE
  - TIME_OF_DAY
  - DATE_TIME
  - AUDIT_
  - CAMPAIGN_GROUP_STATE
  - CONTACT_INFO_TYPE
  - CALL_RESULT
  - DIALING_MODE
  - RECORD_STATUS
  - RECORD_TYPE
  - INTERACTION_RESOURCE_STATE
  - RESOURCE_GROUP_COMBINATION
- Updates the ENTERPRISE_DATE dimension with configured fiscal periods.

Run this job manually once during initial deployment, and again each time that you change the enterprise fiscal period configuration. This job also populates the DECODE tables in the Staging Area database.

The ENTERPRISE_DATE dimension is loaded with complete calendar years from January 1, 1995 through December 31, 2024. Fiscal period columns are updated for complete fiscal years by using the configured enterprise-fiscal-periods options. Fiscal period columns for partial fiscal years at the beginning and the end of the dimension are set to null. The TIME_OF_DAY dimension is loaded with enough rows to represent each minute of the day.

Notes:

- Fiscal period definitions contain many variations. Each enterprise has its own rules for defining fiscal periods, which change from year to year and contain anomalies to account for calendar-year variations, such as leap year and the day of the week that begins the year.

  The options that you select in the enterprise-fiscal-periods configuration section, and the ENTERPRISE_DATE fiscal period column values that Genesys Info Mart populates, are only suggested values. If you want to use fiscal periods for reporting, Genesys strongly recommends that you carefully analyze ENTERPRISE_DATE's column values and customize them to suit your environment. Be sure to customize only the fiscal period columns—that is, those columns that begin with FISCAL_.
- Job_InitializeGIM updates the ENTERPRISE_DATE FISCAL_ period columns each time that you change the Genesys Info Mart application's enterprise-fiscal-periods option values. If you change the option values and run Job_InitializeGIM after you have customized the fiscal period columns to suit your environment, Job_InitializeGIM will overwrite your customized column values.

Unlike the ENTERPRISE_DATE dimension, which is loaded with complete calendar years from January 1, 1995 through December 31, 2024, the DATE_TIME dimension spans January 1, 2006 to December 31, 2013, plus a single row for the last 15-minute interval of December 31, 2025. Active facts use an END_TIME of December 31, 2025, and a DATE_TIME key for the end of an active fact will use that maximum value key from the DATE_TIME dimension.

Job_ExtractICON

This job extracts new and changed data from an IDB, and stores it in the Staging Area database. The Genesys Info Mart Server launches a separate instance of Job_ExtractICON for each extraction role on one or more DAPs during each intraday ETL cycle. After all the extraction jobs have completed

successfully, the Genesys Info Mart Server launches the job that transforms all the extracted ICON, GVP VAR, and Stat Server (if applicable) data.

The Genesys Info Mart Server launches a single instance of `Job_ExtractICON` in the following configurations:

- If multi-IDB merge is configured for the Staging Area DAP, one instance of `Job_ExtractICON` is launched to extract voice data from all DAPs with the role of `ICON_CORE` in order to merge multi-IDB calls.
- If multi-IDB merge has not been configured, but HA for ICON voice details is configured, `Job_ExtractICON` will run once for each HA pair of DAPs with the role of `ICON_CORE`.
- If both multi-IDB merge and HA for ICON voice details are configured, one instance of `Job_ExtractICON` is launched to extract voice data from all DAPs, including HA pairs, with the role of `ICON_CORE`.
- If HA for ICON configuration details is configured, one instance of `Job_ExtractICON` is launched to extract configuration data from the HA pair of DAPs with the role of `ICON_CFG` in order to deduplicate redundant configuration data.
- If HA for ICON Outbound Contact details is configured, one instance of `Job_ExtractICON` is launched to extract Outbound Contact data from the HA pair of DAPs with the role of `ICON_OCS`.

About Extraction Roles

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

- `ICON_CFG`
- `ICON_CORE` (see page 35)
- `ICON_OCS` (see page 36)
- `ICON_MM` (see page 37)

ICON_CFG

For the `ICON_CFG` role, `Job_ExtractICON`:

- Extracts all new and changed data from IDB tables that store the contact center configuration history, and stores the data in the Staging Area database.
- Performs HA deduplication of configuration history data in a configuration with an HA pair of DAPs with the role `ICON_CFG`.

HA Data Extraction

Genesys Info Mart uses redundant HA IDBs to protect valuable data in case of database failure. With HA, two ICON processes work in parallel—gathering

redundant sets of the same details, and storing the data in their respective IDBs independently. You can deploy the HA architecture for ICON configuration details.

If you configure your system with HA IDBs, `Job_ExtractICON` deduplicates the data prior to transformation. To configure ICON processes and their DAPs for HA deduplication, see the *Genesys Info Mart 7.6 Deployment Guide*.

ICON_CORE

For the `ICON_CORE` role, `Job_ExtractICON` extracts:

- Completed virtual queue details.
- Completed voice interaction details, such as calls, attached data, and UserEvent-based 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 Staging Area database. The job also performs the following data streamlining before transformation:

- IDB Merge—Intra-IDB merge and multi-IDB merge
- HA Data Extraction (see page 36)

IDB Merge

Two levels of IDB Merge are intra-IDB merge and multi-IDB merge.

Intra-IDB Merge

Intra-IDB merge is the process in which `Job_ExtractICON` invokes the ICON `gsysIRMerge` stored procedure against the IDB tables before extracting voice interaction data. The stored procedure resolves links between related multi-site calls in IDB.

Note: `Job_ExtractICON` performs intra-IDB merge only if Genesys Info Mart Server is *not* configured to run the `gsysIRMerge` stored procedure periodically (option `ir-merge-interval` in section `[gim-etl]`). For information about this configuration option, see the *Genesys Info Mart 7.6 Deployment Guide*.

Multi-IDB Merge

Multi-IDB Merge is a second merge that is required when your deployment includes voice details that span IDBs (you configured more than one DAP with role = `ICON_CORE`). In this case, `Job_ExtractICON` extracts voice details that have been successfully intra-IDB-merged and stores them in the Merge Staging Area database schema. `Job_ExtractICON` invokes the ICON `gsysIRMerge` stored procedure against the Merge Staging Area database schema to resolve links between related multi-site calls that span IDBs. Finally, `Job_ExtractICON`

extracts successfully multi-IDB-merged voice interactions to the Staging Area database.

HA Data Extraction

Genesys Info Mart uses redundant HA IDBs to protect valuable data in case of database failure. With HA, two ICON processes work in parallel—gathering redundant sets of the same details, and storing the data to their respective IDBs independently. You can deploy the HA architecture for ICON voice details, which provides redundancy for virtual queue, voice interaction, agent activity, attached data, and UserEvent-based KVP data details.

If you configure your system with HA IDBs, `Job_ExtractICON` deduplicates the data prior to transformation. To configure ICON processes and their DAPs for HA deduplication, see the *Genesys Info Mart 7.6 Deployment Guide*.

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 do configure a DAP for `ICON_CORE`, 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 two 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 7.6 Deployment Guide* about configuring Genesys Info Mart and data sources to extract both voice and multimedia details.

ICON_OCS

For the `ICON_OCS` role, `Job_ExtractICON`:

- Extracts all new and changed data from IDB tables that store Outbound Contact data, and stores the data in the Staging Area database.
- Performs switch over from one IDB to the other when IDB metadata indicates there is a gap in source data, in a configuration with an HA pair of DAPs with the role `ICON_OCS`.

HA Data Extraction

Genesys Info Mart uses redundant HA IDBs to protect valuable data in case of database failure. You can deploy an HA architecture for ICON Outbound Contact details. With this HA deployment, the ICON processes that constitute an HA pair receive events from the same Outbound Contact Server (or an HA pair of primary and backup Outbound Contact Servers) and have the same configuration option settings. The only difference is that each stores data in its own IDB. Rather than deduplicating source data, `Job_ExtractICON` extracts from one IDB in the pair until it detects a gap in the source data and switches to the other IDB, as required, to extract Outbound Contact data for particular time periods. The ICONs must be dedicated to storing Outbound Contact details only, and must not store other types of data, such as Configuration details, Voice details, or Multimedia details. To configure ICON processes and their DAPs for HA of Outbound Contact details, see the *Genesys Info Mart 7.6 Deployment Guide*.

Note: Extracting Outbound Contact details from ICON is optional. If you do not wish to extract these details, do not configure a DAP for role `ICON_OCS`. If you do configure a DAP for role `ICON_OCS`, the Genesys Info Mart Server will run `Job_ExtractICON` for each DAP with this role.

ICON_MM

For the `ICON_MM` role, `Job_ExtractICON` extracts:

- All new and changed data from IDB tables that store Multimedia interactions. Both active and completed Multimedia solution interactions are extracted. 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 Staging Area 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 two 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 only one DAP for the `ICON_MM` role, to extract Multimedia details from only one IDB. For more information, see the section in the *Genesys Info Mart 7.6 Deployment Guide* about configuring Genesys Info Mart and data sources to extract both voice and multimedia details.

Restarting a Multimedia ICON

Special considerations and procedures must be followed when restarting a Multimedia ICON to avoid data quality issues within Genesys Info Mart. For detailed information about successfully restarting a Multimedia ICON, see Restarting a Multimedia ICON, page 117.

Improving System Performance During ICON Extraction

To improve system performance, Genesys recommends that you gather statistics on IDB tables after the tables are initially populated, but before you run the initial ICON extraction cycle, and on a regular basis thereafter. Table 2 lists different ICON role configurations and the corresponding IDB tables that are necessary for gathering these statistics.

Table 2: IDB Tables for Gathering Statistics

Genesys Info Mart Extraction Role(s)	ICON Role Option Values	IDB Tables
`ICON_CORE`, `ICON_MM`	`gcc`—Stores call-related and party-related information. `gls`—Stores data regarding agent states and agent login sessions. `gud`—Stores data regarding the attached data that is associated with calls and UserEvent-based data.	• `G_AGENT_STATE_HISTORY` • `G_AGENT_STATE_RC` • `G_CALL` • `G_CUSTOM_DATA_S` (applicable to the `ICON_CORE` role only) • `G_DND_HISTORY` • `G_IR` • `G_IS_LINK` (applicable to the `ICON_CORE` role only) • `G_IS_LINK_HISTORY` (applicable to the `ICON_CORE` role only) • `G_LOGIN_SESSION` • `G_PARTY` • `G_PARTY_HISTORY`

Table 2: IDB Tables for Gathering Statistics (Continued)

Genesys Info Mart Extraction Role(s)	ICON Role Option Values	IDB Tables
ICON_CORE, ICON_MM (continued)		• G_PARTY_STAT • G_ROUTE_RESULT • G_SECURE_USERDATA_HISTORY • G_USERDATA_HISTORY • G_VIRTUAL_QUEUE • GM_F_USERDATA • GM_L_USERDATA • GX_SESSION_ENDPOINT
ICON_OCS	gos—Stores Outbound Contact data regarding outbound calls and campaigns.	• GO_CAMPAIGN • GO_CAMPAIGNHISTORY • GO_CHAIN • GO_CHAINREC_HIST • GO_FIELDHIST • GO_METRICS • GO_SEC_FIELDHIST
ICON_CFG	cfg—Stores the initial configuration state and a history of configuration changes retrieved from Configuration Server.	• GC_AGENT • GC_CALLING_LIST • GC_CAMPAIGN • GC_ENDPOINT • GC_GROUP • GC_IVR • GC_IVRPORT • GC_PLACE • GC_SCRIPT • GC_SKILL • GC_SWITCH • GC_TENANT • GC_TIME_ZONE • GCX_CAMPGROUP_INFO • GCX_CAMPLIST_INFO • GCX_ENDPOINT_PLACE • GCX_GROUP_AGENT

Table 2: IDB Tables for Gathering Statistics (Continued)

Genesys Info Mart Extraction Role(s)	ICON Role Option Values	IDB Tables
ICON_CFG (continued)		• GCX_GROUP_ENDPOINT • GCX_GROUP_PLACE • GCX_SKILL_LEVEL

Planning for Latency

When planning the intraday ETL schedule for Job_ExtractICON, make allowances for data latency. You want to ensure that IDB has had sufficient time to record all the previous day's data before the final extraction begins.

The user-event-data-timeout configuration option can also be set to increase the latency between the end of voice interactions and data extraction. This allows agents to record UserEvent-based key-value pair (KVP) data associated with a voice interaction within a configurable period of time after a voice interaction has ended. For more details, refer to the *Genesys Info Mart 7.6 Deployment Guide*.

The following must be accounted for when configuring the ETL schedule to allow the entire business day's voice interaction data to be processed. If unaccounted for, the voice interactions that occurred around the end of the business day will not be ETL'd until the beginning of the next business day.

Suppose the voice interaction traffic ends at time *x,* then the following two conditions must be met in order to have voice traffic extracted through time *x*:

1. There must be more than *two* intraday ETL cycles after time *x*;
2. (Start time of ETL cycle - user-event-data-timeout) > time *x*

Controlling the Volume of Data Extraction

You can configure Genesys Info Mart to control the volume of the data extracted from an IDB at any one time, for all types of data except configuration data. You need to configure the following dependent configuration options in section [gim-etl] to set the volume of extraction:

- limit-extract-data
- extract-data-after-date
- extract-data-time-range-limit
- extract-data-time-range-units

When configured, the volume of data extraction applies equally to all Genesys databases in the system; you cannot apply this function individually to any one Genesys Info Mart data source. Controlling the volume of data extraction

occurs once per ETL cycle. To set these configuration options to control the volume of data extraction, see the *Genesys Info Mart 7.6 Deployment Guide*.

Job_ExtractGVP

This job extracts new and changed data from one GVP Voice Application Reporter database and stores it in the Staging Area database. The Genesys Info Mart Server launches this job for each GVP VAR database during each intraday ETL cycle. After all the extraction jobs have completed successfully, the Genesys Info Mart Server launches the job that transforms all the extracted IDB, Stat Server (if applicable), and GVP VAR data.

Note: Extracting GVP voice application details from GVP VAR is optional. If you do not wish to extract these details, do not configure a DAP for role `GVP_VAR`. If you do configure a DAP for role `GVP_VAR`, the Genesys Info Mart Server will run `Job_ExtractGVP` for each DAP with this role.

Planning for Latency

When planning the intraday ETL schedule, make allowances for data latency. You want to ensure that the GVP VAR database has had sufficient time to record all the previous day's data before the final extraction begins.

Controlling the Volume of Data Extraction

You can configure Genesys Info Mart to control the volume of the data extracted from a GVP VAR database at any one time. You need to configure the following dependent configuration options in section [`gim-etl`] to set the volume of extraction:

- `limit-extract-data`
- `extract-data-after-date`
- `extract-data-time-range-limit`
- `extract-data-time-range-units`

When configured, the volume of data extraction applies equally to all Genesys databases in the system; you cannot apply this function individually to any one Genesys Info Mart data source. Controlling the volume of data extraction occurs once per ETL cycle. To set these configuration options to control the volume of data extraction, see the *Genesys Info Mart 7.6 Deployment Guide*.

Job_TransformGIM

This job cleanses and transforms data extracted from all Interaction Concentrator, GVP VAR, and Stat Server (if applicable) databases within the Staging Area database. For more information about the scope of the data transformation, see "About the Intraday ETL Cycle" on page 83.

The Genesys Info Mart Server launches this job during each intraday ETL cycle after it has extracted data from all Interaction Concentrator, GVP VAR, and Stat Server (if applicable) databases.

Handling Configuration Errors

By default, when `Job_TransformGIM` encounters unresolved references to configuration objects in the extracted data, it logs an error, stops processing and automatically tries to run again.

During this automatic retry, `Job_ExtractICON` (with the `role ICON_CFG`) and `Job_TransformGIM` attempt to run again. There are two configuration options, `job-retry-count` and `job-retry-wait`, that control the number of attempts to retry a failed job and how many minutes to wait between retries. To set these configuration options, see the *Genesys Info Mart 7.6 Deployment Guide*. This automatic retry will only occur when `Job_TransformGIM` is run from the Genesys Info Mart Server scheduler and not from the Genesys Info Mart Administration Console. For more information, see "Automatically Retrying Failed ETL Jobs with Genesys Info Mart Server" on page 72.

Note: If the `ignore-missing-config-objs` configuration option is set to `TRUE`, then `job-retry-count` and `job-retry-wait` are ignored and no retry is attempted. For more information on the `ignore-missing-config-objs` configuration option, see "Ignoring Configuration Errors" on page 43.

If the automatic retries of `Job_ExtractICON` and `Job_TransformGIM` have completed but were unsuccessful in extracting data, you can still resolve the error using ICON to perform a manual resynchronization of the configuration data between the Configuration Database and the IDB from which Genesys Info Mart extracts the configuration history. If the manual resynchronization does not resolve the error, an optional command-line parameter, `-ignoreMissingConfigObjs`, enables `Job_TransformGIM` to process data that refers to missing configuration objects.

Depending on the type of data being processed, `Job_TransformGIM` takes one of the following actions:

- Transforms the source data, supplying default dimension keys for the missing objects.
- Discards the source data, because the data cannot be transformed without the missing objects.

You must use the Genesys Info Mart Administration Console in order to pass the input parameter to `Job_TransformGIM`. For more information about recovering from missing configuration objects and for instructions on using the `-ignoreMissingConfigObjs` parameter, refer to "Recovering from Unresolved References to Configuration Objects" on page 130.

Ignoring Configuration Errors

Starting with Genesys Info Mart release 7.6.006, you can configure the Genesys Info Mart application so that `Job_TransformGIM` will ignore missing configuration objects and will not fail when it encounters unresolved references to configuration objects in the extracted data.

To configure the Genesys Info Mart application to ignore missing configuration objects, set the `ignore-missing-config-objs` option in the `[gim-transformation]` configuration section to `TRUE`.

From the point of view of an individual instance of `Job_TransformGIM`, configuring the application to automatically ignore missing configuration objects is the equivalent of manually running `Job_TransformGIM` with the `-ignoreMissingConfigObjs` command-line parameter. However, `Job_TransformGIM` does not log any message about the error; there is no possibility of resynchronizing ICON and Configuration Server; and there is no possibility of correcting the data before transforming this set of extracted data.

Accordingly, Genesys recommends that you do not enable this option unless your environment is:

- A lab, where it is acceptable to discard or improperly transform some source data.
- A high-volume deployment, where a separate intraday ETL is installed and it is preferable to discard or improperly transform some source data, rather than have the ETL stop running until the problem is corrected manually. In very high-volume contact centers, an extended outage can cause the ETL to fall too far behind to be able to catch up to the current intraday reporting interval.

Job_LoadRecent

This job loads previously transformed data from the Staging Area database into the dimension tables and intraday fact tables in the Info Mart database.

- For the fact tables that support only completed facts, `Job_LoadRecent` inserts new rows that represent completed facts in the intraday fact tables.

- For the fact tables that support both active and completed facts, Job_LoadRecent inserts new rows that represent those active or completed facts in the intraday fact tables. It also updates the active flags, end times, durations, and other information of the active facts that it loaded previously into the intraday fact tables.
- If you enabled intraday aggregation, Job_LoadRecent uses the new and changed intraday fact table data to calculate (or recalculate) the intraday aggregates in the Info Mart database at the subhour and hour levels. (Job_AggregateGIM calculates day- and month-level aggregates.)

  If your intraday and/or historical fact tables already have data in them, the following situations require special handling in order to prevent a single execution of Job_LoadRecent from aggregating those facts over a large time span, and in so doing, taking many hours to complete:
  - You are enabling some particular aggregate for the first time, or re-enabling an aggregate that you have temporarily disabled for some period of time. For example, you are going to set [gim-aggregates-tenant] populate-ixn-service-type-aggregates=TRUE.
  - You are enabling intraday aggregation for the first time, or re-enabling intraday aggregation after it has been temporarily disabled, and Job_AggregateGIM has not produced historical aggregates for an extended period of time. Intraday aggregation is enabled by setting [schedule] populate-intraday-aggregates=TRUE.

  See “Calculating Aggregates Within a Time Span” on page 53 for the procedure you should follow to prevent a single execution of Job_LoadRecent from aggregating all eligible fact data.

  Starting with Genesys Info Mart release 7.6.005, you can improve ETL performance by running the intraday aggregation portion of Job_LoadRecent less frequently than the rest of the job. For more details about configuring this functionality, see the information about the intraday-aggregates-frequency option in the *Genesys Info Mart 7.6 Deployment Guide.*

  Starting with Genesys Info Mart release 7.6.006, you can improve ETL performance by limiting how far back in time the aggregation processing in Job_LoadRecent will re-aggregate when it encounters newly loaded, late-arriving facts. For more details about configuring this functionality, see the information about the max-late-arriving-fact-time-limit option in the *Genesys Info Mart 7.6 Deployment Guide.*
- Job_LoadRecent also updates information in the historical fact tables to account for late-arriving source data, such as After Call Work counts and durations for voice interaction facts, and interaction ID for contact attempt facts and GVP call facts.

To ensure that users can query intraday fact tables during the load process, Job_LoadRecent does not use bulk loading. To minimize loading time, do *not* add a lot of indexes to the intraday fact and aggregate tables. Generally, the intraday tables are relatively small and should not require many indexes.

If you plan to use intraday loading, configure Genesys Info Mart Server to launch Job_LoadRecent as part of the intraday ETL cycle. If you do not plan to use intraday loading, configure the time of day that you want Genesys Info Mart Server to launch Job_LoadRecent. This is generally once a day after the final intraday ETL cycle has completed and before you run Job_LoadGIM.

Note: You must use Job_LoadRecent even if you do not plan to use intraday loading. Job_LoadRecent is the only job that moves transformed data from the Staging Area database to the Info Mart database.

For information about loading the historical tables, refer to the "Job_LoadGIM" and "Job_AggregateGIM" sections on page 50.

Table 3 summarizes how Job_LoadRecent loads the transformed data into the Info Mart database.

Table 3: Results of Job_LoadRecent, by Category of Data

Category of Data	Data Source	Target Info Mart Database Tables	
		Dimension Tables	Intraday Fact Tables
Contact center configuration	IDB	Loads new and changed data into the following dimension tables: • CALLING_LIST • CAMPAIGN • GROUP_ • PLACE • RESOURCE_ • SKILL • TENANT • TIME_RANGE • TIME_ZONE	• Loads new configuration facts (such as resource group fact, resource group combination fact, place group fact, resource skill fact, campaign group to campaign fact, and calling list to campaign fact) into the intraday fact tables. • Updates the active flags, end times, and durations of the active facts that it loaded previously into the intraday fact tables.
Note: Job_LoadRecent can optionally calculate or recalculate data for the intraday aggregate tables to reflect the most recent changes made to the underlying fact tables.			

Table 3: Results of Job_LoadRecent, by Category of Data (Continued)

Category of Data	Data Source	Target Info Mart Database Tables	
		Dimension Tables	Intraday Fact Tables
Voice and Multimedia interactions	IDBs	Loads new data into the following dimension tables: • CUSTOMER • INTERACTION_DESCRIPTOR • MEDIA_TYPE (Multimedia only) • REQUESTED_SKILL • REQUESTED_SKILL_COMBINATION • RESOURCE_GROUP_COMBINATION • ROUTING_TARGET • STOP_ACTION (Multimedia only) • STRATEGY • USER_DATA • USER_DATA_2	• Loads completed voice interaction facts and their associated interaction segment facts, interaction resource facts, and interaction resource state facts (including their corresponding media-specific extensions) into the intraday fact tables. Loads new (active and completed) Multimedia interaction facts and their associated interaction segment facts (including their corresponding media-specific extensions) into the intraday fact tables. • Loads completed voice ACD (automated call distribution) queue and virtual queue segment facts into the intraday fact table. Loads new (active and completed) Multimedia virtual queue segment facts into the intraday fact table.
Voice and Multimedia interactions (continued)	IDBs		• Updates the active flags, end times, durations, and other information of the active facts that it loaded previously into the intraday fact tables. Also updates completed facts with late-arriving information.
Resource activity	IDB and Stat Server database (in legacy reporting environments only)	Loads new data into the RESOURCE_STATE_REASON dimension table.	• Loads completed resource state facts and resource state reason facts into the intraday fact tables. Active resource state facts and resource state reason facts are not loaded. • Loads new (active and completed) resource session and DND facts into the intraday fact tables. • Updates the active flags, end times, and durations of the active facts that it loaded previously into the intraday fact tables.
Note: Job_LoadRecent can optionally calculate or recalculate data for the intraday aggregate tables to reflect the most recent changes made to the underlying fact tables.			

Table 3: Results of Job_LoadRecent, by Category of Data (Continued)

Category of Data	Data Source	Target Info Mart Database Tables	
		Dimension Tables	Intraday Fact Tables
Outbound Contact	IDB	Loads new data into the `RECORD_FIELD_GROUP_1` and `RECORD_FIELD_GROUP_2` dimension tables.	• Loads new calling list metric facts into the intraday fact tables. These facts are not updated. • Loads new (active and completed) campaign group session and campaign group state facts into the intraday fact tables. • Loads completed contact attempt facts into the intraday fact table. Active facts are not loaded. • Updates the active flags, end times and durations of the active facts that it loaded previously into the intraday fact tables. Also updates completed facts with late-arriving information.
GVP Voice Application Reporter (VAR)	GVP VAR database	Loads new data into the following resource dimension tables: • `GVP_APPLICATION` • `GVP_SUBCALL_FLOW` • `GVP_VOICE_MEDIA_SERVER` • `GVP_WEB_APPL_SERVER`	• Loads completed GVP call and GVP subcall facts into the intraday fact tables. Active GVP call and GVP subcall facts are not loaded. • Updates completed facts with late-arriving information.
Note: `Job_LoadRecent` can optionally calculate or recalculate data for the intraday aggregate tables to reflect the most recent changes made to the underlying fact tables.			

Table 3: Results of Job_LoadRecent, by Category of Data (Continued)

Category of Data	Data Source	Target Info Mart Database Tables	
		Dimension Tables	Intraday Fact Tables
General context		• Loads the TENANT_DATE dimension with complete calendar years from January 1, 1995, through December 31, 2024 for each tenant. • Updates FISCAL_ period columns for complete fiscal years, using the configured tenant-fiscal-periods options. For partial fiscal years at the beginning and the end of the dimension, it sets FISCAL_ period columns to null. • Updates the AUDIT_ dimension.	
Note: Job_LoadRecent can optionally calculate or recalculate data for the intraday aggregate tables to reflect the most recent changes made to the underlying fact tables.			

Notes:

- Fiscal period definitions contain many variations. Each tenant has its own rules for defining fiscal periods, which change from year to year and contain anomalies to account for calendar-year variations, such as leap year and the day of the week that begins the year.

  The option values that you select in tenant-fiscal-periods (or in the gim-tenant-fiscal-periods section under the tenant's Annex tab), and the TENANT_DATE fiscal period column values that Genesys Info Mart populates, are only suggested values. If you want to use fiscal periods for reporting, Genesys strongly recommends that you carefully analyze the TENANT_DATE column values and customize them to suit your environment. Be sure to customize only the fiscal period columns (that is, those columns that begin with FISCAL_).

- Job_LoadRecent updates the TENANT_DATE fiscal period columns each time that you change the option values in the Genesys Info Mart application's tenant-fiscal-periods section (or in the gim-tenant-fiscal-periods section under the tenant's Annex tab).

  If you change the option values after you have customized the fiscal period columns to suit your environment, Job_LoadRecent will overwrite your customized column values.

Job_LoadGIM

This job moves data from each intraday fact table to its counterpart historical fact table.

- For the fact tables that support only completed facts, Job_LoadGIM inserts new rows that represent those completed facts into the historical fact tables.
- For the fact tables that support both active and completed facts, Job_LoadGIM inserts new rows that represent those active or completed facts that it did not insert previously into the historical fact tables.
- Job_LoadGIM also updates the active flags, end times, durations, and other information of the active facts that it loaded previously into the historical fact tables.
- Finally, Job_LoadGIM updates completed facts with late-arriving information.

Job_LoadGIM loads fact data into the historical fact tables in several groups: configuration facts, interaction facts, resource (agent) facts, Outbound Contact facts, and GVP facts.

Multimedia interactions generally take much longer to complete than interactions for voice media type, and therefore will require updates across several ETL cycles. For this reason, Job_LoadGIM does not move them from the intraday fact tables to the historical fact tables until they are complete.

When Job_LoadGIM finishes loading the historical fact tables, it deletes the corresponding data from the intraday fact and aggregate tables.

Regardless of whether you plan to use intraday loading, configure the time of day that you want Genesys Info Mart Server to launch Job_LoadGIM. This is generally once a day after the final intraday ETL cycle has completed and after Job_LoadRecent has run.

Optimizing the Transaction Size

You can configure Genesys Info Mart to control the size of database transactions for Job_LoadGIM. For this job, the transaction size is defined as a

number of rows to be loaded to a single table in a single transaction. You need to configure the following option in the section [gim-etl]:

- load-transaction-size

This option applies only to the fact tables that are loaded by Job_LoadGIM. Finding an optimal setting for this option for your environment should help you handle very large datasets without requiring excessive database resources. The following considerations should be taken into account:

- Setting the option value to 0 (zero) imposes no limits on a transaction size. A single transaction might be the most efficient method of updating a table, but it does raise certain operation issues when the data size becomes very large (gigabytes).
- Setting the option to a low value results in multiple commits being required to perform updates to large tables, thus causing the loading process to take longer.

Job_AggregateGIM

This job 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.

Job_AggregateGIM calculates day- and month-level aggregates. The job does not calculate or update information in the intraday aggregate tables (tables named R_AG_*_HOUR). Job_LoadRecent calculates intraday aggregate tables.

Aggregation Types

From a technical perspective, Genesys Info Mart provides two types of aggregates:

- Disposition-based, when aggregation is performed by assigning fact counts and entire durations to the interval in which the fact started.
- Interval-based, when aggregation is performed by apportioning fact counts and durations to each interval in which the fact occurred.

From a reporting application perspective, Genesys Info Mart provides separate out-of-box aggregates for the following Genesys applications:

- CCPulse+ (These aggregates were first introduced in Genesys Info Mart 7.2.)
- Genesys Interactive Insights (GI2) (These aggregates are first introduced in Genesys Info Mart 7.6.)

Your custom reporting applications may also use these out-of-box aggregates. This guide refers to either the technical or the reporting application perspective when it describes the functionality of the out-of-box aggregates.

Aggregation Levels

Disposition-Based Aggregates

Disposition-based aggregates are upwardly additive from one level to the next. Each level of aggregation is based only on the previous level. For example—day aggregates are based on hour aggregates, and month aggregates are based on day aggregates. This means:

- Each fact is assigned to one and only one time interval. The interval assignment is based on the standard tenant start time of the fact.
- The entire fact, including all durations, is completely counted in the interval in which the fact started. In other words, the fact is counted only in the interval in which it started and not in any other intervals that it might span.

Interval-Based Aggregates

Interval-based aggregates are not additive from one level to the next. Interval-based aggregates measure facts about objects in each interval in which the facts occur. These aggregates capture only the part of the fact that fits in the interval. For this reason, interval-based aggregates are populated from details, and are not based on the previous level.

The following aggregate tables contain interval-based aggregates:

- `AG2_INB_V_I_IXN_AGENT_*`
- `AG2_INB_V_I_SESS_STATE_*`
- `AG2_INB_V_I_STATE_RSN_*`

Storage

For each aggregation level, both disposition-based and interval-based aggregates store only additive aggregates (such as `SUM` and `COUNT`) or aggregates that are derivable from a lower level (such as `MIN` and `MAX`). Non-additive aggregates such as `AVG` or `RATIO` are not explicitly stored, but can be derived.

The Info Mart database supplies separate physical tables and views for various aggregate levels, as follows:

- For disposition-based aggregates:
  - Tables provide hour, day, and month aggregates.
  - Views provide subhour, week, quarter, and year aggregates.

> **Note:** Genesys Info Mart provides subhour aggregates for GI2 out-of-box aggregates only.

- For interval-based aggregates:
  - Tables provide subhour, hour, and day aggregates.
  - No views are provided.

> **Note:** Genesys Info Mart provides interval-based aggregates for GI2 out-of-box aggregates only.

Aggregation Intervals

Job_AggregateGIM does not perform incremental aggregation. Instead, it recalculates all the intervals for all the aggregate levels that are affected by new or changed fact data. For example—if fact data changed for the hour that represents the interval between 11:00 and 11:59 on June 1, 2006, the following aggregate intervals are reloaded:

- 30 min Subhour – 11:00 - 11:29, June 1, 2006
- 30 min Subhour – 11:30 - 11:59, June 1, 2006
- Hour – 11:00, June 1, 2006
- Day – June 1, 2006
- Month – June, 2006

Starting with Genesys Info Mart release 7.6.006, you can improve ETL performance by limiting how far back in time Job_AggregateGIM will re-aggregate when it encounters newly loaded, late-arriving facts. For more details about configuring this functionality, see the information about the max-late-arriving-fact-time-limit option in the *Genesys Info Mart 7.6 Deployment Guide.*

Agent Group Aggregates

When calculating agent group aggregates, Job_AggregateGIM considers that agent group membership can change during the aggregation interval.

With CCPulse+ out-of-box aggregates for agent groups, Job_AggregateGIM assigns each interaction to the appropriate hour aggregate interval by using the standard tenant start time for each interaction. Using the agent group membership information that was in effect when each interaction started (from RESOURCE_GROUP_FACT), the job performs aggregation along the Group dimension. The standard tenant start time is used as the time reference.

- If an agent that handled the interaction was a member of multiple agent groups, the interaction is included in each agent group's aggregate.
- If an interaction was handled by multiple agents that were members of the same agent group, the interaction is included only once in that agent group's aggregate.
- If an interaction was handled by multiple agents that were members of different agent groups, the interaction is included once in each distinct agent group's aggregate.

The agent group aggregates are accurate at each aggregation level because Job_AggregateGIM uses the agent group membership details that were in effect at the time each interaction occurred, rather than using those that were in effect at the beginning of the aggregation interval.

With GI2 out-of-box aggregates, the RESOURCE_GROUP_COMBINATION dimension contains the group membership information for an agent (or a queue) resource

at the time when the fact started. Job_AggregateGIM performs the aggregation based on this information. When an aggregate combines information from multiple fact tables, where each fact table has a different representation of the group membership, Job_AggregateGIM uses the RESOURCE_GROUP_COMBINATION of the lowest-grain fact table. For example—when data from the INTERACTION_RESOURCE_FACT table is combined with the data from the SM_RES_STATE_FACT table, the job uses the group membership from the RESOURCE_GROUP_COMBINATION dimension of SM_RES_STATE_FACT.

Scheduling

Job_AggregateGIM is an optional job. Configure the time of day that you want Genesys Info Mart Server to launch Job_AggregateGIM, if you plan to do one or more of the following:

- Use CCPulse+ reports that read data from the Info Mart database.
- Use GI2 reports.
- Build custom reports using the aggregate tables that are provided in the Info Mart database.

The good practice is to run Job_AggregateGIM once a day after Job_LoadGIM completes.

If you build custom reports that do *not* use the aggregate tables that are provided in the Info Mart database, configure Genesys Info Mart Server so that Job_AggregateGIM does not launch.

Calculating Aggregates Within a Time Span

You can use input parameters to specify the time span for which you want Job_AggregateGIM to calculate or recalculate the aggregate tables. In order to pass these input parameters to Job_AggregateGIM, you must use Genesys Info Mart Administration Console. For more information, see "Managing ETL Jobs with Genesys Info Mart Administration Console" on page 74. For information about specifying input parameters to Job_AggregateGIM, see "Specifying Input Parameters to Job_AggregateGIM" on page 56.

Consider specifying a time span for any of the following situations:

a. You run Job_LoadGIM to load fact table data for an extended period of time without running Job_AggregateGIM to load aggregate tables (or if you have never run Job_AggregateGIM). See "Enabling or Re-Enabling Aggregation" on page 54 for the detailed procedure you should follow.

b. You run Job_AggregateGIM on a regular basis, but you are enabling some particular aggregate for the first time, or re-enabling some particular aggregate after it has not been loaded for an extended period of time. See "Enabling or Re-Enabling Particular Aggregates" on page 54 for the detailed procedure you should follow.

c. You discover an error in the aggregates and need to recalculate them after installing a Genesys Info Mart software update. See "Recalculating Aggregates" on page 55 for the detailed procedure you should follow.

Enabling or Re-Enabling Aggregation

It is important to populate the aggregates for only the most recently loaded data. This will allow `Job_AggregateGIM` to complete in a reasonable amount of time, and will allow future ETL cycles to run normally and produce aggregates for any newly loaded or updated facts. The following procedure will direct `Job_AggregateGIM` to populate the aggregates for only the most recently loaded data, and will prevent a single run of `Job_AggregateGIM` from populating aggregates for all the fact data that has been loaded since the last time aggregates were produced.

Plan to use this procedure during a daily maintenance window before you enable intraday aggregation (by setting [`scheduler`] `populate-intraday-aggregates=TRUE`) or historical aggregation (by setting [`schedule`] `run-aggregates=TRUE`):

1. Set [`schedule`] `run-scheduler=FALSE` to prevent the scheduler from starting any ETL jobs.
2. Using Genesys Info Mart Administration Console, run a final ETL cycle (all extract jobs, `Job_TransformGIM`, and `Job_LoadRecent`).
3. Using Genesys Info Mart Administration Console, run `Job_LoadGIM`.
4. Using Genesys Info Mart Administration Console, run `Job_AggregateGIM`, and specify a date range that encompasses only the date of the most recently loaded data (that is, set both the `startDate` and `endDate` to the same date). This calculates, for the specified date only, all enabled aggregates for all aggregate-enabled tenants, including the newly enabled aggregate.
5. Set [`schedule`] `run-aggregates=TRUE` to enable scheduling of `Job_AggregateGIM`.
6. Set [`schedule`] `populate-intraday-aggregates=TRUE` if you want `Job_LoadRecent` to populate intraday aggregates.
7. Set [`schedule`] `run-scheduler=TRUE` allowing the scheduler to resume operation.
8. If you must calculate aggregates for reporting intervals earlier than the most recently loaded fact data, follow the procedure for "Recalculating Aggregates".

Enabling or Re-Enabling Particular Aggregates

For a newly enabled aggregate, it is important to populate aggregate data for only the most recently loaded data. This will allow `Job_AggregateGIM` to complete in a reasonable time, and will allow future ETL cycles to run normally and produce aggregates for any newly loaded or updated facts. The following procedure will direct `Job_AggregateGIM` to populate the newly enabled aggregate for only the most recently loaded data, and will prevent a

single run of Job_AggregateGIM from populating the newly enabled aggregate for all the available fact data.

Plan to use this procedure during a daily maintenance window before you enable some particular aggregate for the first time (or re-enable some aggregate that has been disabled for an extended period of time):

1. Set [schedule] run-scheduler=FALSE to prevent the scheduler from starting any ETL jobs.
2. Using Genesys Info Mart Administration Console, run a final ETL cycle (all extract jobs, Job_TransformGIM, and Job_LoadRecent).
3. Using Genesys Info Mart Administration Console, run Job_LoadGIM.
4. Enable the particular aggregate; for example, set [gim-aggregates-tenant] populate-ixn-service-type-aggregates=TRUE.
5. Using Genesys Info Mart Administration Console, run Job_AggregateGIM, and specify a date range that encompasses only the date of the most recently loaded data (that is, set both the startDate and endDate to the same date). This will recalculate, for the specified date only, all enabled aggregates for all aggregate-enabled tenants, including the newly enabled aggregate.
6. Set [schedule] run-scheduler=TRUE allowing the scheduler to resume operation.
7. If you must calculate the newly-enabled aggregate for reporting intervals earlier than the most recently loaded fact data, follow the procedure for "Recalculating Aggregates".

Recalculating Aggregates

Repeat this process over one or more daily maintenance windows to calculate the aggregates for the desired time span, using date ranges that allow Job_AggregateGIM to complete in a reasonable time, without interfering with the operation of the next day's intraday ETL cycles.

1. Set [schedule] run-scheduler=FALSE to prevent the scheduler from starting any ETL jobs.
2. Using Genesys Info Mart Administration Console, run a final ETL cycle (all extract jobs, Job_TransformGIM, and Job_LoadRecent).
3. Using Genesys Info Mart Administration Console, run Job_LoadGIM.
4. Using Genesys Info Mart Administration Console, run Job_AggregateGIM, specifying a reasonable date range. This will recalculate, for the specified date range, all enabled aggregates for all aggregate-enabled tenants. The date range that will complete in a reasonable time depends on the performance of your Info Mart database. To find a reasonable date range for your environment, it is a good practice to start with a very small date range, such as one day, and increase it over subsequent runs.
5. Repeat Step 4 as time permits within your maintenance window.

6. Set [schedule] run-scheduler=TRUE allowing the scheduler to resume its operation.

Specifying Input Parameters to Job_AggregateGIM

In order to specify an aggregation time span, you must use Genesys Info Mart Administration Console to run (or schedule) Job_AggregateGIM. Follow the appropriate procedure to run the job immediately or to schedule the job. For more information, see "Executing and Scheduling ETL Jobs" on page 76".

Before you click Run to execute or schedule the job, you may supply values for the following optional parameters:

- -tenantKey—the surrogate key (TENANT_KEY) from the TENANT dimension table that represents the tenant for which aggregates will be calculated.
- -startDate—the start of the time span. This parameter has the format YYYYMMDD.
- -endDate—the end of the time span. This parameter has the format YYYYMMDD. Make sure endDate is equal to or greater than startDate.

For more information about how to enter ETL job parameters, see Step 6 of the procedure to run an ETL job immediately (page 79).

Notes:

- If you do supply a value for the optional -tenantKey parameter, then you must also supply values for the -startDate and-endDate parameters.
- If you do not supply a value for the optional -tenantKey parameter, aggregates are calculated (or recalculated) for all tenants for which aggregation is enabled.
- If you do not supply a value for the optional -startDate and -endDate parameters, aggregates are calculated (or recalculated) for all time spans that include new or changed data. Depending on the volume of new and changed data, Job_AggregateGIM might run for a very long time. Genesys recommends supplying values for these parameters if you are enabling or re-enabling aggregation, and have been loading new or changed facts for a significant time span.
- The time range that you specify when running the job from the Genesys Info Mart Administration Console overrides any value you may have set for the max-late-arriving-fact-time-limit option, which specifies how far back in time Job_AggregateGIM will routinely re-aggregate when it encounters newly loaded, late-arriving facts.

Optimizing the Transaction Size

You can configure Genesys Info Mart to control the size of database transactions for `Job_AggregateGIM`. For this job, the transaction size is defined as a time range for the data that is aggregated by the aggregation queries in a single transaction. You need to configure the following options in the section `[gim-etl]`; these options indirectly control the transaction sizes of the aggregation queries:

- `aggregate-time-range-units`
- `aggregate-time-range-limit`

These options apply only to the aggregation queries that populate aggregation tables with GI2 out-of-box aggregates; the options do not apply to CCPulse+ out-of-box aggregates or to any aggregates that are presented through database views.

Finding optimal settings for these options for your environment should help you balance aggregate execution time against the RDBMS resources that are required to produce the aggregates:

- Setting the `aggregate-time-range-limit` option to a value of `0` (zero) imposes no limits on a transaction size. Users who have sufficient database resources can set this time range to a large value (or to `0`).
- Setting the `aggregate-time-range-limit` option to a low value results in multiple runs of the performance-intensive aggregation queries that are required to aggregate the eligible data, and this causes the aggregation process to take longer.

Job_MaintainGIM

This job performs the following tasks:

- Identifies old fact table rows in the Genesys Info Mart database that are eligible to be purged.
- Identifies old aggregation table rows in the Genesys Info Mart database that are eligible to be purged.
- Depending on the way that the Genesys Info Mart application has been configured, either deletes the eligible table rows or flags them as eligible to be purged.
- Purges the historical information about steps that the ETL has performed related to job execution, source data extraction, target table loading, table purging, and data aggregation from the Staging Area database.

Retention Period

The Genesys Info Mart configuration specifies the number of days that facts and aggregates are retained in the Info Mart database. Facts and aggregate time intervals that occurred earlier than the specified retention period are eligible to be purged from the Info Mart database. You can configure separate configuration options for facts and aggregates; most likely, aggregate data needs to be retained longer than the underlying facts. In a multi-tenant deployment, you can configure separate configuration options for each tenant. This allows each tenant to specify its own data retention period, while ensuring consistency among fact tables.

The configuration option 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.

An additional configuration option specifies how long the Info Mart database will retain fact data for tenants removed from the Genesys Info Mart configuration.

Note: For more information about the retention period configuration options, see the chapter about the Genesys Info Mart application configuration in the *Genesys Info Mart 7.6 Deployment Guide*.

Purging Rules

The following rules apply to purging:

- Facts are evaluated for purging based on the date portion of their standard tenant end time.
- Active facts are not eligible for purging (see Table 4 on page 84).
- To maintain consistency among fact tables, related fact table rows within the same fact table, as well as across fact tables, are treated as a unit and are eligible only when all parts are eligible to be purged. This means that all related facts must end earlier than the data-retention period in order to be eligible. The following are examples of related fact tables:
  - `INTERACTION_SEGMENT_FACT` to `INTERACTION_FACT`
  - `CAMPAIGN_GROUP_STATE_FACT` to `CAMPAIGN_GROUP_SESSION_FACT`
  - `CONTACT_ATTEMPT_FACT` to `CAMPAIGN_GROUP_SESSION_FACT`
  - `CONTACT_ATTEMPT_FACT` to `INTERACTION_FACT`
  - `CALLING_LIST_METRIC_FACT` to `CAMPAIGN_GROUP_SESSION_FACT`

  - GVP_SUBCALL_FACT to GVP_CALL_FACT
  - GVP_CALL_FACT to INTERACTION_FACT

  For more information, see "Purging the Genesys Info Mart Database" on page 94.
- GI2 out-of-box aggregates (whose table names start with AG2_INB_V_) and outbound aggregates (whose table names start with AG2_OUT_V_) are evaluated for purging based on the standard tenant date (DATE_TIME dimension), which represents the start of the aggregate interval.
- CCPulse+ out-of-box aggregates are evaluated for purging based on their standard tenant date (TENANT_DATE dimension), which represents the start of the aggregate interval.
- The two basic types of aggregates (interval-based and disposition-based) are purged independently; each time level that is kept in a physical table, instead of in a view, is purged independently.
- Dimension table rows are not eligible to be purged.

Deleting or Flagging Rows

When you configure the Genesys Info Mart application, you choose whether you want Job_MaintainGIM to delete old fact and aggregate table rows or just to flag them as eligible to be purged.

- When you configure Job_MaintainGIM to delete old fact and aggregate table rows, the job issues simple 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 you configure Job_MaintainGIM to flag fact and aggregate table rows that are eligible to be purged, the job sets the PURGE_FLAG to 1 in those table rows. Flagging table rows, instead of deleting them, allows a data-warehousing specialist to control when and how to delete the data most efficiently.

Job_MaintainGIM does not delete or flag any rows from dimension tables, as these might be referenced by external fact tables and aggregate tables.

Scheduling

Job_MaintainGIM is an optional job. If you do not want to purge data, or if you want to implement your own purging, configure Genesys Info Mart server not to launch Job_MaintainGIM. If you do want to use Job_MaintainGIM to purge data, configure the time of day you want Genesys Info Mart Server to launch this job. This is generally once a day, after Job_AggregateGIM completes.

Optimizing the Transaction Size

You can configure Genesys Info Mart to control the size of database transactions for `Job_MaintainGIM`. For this job, the transaction size is defined as a time range for the data that is updated in, or deleted from, a given table in a single transaction. You need to configure the following options in the section `[gim-etl]`; these options indirectly control the transaction sizes of the aggregation queries:

- `maintain-time-range-units`
- `maintain-time-range-limit`

These options define the time range of fact data that is operated on by a single purging query.

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 facts and aggregates:

- Setting the `maintain-time-range-limit` option to a value of `0` (zero) imposes no limits on a transaction size. Users who have sufficient database resources can set this time range to a large value (or to `0`).
- Setting the `maintain-time-range-limit` option to a low value results in multiple runs of the purging queries, and this causes the purge process to take longer.

Job_MigrateGIM

This job migrates the data from the Staging Area and Info Mart databases of release 7.5 to release 7.6.

During the data migration process, the job:

- Controls how much data is updated and committed to the database.
- Tracks the progress in the special administrative tables.

Depending on the amount of data that has been accumulated by a previous release of Genesys Info Mart, and the time that you have to complete the migration, you have an option to:

- Run `Job_MigrateGIM` once, until all data is migrated.
- Schedule `Job_MigrateGIM` to run outside of the intraday ETL cycle, to update data in small chunks over a period of time.

You can either configure the migration job at Genesys Info Mart application level or run it from Genesys Info Mart Administration Console.

Note: In a migrating environment, Genesys Info Mart Server ensures that no job can run until the data migration of the critical tables has been completed.

Prerequisites for the Migration Job

Run the migration SQL scripts prior to starting `Job_MigrateGIM`. These scripts create and alter the database tables as required for the Genesys Info Mart release 7.6. Refer to Genesys Info Mart section of the *Genesys Migration Guide* for instructions.

Order of Data Migration

The migration job uses a two-fold approach to the order in which to migrate data. First, the job evaluates the table criticality; then, the job evaluates the data time.

Critical Tables

The migration job identifies critical tables—that is, the tables that are considered necessary to migrate before any other job can run. A table, whether it belongs to the Staging Area or Info Mart database, is critical for migration when any job is dependent on the data in the table being migrated before this job can run. A table is non-critical if no jobs are dependent on the data in the table that is being migrated. The migration job completes the data migration of the critical tables before it migrates the data from non-critical tables.

Note: A decision whether a table is critical or non-critical is based on what type of update is being made to the table in release 7.6. For example—if a job is dependent on a new table column being initialized (in other words, not set to null), the migration job identifies this table as critical.

The following tables are considered critical when you migrate from release 7.5 to release 7.6:

- `R_CHAT_IXN_FACT_EXT`
- `R_CHAT_SEG_FACT_EXT`
- `R_EMAIL_IXN_FACT_EXT`
- `R_EMAIL_SEG_FACT_EXT`

Data Time

For both critical and non-critical tables, the most recent data is migrated first. The migration job orchestrates data migration across multiple tables, to ensure that data migration is completed for each table for the current (most recent) data time range, before the job starts to process the data from the next (less recent) data time range.

Scheduling

You can configure Genesys Info Mart so that `Job_MigrateGIM` starts at a certain time and runs either on a daily basis, for a specific number of hours, or once, until all data has been migrated.

Genesys Info Mart Server schedules the migration job to run if the `run-migration` option is configured in the `[schedule]` section. The job runs at the time that is specified by the `migration-start-time` option, for the time period that is specified by the `migration-duration-in-hours` option, in the same

section. If the migration-duration-in-hours option is set to 0, the migration job runs until all data is migrated.

Note: Be sure to set the value for the migration-duration-in-hours option to be outside of the interval that is defined for the intraday ETL cycle by the etl-start-time and etl-end-time options. Genesys also recommends that you avoid running other daily jobs while Job_MigrateGIM runs, to prevent them from accessing the same database tables.

Performance Considerations

Three factors affect performance of the Job_MigrateGIM:

- The amount of time that the job is to run (configured in the migration-duration-in-hours option in the [schedule] section).
- Whether other jobs are running at the same time.
- The amount of data that is being committed to the database.

Transaction Size

The Job_MigrateGIM controls how much data is committed to the database in a single transaction during the updates, based on the values of two options in the [gim-etl] section:

- data-migration-time-range-units
- data-migration-time-range-limit

Finding optimal settings for these options for your environment should help you balance migration execution time against the RDBMS resources that are required to migrate the data:

- Setting the data-migration-time-range-limit option to a large value results in large database transactions. Users who have sufficient RDBMS resources can choose a relatively large value to allow the data migration to complete more quickly.
- Setting the data-migration-time-range-limit option to a small value results in smaller database transactions. Users who have more modest RDBMS resources can choose a relatively small value, which permits the data migration to complete over many executions of Job_MigrateGIM.

For single-tenant deployments, Job_MigrateGIM determines the time range of data that is to be migrated, based on the standard tenant date and time-of-day dimensions, which are indexed for better performance. For multi-tenant deployments, Job_MigrateGIM does not issue a separate query per tenant, but uses the standard enterprise date and time-of-day dimensions to ensure that the job selects a uniform slice of data to be migrated across all tenants.

Note, however, that for the tables (such as dimension tables) that do not have the date and time-of-day dimension keys, all data that is updated in each table is committed in one database transaction.

Tracking the Progress

The `Job_MigrateGIM` uses the `STG_DATA_MIGRATION` and `GIM_DATA_MIGRATION` tables in the Staging Area and Genesys Info Mart databases, respectively, to track the time ranges of data that has already been migrated for the tables in each database. The job adds a record to the `STG_DATA_MIGRATION` or `GIM_DATA_MIGRATION` table when it commits data updates for a given table to a respective database.

The amount of data that has been migrated for every table is persisted, so that if the migration job is restarted, it continues the migration process where it left off.

If the migration job encounters a database error during data migration, Genesys Info Mart Server generates a log event. You can set up an alarm condition to receive a notification of such an error. (Refer to the *Framework 7.6 Solution Control Interface Help* for instructions on how to set alarm conditions.) When you restart the job after correcting the database error, the migration job will continue data migration where it left off.

Chapter

3 Working with ETL Jobs

The Genesys Info Mart has ETL jobs that, when executed, extract data from your source databases and load the transformed data into the Info Mart database. For more information about the ETL jobs, see “Understanding Genesys Info Mart ETL Jobs” on page 27.

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

This chapter provides information on the following:

Note: Before you can execute any ETL job, you must complete the tasks in the *Genesys Info Mart 7.6 Deployment Guide*.

Launching ETL Jobs with Genesys Info Mart Server

The Genesys Info Mart Server is the component that launches the ETL jobs based on the schedule options that you configure in the Genesys Info Mart application.

About the Intraday ETL Cycle

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

- Extracts data from each data source.
- Transforms the extracted data.
- Loads the transformed data into the dimension tables and the intraday fact and aggregate tables.

Note: Intraday loading is optional. If you want to load Info Mart data only once a day, you can postpone this step until after the final intraday ETL cycle has occurred.

Intraday ETL Cycle Configuration Options

You configure options that control the intraday ETL cycle, including:

- Time of day the first intraday ETL cycle should begin.
- Time of day the final intraday ETL cycle should begin.
- Frequency of the intraday ETL cycle.
- Whether loading of dimension tables, intraday fact tables, and calculation of intraday aggregates occurs as part of the intraday ETL cycle.

You configure additional options to specify the time of day when:

- Data moves from the intraday fact tables to the historical fact tables.
- Historical aggregate tables are updated based on the new and changed fact table data.
- Old fact table data is purged.

You can also set configuration options to:

- Temporarily stop Genesys Info Mart Server from launching ETL 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 fact tables.
- Stop Genesys Info Mart from launching the job that migrates fact table data from a previous release.

About the Genesys Info Mart Server

This section provides the information about the following topics:

Job Interdependencies

Genesys Info Mart Server takes into account job interdependencies. It does not start a job that is predetermined to fail because of an error in a preceding job. Genesys Info Mart Server also does not start a job if there is another instance of that job already running. Genesys Info Mart Server monitors job status tables in the Staging Area database to determine whether jobs are running or have failed.

Note: In a migrating environment Genesys Info Mart Server does not start:

- Any job until the SQL migration scripts run successfully.
- Any other job until `Job_MigrateGIM` completes the data migration for the critical tables.

Using the Solution Control Interface

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

Starting Genesys Info Mart Server

When you start Genesys Info Mart Server by using SCI, the Genesys Info Mart ETL application status reflects the status of the Genesys Info Mart Server itself, instead of the status of any ETL jobs. In other words, the `STARTED` status that is reported by SCS indicates that Genesys Info Mart Server is operational, but it does not indicate whether ETL 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 it 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 ETL 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 ETL 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 will automatically launch Genesys Info Mart. For more information on Windows Service, see the section about installing the Genesys Info Mart component in the *Genesys Info Mart 7.6 Deployment Guide.*

Connecting to Configuration Server

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

If Genesys Info Mart Server cannot connect to Configuration Server upon startup, Genesys Info Mart Server reads the values of configuration options previously stored in a local file and reattempts 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 is already connected, the Genesys Info Mart Server will exit 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. Should Genesys Info Mart Server then later lose its connection to Configuration Server, Genesys Info Mart Server repeats the following process until a connection is established. Genesys Info Mart Server:

- Attempts to connect to Configuration Server.
- Attempts to connect to the backup Configuration Server, if configured.
- Waits 30 seconds specified for the reconnect timeout.

Setting Scheduling Options for Genesys Info Mart Server

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

ETL schedules are defined in 24-hour time spans with 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 ETL jobs, see "ETL Job Summary Table" on page 28. For information on how the jobs transform data, job interdependencies, and sample schedules, see "About the Intraday ETL Cycle" on page 83.

To set the ETL schedule:

1. From the Configuration Manager, navigate to the `Options` tab of the Genesys Info Mart application.
2. Navigate to the `[schedule]` section.
3. Enter a value for `etl-start-time` to specify the time of day the first ETL cycle begins.
4. Enter a value for `etl-end-time` to specify the time of day the final intraday ETL cycle begins. This should be a time of day when no other intraday ETL cycles will begin.
5. Enter a value for `etl-frequency` to specify the number of minutes between the start times of adjacent intraday 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.

   Genesys Info Mart Server extracts data from Stat Server (in legacy environments only), GVP VAR, and IDB. Extractions for all instances of these databases are done concurrently.
6. Enter a value for `max-concurrent-extract-jobs`. This value specifies how many extraction jobs launch at one time. For example—if you have more than five data sources and you set `max-concurrent-extract-jobs` to `5`, Genesys Info Mart Server launches five extraction jobs. When one extraction job finishes, Genesys Info Mart Server launches another extraction job to maintain the value of five. Once all the extractions complete, transformation begins.

Note: Genesys Info Mart Server launches a separate instance of `Job_ExtractICON` for each of the extraction roles that you configured in the database access point (DAP). Make sure to take this into consideration when configuring `max-concurrent-extract-jobs`. For more information about configuring extraction roles, see the chapter about configuring DAPs for Genesys Info Mart in the *Genesys Info Mart 7.6 Deployment Guide*.

7. If you plan to use intraday loading, set `load-recent-with-extract-and-transform` to `TRUE`. This option specifies whether the Genesys Info Mart Server launches the job, `Job_LoadRecent`, that loads data into dimension tables and intraday fact tables during each intraday ETL cycle.

   If you plan to use intraday loading and intraday aggregates, set `populate-intraday-aggregates` to `TRUE`. To improve ETL performance, consider setting the `intraday-aggregates-frequency` option to a value greater than the `etl-frequency` option. Otherwise, if you have configured the ETL to populate intraday aggregates, intraday aggregation will be performed every time `Job_LoadRecent` runs, and the repeated re-aggregation of overlapping time ranges degrades job performance.

   Note: If intraday loading is configured, `Job_LoadRecent` may run while extract jobs are running. `Job_TransformGIM` will not run until all extract jobs have completed and all load steps for `Job_LoadRecent` have completed successfully. `Job_TransformGIM` can start and run in parallel with the intraday aggregation step of `Job_LoadRecent`.

8. If you do *not* plan to use intraday loading:
   a. Set `load-recent-with-extract-and-transform` to `FALSE`.
   b. Set `load-recent-start-time` to the time of day that you want Genesys Info Mart to launch the job that loads the dimension tables and the intraday fact tables. The time of day must be outside the range that is specified by `etl-start-time` and `etl-end-time`. `Job_LoadRecent` will not start until the `load-recent-start-time` has been reached and all jobs running in the ETL cycle have completed successfully.
9. Enter a value for `load-start-time`. This value specifies the time of day that you want Genesys Info Mart Server to launch the job, `Job_LoadGIM`. This job moves data from the intraday fact tables to the historical fact tables. The time of day must be outside the range that is specified by `etl-start-time` and `etl-end-time`. `Job_LoadGIM` will not start until the `load-start-time` has been reached and `Job_LoadRecent` has completed successfully.

10. 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-start-time` to the time of day that you want Genesys Info Mart to launch `Job_AggregateGIM`. The time of day must be outside the range that is specified by `etl-start-time`, `etl-end-time`, and after `load-start-time`. `Job_AggregateGIM` will not start until the `aggregate-start-time` has been reached and `Job_LoadGIM` has completed successfully.
11. If you do *not* plan to use the Info Mart aggregate tables, set `run-aggregates` to `FALSE`. This option specifies whether the Genesys Info Mart Server launches the job, `Job_AggregateGIM`, that loads the aggregate tables based on newly added or changed fact table data.
12. If you plan to purge old data from fact and aggregate tables:
    a. Set `run-maintain` to `TRUE`. This option specifies whether Genesys Info Mart Server launches the job, `Job_MaintainGIM`, that purges old fact and aggregate table data.
    b. 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`, and after `load-start-time` and `aggregate-start-time`. `Job_MaintainGIM` will not start until the `maintain-start-time` has been reached, and `Job_LoadGIM` has completed and (if configured to run) `Job_AggregateGIM` has completed successfully.
13. If you do *not* plan to purge old data from fact or aggregate tables, set `run-maintain` to `FALSE`. This option specifies whether Genesys Info Mart Server launches the job, `Job_MaintainGIM`, that purges old fact and aggregate table data.
14. To start or resume the Genesys Info Mart Server schedule, set `run-scheduler` to `TRUE`. You can set this option to `FALSE` to temporarily stop Genesys Info Mart Server from launching ETL jobs.

Automatically Retrying Failed ETL Jobs with Genesys Info Mart Server

Genesys Info Mart can automatically retry failed ETL jobs. This allows the Genesys Info Mart Server to recover from many types of failures, such as temporary network outages, temporary database outages, and SQL Server deadlocks. If the problem occurs within the context of a running job, Genesys Info Mart Server will retry the job. This automatic retry will only occur when the job is run from Genesys Info Mart Server scheduler and not from the Genesys Info Mart Administration Console. If the problem occurs with Genesys Info Mart Server, but does not occur within the context of a failed job, Genesys Info Mart Server will clean up and retry the failed operation.

Note: Genesys Info Mart Server will not automatically retry jobs that failed before Genesys Info Mart Server was restarted. If a job fails, and you restart Genesys Info Mart Server, you must use the Genesys Info Mart Administration Console to manually rerun the failed job.

Configuring Automatic Retry of ETL Jobs

Two new configuration options have been added to the Genesys Info Mart ETL application to allow the automatic retry of failed jobs:

- `job-retry-count`—Determines how many times to retry a failed job started by Genesys Info Mart Server.
- `job-retry-wait`—Determines the number of minutes to wait before retrying a failed job started by Genesys Info Mart Server.

For more information about setting these configuration options, see the *Genesys Info Mart 7.6 Deployment Guide*.

Automatically Retrying Failed ETL Jobs

When Genesys Info Mart Server detects a failed job that was started by the scheduler, and `job-retry-count and job-retry-wait` are configured, Genesys Info Mart Server retries the failed job by performing the following:

1. Waiting the amount of time configured in the `job-retry-wait` configuration option before proceeding to Step 2.
2. Releasing all connections to the Staging Area and Info Mart databases. If the failed job is an extract job, all connections to the source databases (IDBs) are also released.
3. Recreating all required database connections.
4. Restarting the failed job.

If the job fails again, Steps 1 to 4 on page 72 are repeated until the configured number of retries has been reached. During this retry phase, if Genesys Info Mart Server detects running jobs, (for example—a job started from the Genesys Info Mart Administration Console) it waits until all jobs are finished before continuing the retry process. While the Genesys Info Mart Server is waiting to retry a job, it does not start any scheduled jobs. This mainly affects the extract jobs and `Job_LoadRecent`. These jobs can run in parallel after `Job_TransformGIM` completes. Jobs can still be started from the Genesys Info Mart Administration Console, but these jobs must finish before Genesys Info Mart Server releases the database connections and retries the failed job. If a job started from Genesys Info Mart Administration Console is being retried and completes successfully, then Genesys Info Mart Server stops retrying that job.

Although automatic retry handles all types of failures, it cannot recover from, or correct data related problems, such as data value overflows, duplicate key or other database constraint violations. Retries of these types of failures result in another occurrence of the same failure. Recovery still requires the manual correction of incorrect data and the manual retry of the failed job.

Note: If the problem is not being solved by the automatic job retry, and you want Genesys Info Mart Server to stop the retry so you can manually correct the underlying issue, temporarily set the `run-scheduler` option in the `[schedule]` section to `FALSE`. Once you correct the underlying problem (for example—manually deleting a duplicate record that is causing a constraint violation), run the failed job manually from the Genesys Info Mart Administration Console. When it completes successfully, set the `run-scheduler` option in the `[schedule]` section to `TRUE` to resume normal job scheduling.

Handling Unresolved References to Configuration Objects

When `Job_TransformGIM` fails due to unresolved references to configuration objects, the retry of `Job_TransformGIM` is handled differently. When new objects, such as `DNs`, `AGENTS` and `PLACES`, are added to the contact center configuration and used immediately, the newly added objects are not present on the most recent run of `Job_ExtractICON` for `role=ICON_CFG`. However, these configuration objects are referenced by the interactions or agent states extracted by the most recent run of `Job_ExtractICON` for other `roles`, such as `ICON_CORE`, `ICON_MM` or `ICON_OCS`. If `Job_TransformGIM` cannot find a configuration object and the `ignore-missing-config-objs` option in the `[gim-transformation]` section is set to `FALSE`, then `Job_TransformGIM` fails.

Genesys Info Mart Server detects this failure and performs the following:

1. Without releasing database connections, it starts `Job_ExtractICON` for `role=ICON_CFG` to extract configuration data. No other extract jobs are started.

2. If the job extracts configuration data successfully, it restarts `Job_TransformGIM`.
3. If `Job_TransformGIM` fails again with missing configuration objects, Step 1 and Step 2 are repeated until the number of retries configured in `job-retry-count` has been reached.
4. If all retries of `Job_TransformGIM` fail because of missing configuration objects, the ICON IDB may require resynchronization with the Configuration Server database. Perform the procedure in Recovering from Unresolved References to Configuration Objects, page 130 to complete the recovery procedure.

Once `Job_TransformGIM` completes successfully, Genesys Info Mart Server returns to scheduling other jobs.

Managing ETL Jobs with Genesys Info Mart Administration Console

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

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

- Execute or schedule the ETL job that initializes the Staging Area and Info Mart databases.
- Execute or schedule one or more ETL 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.
- Execute or schedule the ETL job that aggregates Info Mart data when you want to recalculate aggregates for the time span that you specify.
- Execute or schedule `Job_MigrateGIM` to perform critical data migration when you are migrating from a previous release of Genesys Info Mart.
- Display the current ETL job execution status.
- Display a history of ETL job execution, such as start time, stop time, duration, and final status.
- Filter the jobs based on time and status.
- Execute a single ETL job as needed (either now or at some future specified time and date), including passing a required input parameter.
- Remove a scheduled job.
- Selectively shut down a running ETL 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 Genesys Info Mart Administration Console:

1. Open Configuration Manager.
2. Select the Genesys Info Mart ETL `Application` object.
3. Right-click and select `Wizard > Configure` options. This will invoke the Genesys Info Mart Administration Console for the Genesys Info Mart ETL `Application` object.

Monitoring ETL Jobs

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

To monitor your ETL jobs, view the `Status` tab window in the Genesys Info Mart Administration Console. Figure 1 shows the `Status` tab window.

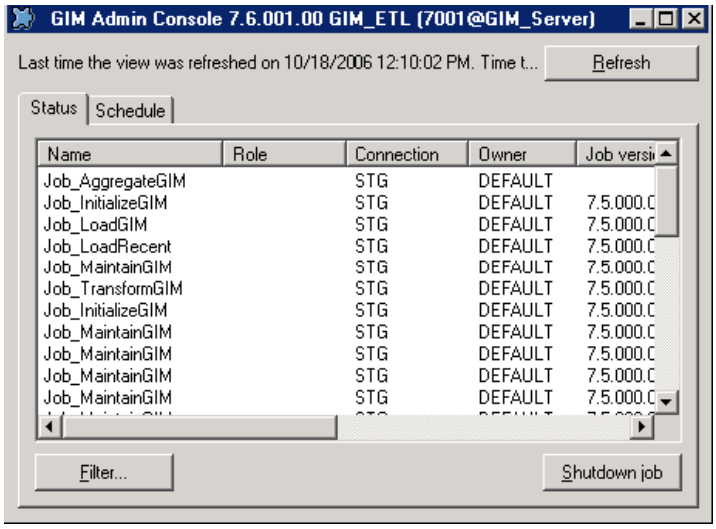

Figure 1: Genesys Info Mart Administration Console Status Tab

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`

- For data extraction jobs, the database connection (under the `Connection` column) and database owner (under the `Owner` column). The following database connection information is displayed for the supported database types:
  - Oracle—`<TNS Name>`
  - Microsoft SQL—`<DB Name>@<Server>:<Port>`
  - DB2—`<Server>:<Port>`
- For the ICON data extraction job (`Job_ExtractICON`), the extraction role.
- The time of the last update and the time until the next update. To determine when a job completed, see the `End Time` column.

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.

Note: 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 also filter the view to determine which jobs will appear. See "Filtering Genesys Info Mart Administration Console View" on page 81 for more information.

Executing and Scheduling ETL Jobs

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

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

- Run the ETL job that initializes the Staging Area and Info Mart databases (`Job_InitializeGIM`).
- Run one or more ETL jobs to recover from job failures.
- Run the ETL job that aggregates Info Mart data when you want to recalculate aggregates for a time span that you specify.
- Run `Job_ExtractICON` following the update of configuration data in IDB with the ICON on-demand resynchronization feature.
- Run `Job_MigrateGIM` to migrate data from the Staging Area and Info Mart databases of release 7.5 to release 7.6.

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

Running ETL Jobs One-by-One

Initially, you may want to run the ETL jobs one-by-one, rather than scheduling them to run. To run the jobs one-by-one:

1. Set the `run-scheduler` configuration option (in the `[schedule]` section on the Genesys Info Mart `Options` tab) to `FALSE`.
2. Run the jobs in the order in which they appear in the following list:
   - `Job_InitializeGIM`
   - `Job_ExtractICON` (see the following Note)
   - `Job_ExtractSS` (see the following Note)
   - `Job_ExtractGVP` (see the following Note)
   - `Job_TransformGIM`
   - `Job_LoadRecent`
   - `Job_LoadGIM`
   - `Job_AggregateGIM`
   - `Job_MaintainGIM`

Notes:

- In a migrating environment, `Job_MigrateGIM` must run first, to migrate the data for the tables that are critical for other jobs to run.
- Execute the extraction job for each source database from which you want to extract data, or specify `ALL SOURCES`. For `Job_ExtractICON`, you must select `ALL SOURCES`. Generally, each DAP represents a source database.
- When you run `Job_ExtractICON` for `All SOURCES`, you may see a popup box that indicates that some extraction jobs are not eligible to run until after `Job_TransformGIM` has completed. This popup box does not indicate an error, but is a reminder that a subset of extraction jobs has already ran after the last time `Job_TransformGIM` completed, and only the remaining extraction jobs will now run.

After the initial run, you can set the `run-scheduler` configuration option to `TRUE`, then use the Genesys Info Mart Server to launch ETL jobs based on the schedule you configured in the Genesys Info Mart application. For more information, see "Launching ETL Jobs with Genesys Info Mart Server" on page 66.

Running an ETL Job Immediately

Use the following steps to run an ETL 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 2.

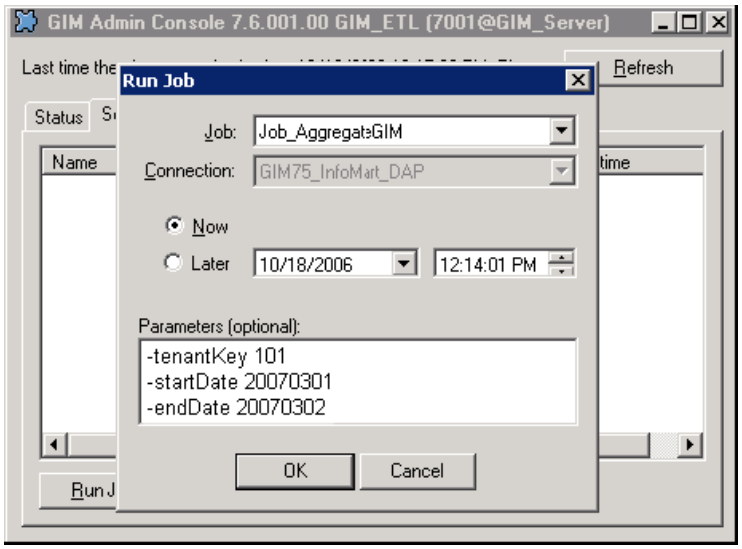

Figure 2: Genesys Info Mart Administration Console Run Job Window

3. From the `Job` drop-down list, select the job that you wish to execute.
4. For a data extraction job, select a DAP or the `ALL SOURCES` option from the `Connection` drop-down list. The list contains all DAPs attached to the Genesys Info Mart ETL application that are valid for the selected job. Selecting a DAP is optional. If you specify a DAP, the job will extract data only from that DAP. If you select `ALL SOURCES`, the job will extract data from all configured DAPs that are connected to the Genesys Info Mart application.

Notes:

- When you run the ICON data extraction job (`Job_ExtractICON`), you must specify `ALL SOURCES`. Genesys Info Mart Server will automatically launch a separate instance of `Job_ExtractICON` for each extraction role you configured in the DAP.

- When you run Job_ExtractICON for All SOURCES, you may see a popup box that indicates that some extraction jobs are not eligible to run until after Job_TransformGIM has completed. This popup box does not indicate an error, but is a reminder that a subset of extraction jobs has already ran after the last time Job_TransformGIM completed, and only the remaining extraction jobs will now run.

5. Select the option to run the job Now.
6. If any parameters are needed to execute the job, enter the required parameters in the Parameters text box.

   Precede the parameter name with a dash (–), and enter a space before the parameter value. You can enter multiple parameters on a single line or use multiple lines. The parameter names are not case sensitive.

   Example 1

   ```
   -tenantKey 1 -startDate 20061101 -endDate 20061102
   ```

   Example 2

   ```
   -tenantKey 1
   -startDate 20061101
   -endDate 20061102
   ```

Note: Job_AggregateGIM and Job_TransformGIM are the only jobs that accept input parameters. For information about the optional parameters, see "Calculating Aggregates Within a Time Span" on page 53 and "Job_TransformGIM" on page 42.

7. Click the OK button.
8. 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 startup, the status will be Failed.

Scheduling an ETL 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 an ETL 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 2 on page 78.
3. From the Job drop-down list, select the job that you wish to execute.
4. For a data extraction job, select a connection from the Connection drop-down list. The list contains all DAPs that are attached to the Genesys Info Mart ETL application and are valid for the selected job. Selecting a DAP is optional. If you specify one, the job will extract data only from that DAP. If you do not specify one, the job will extract data from all configured DAPs that are connected to the Genesys Info Mart application.

   Note: When you run the ICON data extraction job (Job_ExtractICON), you must specify ALL SOURCES. Genesys Info Mart Server will automatically launch a separate instance of Job_ExtractICON for each extraction role that you configured on the DAP.

5. 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.
6. If any parameters are needed to execute the job, enter the required parameters in the Parameters text box.

   Precede the parameter name with a dash (–), and enter a space before the parameter value. You can enter multiple parameters on a single line or use multiple lines. The parameter names are not case sensitive. For examples—see page 79.

   Note: Job_AggregateGIM and Job_TransformGIM are the only jobs that accept input parameters. For information about the optional parameters, see "Calculating Aggregates Within a Time Span" on page 53 and "Job_TransformGIM" on page 42.

7. 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 an ETL Job

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

1. In the Genesys Info Mart Administration Console, select the `Status` tab. The window displays the status of current and completed jobs (see Figure 1 on page 75).
2. Select the job that you want to stop. You can stop jobs that have a status of `Running` or `Waiting`.
3. Click the `Shutdown Job` button at the bottom of the window.

   A confirmation window appears (see Figure 3).

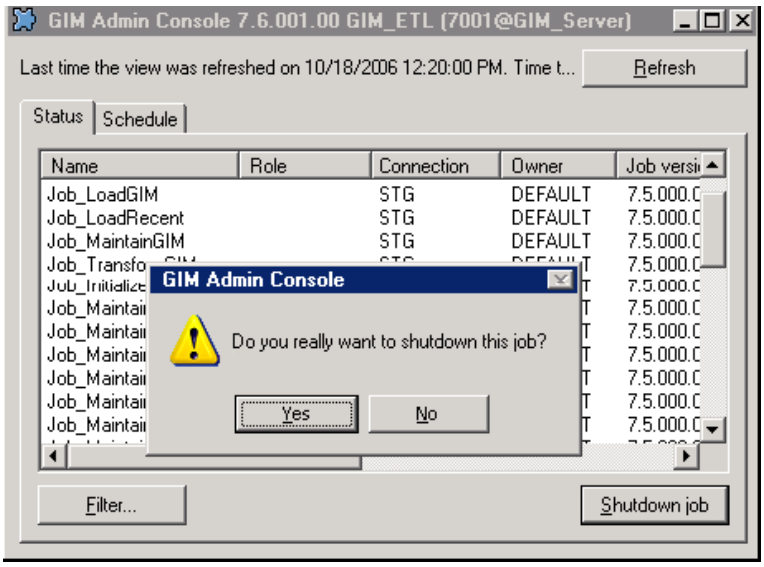

Figure 3: Shutdown Job Confirmation

4. 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. The job will transition to `Shutdown` state once it has completely shut down. The job status that is reported on the Genesys Info Mart Administration Console `Status` tab will change accordingly.

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 1 on page 75).
2. Click the `Filter` button. The `Filter` dialog box appears (see Figure 4).

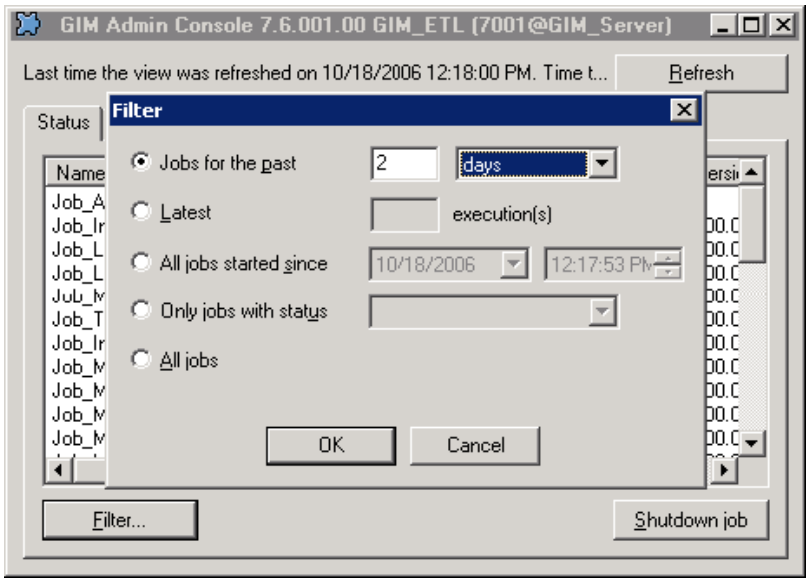

Figure 4: Filtering Options for Genesys Info Mart Administration Console

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.

About the Intraday ETL Cycle

This section describes the scope of the data extraction, transformation, and loading and provides sample ETL job schedules.

Extraction

Each time Genesys Info Mart Server runs an extraction job (`Job_ExtractICON`, `Job_ExtractGVP`, or `Job_ExtractSS`), data that has been added or updated since the last time you ran that job is extracted and stored in the Staging Area database. The extraction algorithm that Genesys Info Mart Server uses depends on the type of data.

If configured, the extraction job launches the following operations during each ETL cycle:

- Controlling the rate of data extraction to prevent too much data from being processed on each ETL cycle, and to delay extraction of voice interactions to facilitate synchronization of associated UserEvent-based KVP data.
- IDB merge (`Job_ExtractICON` only)
- High availability (HA) Deduplication for Configuration or Voice details or HA switchover for Outbound Contact details (`Job_ExtractICON` only)

For more information, see "Job_ExtractICON" on page 33, "Job_ExtractGVP" on page 41, or "Job_ExtractSS" on page 135.

Note: To improve system performance, Genesys recommends that you gather statistics from ICON before running the initial ICON extraction cycle, and on a regular basis thereafter. See "Improving System Performance During ICON Extraction" on page 38 for more information.

Transformation

The transformation job (`Job_TransformGIM`) transforms data that is ready to be processed (in other words, that is not to be delayed until a subsequent run), and stores the transformed data in the Staging Area database until the next time that the Genesys Info Mart Server runs the intraday loading job (`Job_LoadRecent`). `Job_TransformGIM` saves the data that is not ready to be processed (in other words, that is to be delayed until a subsequent run) in the Staging Area database and attempts to transform it the next time that Genesys Info Mart Server runs the transformation job.

The job determines whether or not data is ready to be processed, based on the following considerations:

- The status of contact center activities (active or complete)—The job delays transforming some types of data until the activities they represent are complete. Tables 4 and 5 indicate the status required for different kinds of activities, in order for the associated facts to be transformed.
- The data source lag time

- The job delays transforming some of the most recent Outbound Contact data to ensure that IDB has had enough time to store all related source data that the job must process as a unit. The job calculates the end of the data transformation window by subtracting fifteen minutes from the most recent data extraction job start time.
- The job delays transforming some of the most recent agent activity–related data to ensure that IDB has had enough time to store all related source data that the job must process as a unit. The job calculates the end of the data transformation window by subtracting five minutes from the start time of the most recent agent activity–related data recorded in IDB by a given Interaction Concentrator instance. A new configuration option in the `[gim-transformation]` section, `complex-voice-agent-env`, can be used to bypass this five minute delay for voice agent activity in certain deployments. (See the *Genesys Info Mart 7.6 Deployment Guide* for details.)

Table 4: Transformation and Loading by Activity Status

Info Mart Fact Tables		Activity Status
Category	**Table**	
Contact Center Configuration History	PLACE_GROUP_FACT	Active
	RESOURCE_GROUP_FACT	Active
	RESOURCE_SKILL_FACT	Active
	CAMPAIGN_GROUP_TO_CAMPAIGN_FACT	Active
	CALLING_LIST_TO_CAMPAIGN_FACT	Active
Outbound Contact	CAMPAIGN_GROUP_SESSION_FACT	Active
	CAMPAIGN_GROUP_STATE_FACT	Active
	CALLING_LIST_METRIC_FACT	Complete
	CONTACT_ATTEMPT_FACT	Complete
GVP VAR	GVP_CALL_FACT	Complete
	GVP_SUBCALL_FACT	Complete

Table 5: Transformation and Loading by Activity Status and Media Type

Info Mart Fact Tables		Activity Status	
Category	**Table**	**Voice**	**Multimedia**
Interaction Activity	INTERACTION_FACT	Complete	Active
	[Media]_IXN_FACT_EXT	Complete	Active
	INTERACTION_RESOURCE_FACT	Complete	Not Applicable
	[Media]_RES_FACT_EXT	Complete	Not Applicable
	IXN_RESOURCE_STATE_FACT	Complete	Not Applicable
	INTERACTION_SEGMENT_FACT	Complete	Active
	[Media]_SEG_FACT_EXT	Complete	Active
Resource Activity	DT_DND_FACT	Active	Active
	DT_RES_STATE_FACT	Complete	Complete
	DT_RES_STATE_REASON_FACT	Complete	Complete
	RESOURCE_SESSION_FACT	Active	Active
	RESOURCE_STATE_FACT	Complete	Complete
	RESOURCE_STATE_REASON_FACT	Complete	Complete
	SM_RES_SESSION FACT	Active	Active
	SM_RES_STATE_FACT	Complete	Complete
	SM_RES_STATE_REASON_FACT	Complete	Complete
ACD Queue and Virtual Queue Activity	MEDIATION_SEGMENT_FACT	Complete	Active

Intraday Loading Each time that Genesys Info Mart runs the intraday load job (`Job_LoadRecent`), all the transformed data in the Staging Area database is loaded into the dimension and intraday fact tables in the Info Mart database. In addition, if you enabled intraday aggregation, the intraday aggregate tables are updated in the Info Mart database in accordance with the configured frequency for running the intraday aggregation part of `Job_LoadRecent` (the `intraday-aggregates-frequency` option).

Historical Loading Each time that Genesys Info Mart runs the historical load job (`Job_LoadGIM`), the data that was previously loaded by `Job_LoadRecent` into each intraday fact table is moved to its counterpart historical fact table. Multimedia interactions

generally take much longer to complete than interactions for voice media type, and therefore will require updates across several ETL cycles. For this reason, `Job_LoadGIM` does not move them from the intraday fact tables to the historical fact tables until they are complete. Then, `Job_LoadGIM` deletes the data from the intraday fact and intraday aggregate tables.

Schedule Considerations

The number of times that you schedule the jobs to run each day depends on whether you plan to use intraday loading, and such contact center factors as the daily call volume, peaks in the call rate, network latency between components, and the capacity of the Genesys Info Mart Server or Staging Area database server:

- If you have a low daily call volume and do not plan to use intraday loading, you may choose to schedule all the jobs to run once per day, after the end of the reporting day. Genesys does not recommend this type of scheduling, because processing a full day's data in a single ETL cycle requires a lot of database resources and takes a long time to complete.
- If there are slow network connections between the data sources and the Genesys Info Mart Server, you may choose to schedule the extraction and transformation jobs multiple times per day.
  - If you plan to use intraday loading, schedule the intraday loading job along with the extraction and transformation jobs.
  - If you do not plan to use intraday loading, schedule the intraday load job to run once per day before the historical load job.
- If you have a high daily call volume or if you experience high peak call rates that result in a significant backlog of extracted data, you may choose to schedule the extraction and transformation jobs multiple times per day.
  - If you plan to use intraday loading, schedule the intraday loading job along with the extraction and transformation jobs.
  - If you do not plan to use intraday loading, schedule the intraday load job to run once per day before the historical load job.
- Genesys recommends that you run the historical load job (`Job_LoadGIM`) only once per day, after the final run of the extraction, transformation, and intraday load jobs. Running the historical load job multiple times a day can severely degrade the business users' ability to query the Info Mart database data.

Job Interdependencies

Figure 5 depicts the interdependencies among jobs.

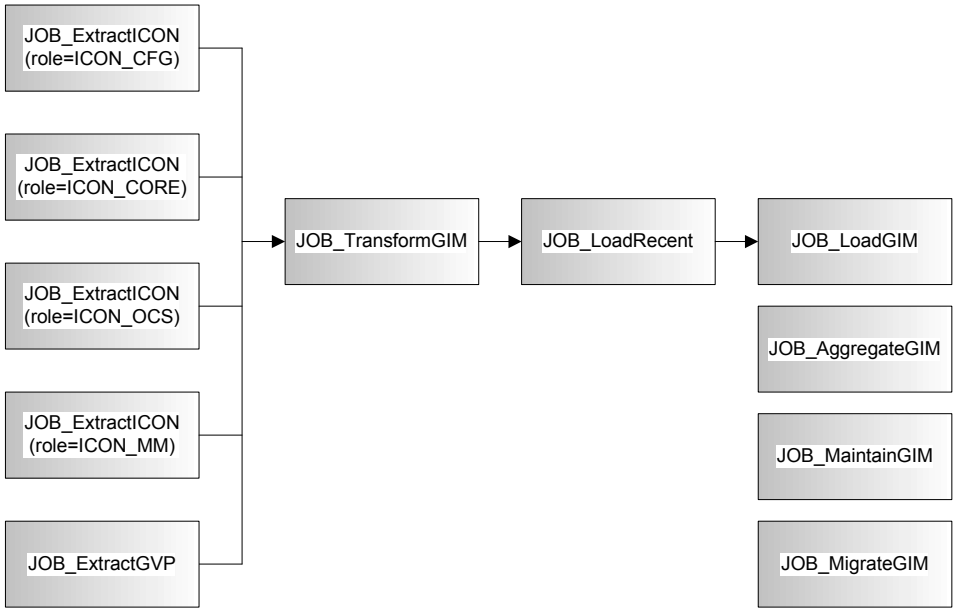

Figure 5: Interdependencies Among Jobs

Genesys Info Mart Server manages the job interdependencies automatically as follows:

- Genesys Info Mart Server will not launch the transformation job until all the extraction jobs complete successfully. If you enabled intraday loading, Genesys Info Mart Server will not launch the transformation job until the load steps of the previous intraday load job also complete successfully.

  Note: One of the extraction jobs must be the extraction of configuration data from IDB. Otherwise, the transformation job will not start.

- Genesys Info Mart Server will not launch the intraday load job until the transformation job completes successfully.
- Genesys Info Mart Server will not launch the historical load job until the intraday load job completes successfully.
- Genesys Info Mart Server ensures that the historical load job, the aggregation job, and the purging job do not run at the same time.

Note: Genesys Info Mart does not prevent you from running `Job_MigrateGIM` simultaneously with other daily jobs. Consider scheduling these jobs at different times, to prevent them from accessing the same database tables.

The ETL jobs manage their interdependencies automatically, even when you use the Genesys Info Mart Administration Console to run an ETL job immediately or schedule it to run at some future time.

Sample Schedules

This section provides the following sample schedules:

- "Intraday Loading in the Same Time Zone" on page 88
- "Intraday Loading in Different Time Zones" on page 90
- "Loading Once at End of Day" on page 91
- "Migrating Data Within a Daily Schedule" on page 93

Note: The term "End of Reporting Day According to Business Users" appears on the bottom of Figure 6, Figure 7 on page 90, Figure 8 on page 92, and Figure 9 on page 93. This term is based on the configured time-zone setting (`std-enterprise-time-zone`). When midnight occurs in the configured time zone, this is considered to be the end of the reporting day.

Intraday Loading in the Same Time Zone

Figure 6 depicts a sample Genesys Info Mart Server ETL job schedule that uses intraday loading.

The schedule options are set as follows:

- `etl-start-time` = `06:30`
- `etl-end-time` = `00:30`
- `etl-frequency` = `30` minutes
- `run-load-recent-with-extract-and-transform` = `TRUE`
- `populate-intraday-aggregates` = `TRUE`
- `intraday-aggregates-frequency=0`
- `load-start-time` = `02:00`
- `aggregate-start-time` = `03:00`
- `maintain-start-time` = `04:00`

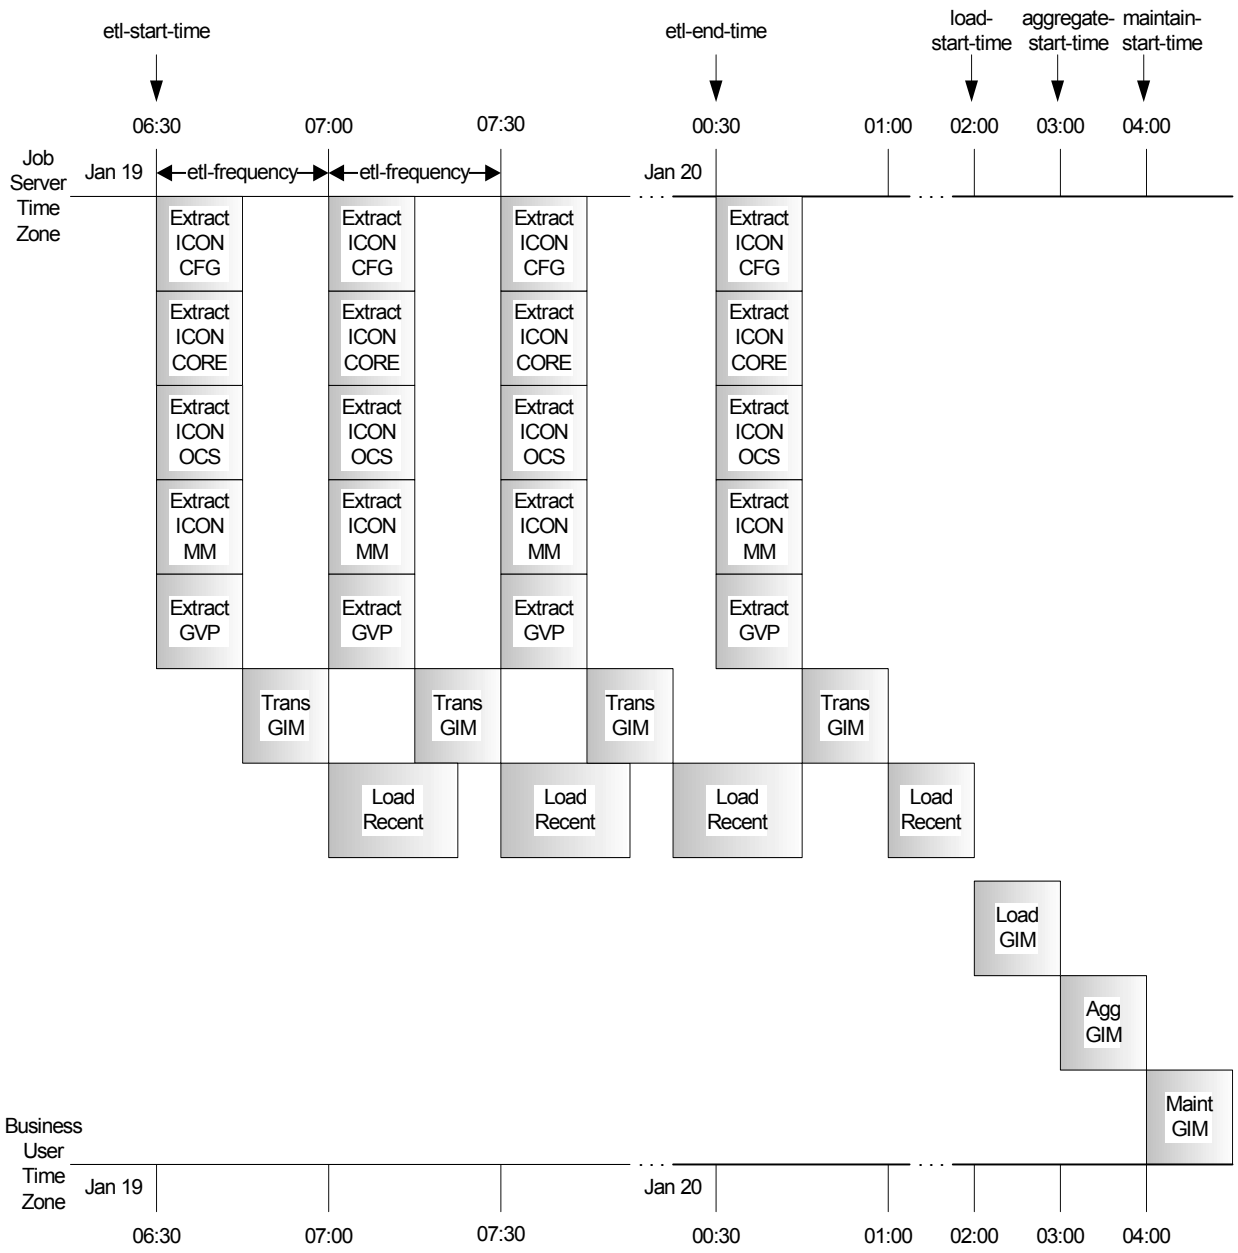

Figure 6: Sample Genesys Info Mart Server Schedule—Intraday Loading in the Same Time Zone

In the sample schedule, the Genesys Info Mart Server and the business users that will query the Info Mart data are in the same time zone, so the final extraction that begins at 00:30 contains all the reporting data from the previous day from the perspective of the business users.

Intraday Loading in Different Time Zones

Figure 7 depicts another sample Genesys Info Mart Server ETL job schedule that uses intraday loading.

The schedule options are set as follows:

- `etl-start-time` = `03:30`
- `etl-end-time` = `21:30`
- `etl-frequency` = `30` minutes
- `run-load-recent-with-extract-and-transform` = `TRUE`
- `populate-intraday-aggregates` = `TRUE`
- `intraday-aggregates-frequency=0`
- `load-start-time` = `23:00`
- `aggregate-start-time` = `00:00`
- `maintain-start-time` = `01:00`

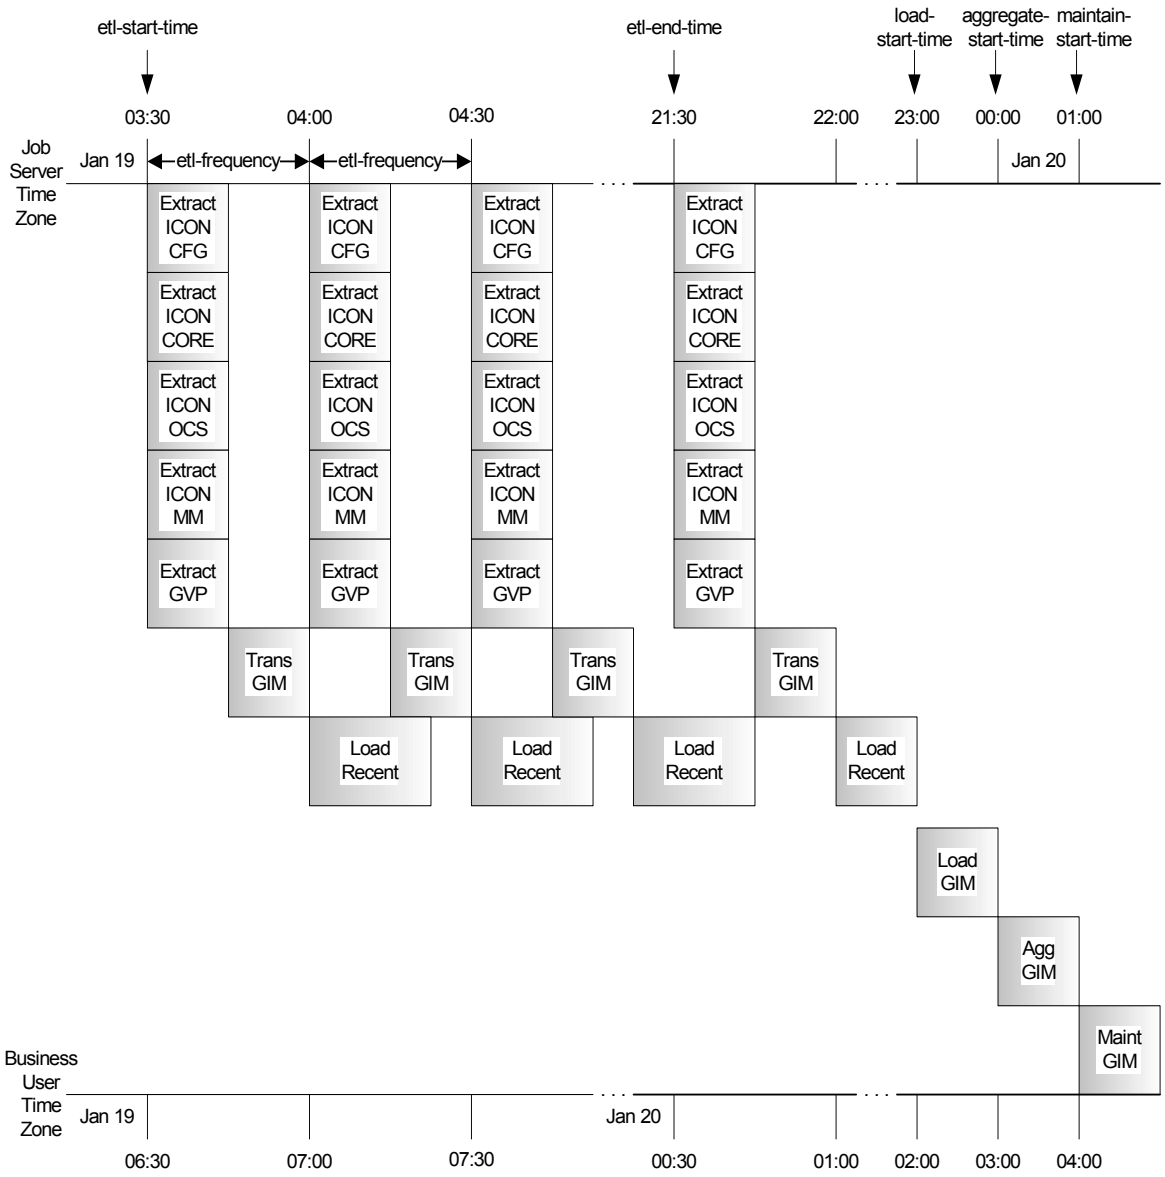

Figure 7: Sample Genesys Info Mart Server Schedule—Intraday Loading in Different Time Zones

In the sample schedule, the Genesys Info Mart Server and the business users that will query the Info Mart data are in different time zones, so the final extraction that begins at 21:30 contains all the reporting data from the previous day from the perspective of the business users.

Loading Once at End of Day

Figure 8 depicts a sample Genesys Info Mart Server ETL job schedule that does not use intraday loading. The schedule options are set as follows:

- `etl-start-time` = `06:30`
- `etl-end-time` = `00:30`
- `etl-frequency` = `30` minutes
- `run-load-recent-with-extract-and-transform` = `FALSE`
- `populate-intraday-aggregates` = `FALSE`
- `load-recent-time` = `01:00`
- `load-start-time` = `02:00`
- `aggregate-start-time` = `03:00`
- `maintain-start-time` = `04:00`

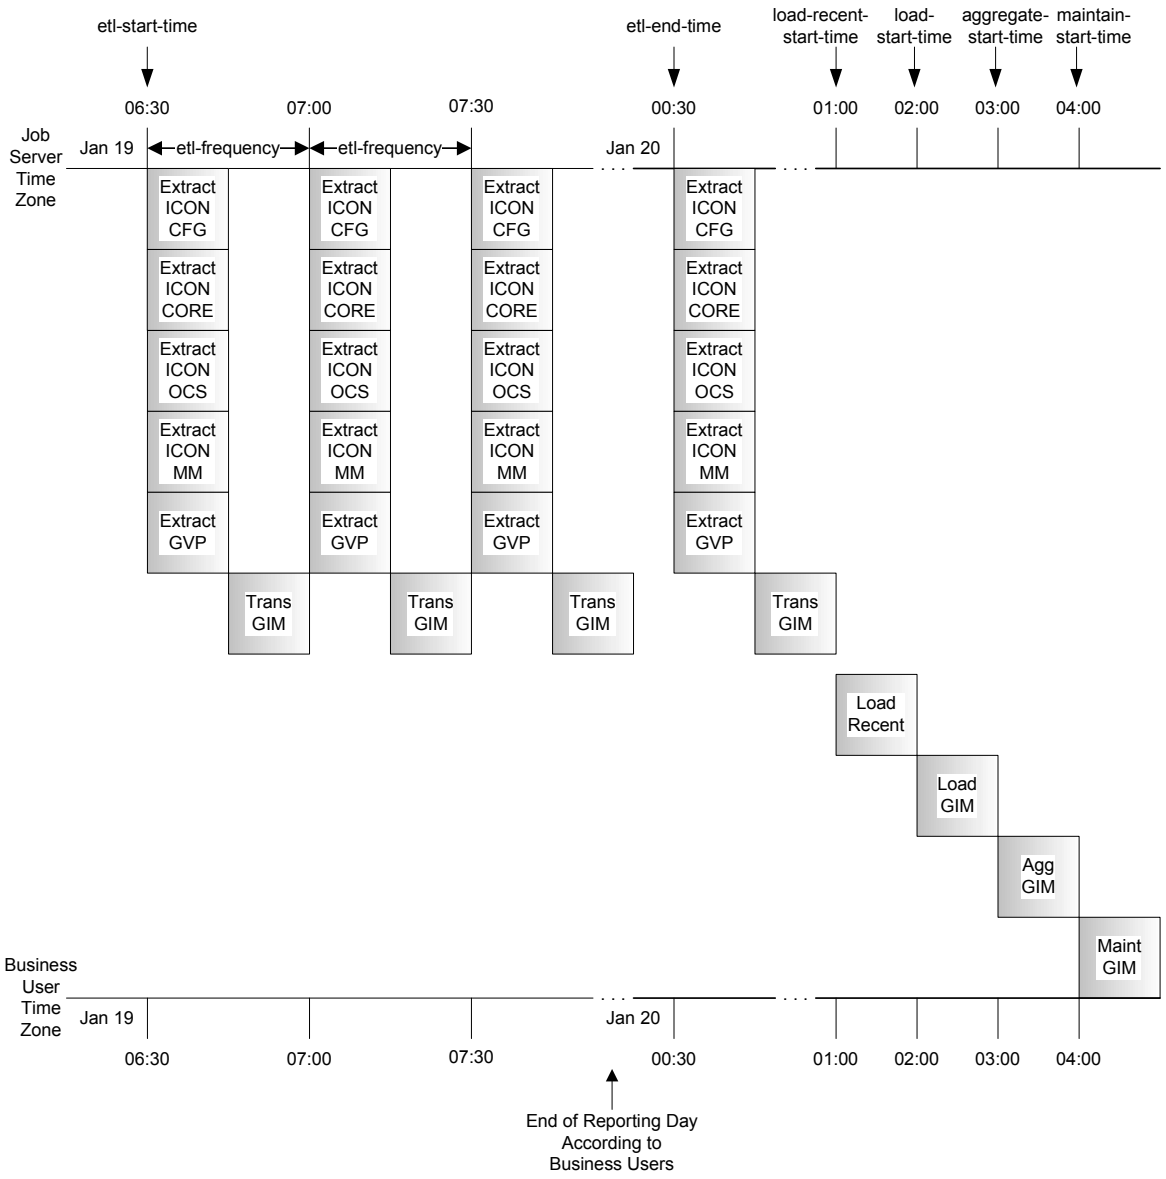

Figure 8: Sample Genesys Info Mart Server Schedule—Loading Once at End of Day

Migrating Data Within a Daily Schedule

Figure 9 depicts a sample Genesys Info Mart Server ETL job schedule that accounts for the migration job, `Job_MigrateGIM`, to be run on a daily basis, outside of the intraday ETL cycle. The configuration options are set as follows:

`[schedule]`

- `run-migration = TRUE`
- `migration-start-time = 05:00`
- `migration-duration-in-hours = 1`
- `etl-start-time = 06:30`
- `etl-end-time = 00:30`

`[etl]`

- `data-migration-time-unit = DAYS`
- `data-migration-time-range = 2`

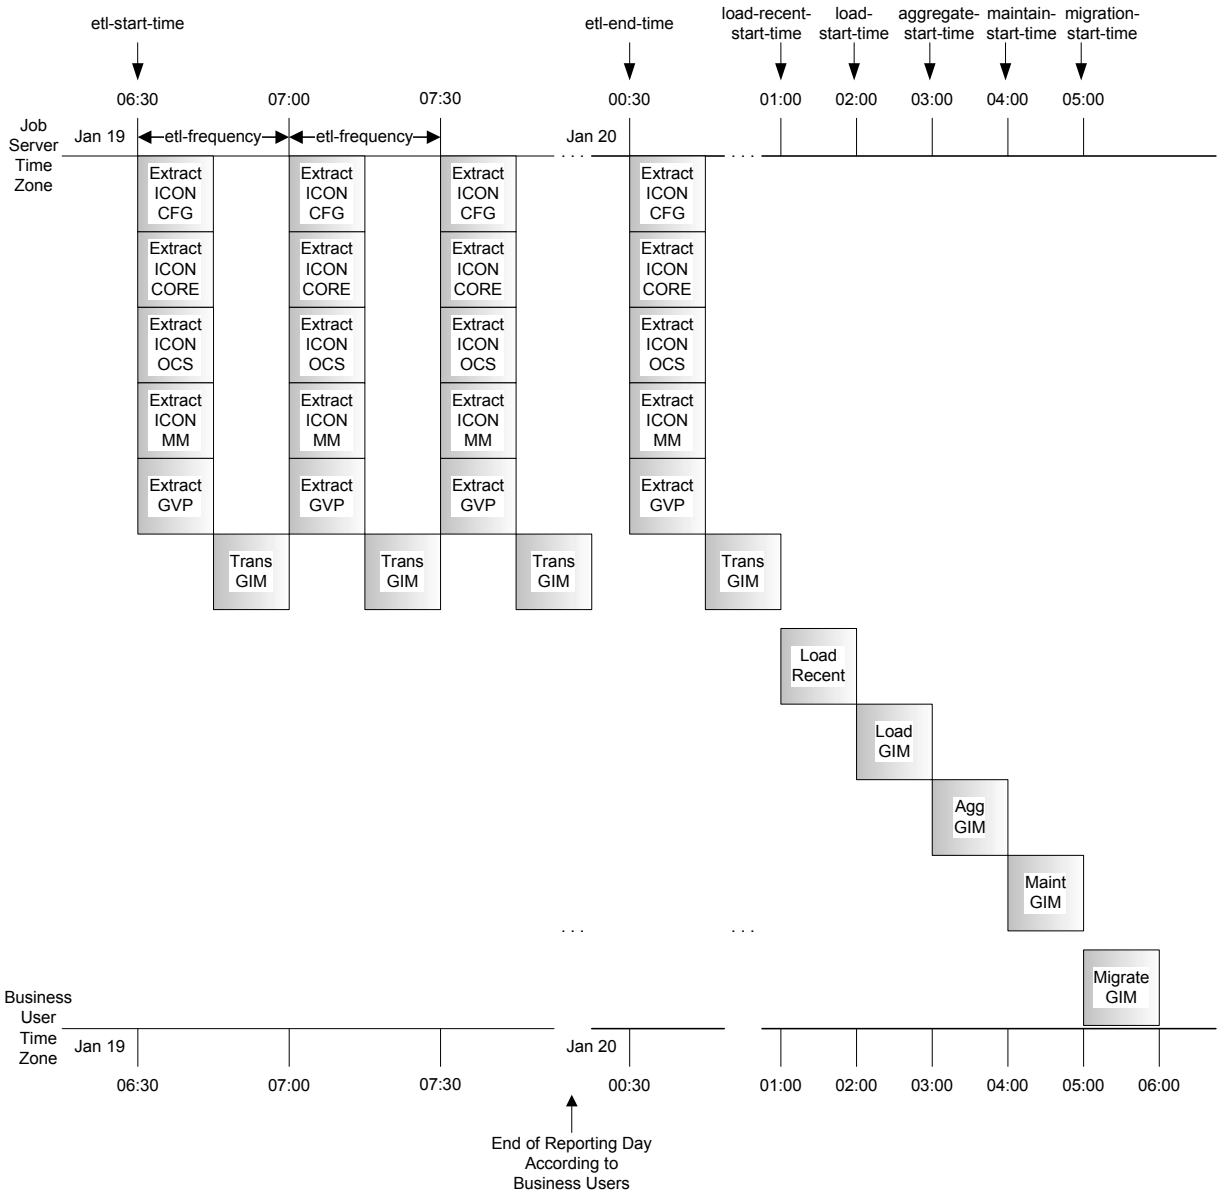

Figure 9: Sample Genesys Info Mart Server Schedule—Migrating Data Within a Daily Schedule

In this sample schedule, Genesys Info Mart Server is to schedule the migration job to run for one hour, starting at 5:00 AM, outside of the intraday ETL cycle. The job is to migrate the data in tables that have standard enterprise date-time keys, in chunks of two days' worth of data.

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 contains the Genesys recommendations for purging old data from the Genesys Info Mart database.

For information about purging old data from the Interaction Database (IDB) and Stat Server database, see "Purging Data Sources" on page 112, and "Purging the Stat Server Database" on page 136.

Purging Data from Fact Tables

`Job_MaintainGIM` determines the maximum standard tenant end date for the facts that you want to purge as follows:

- Converts the current system date and time in the Genesys Info Mart Server's time zone to the equivalent date and time in the tenant's standard reporting time zone.
- Subtracts the tenant's configured `days-to-keep-gim-facts` (or `days-to-keep-dt-resource-activity-facts`, for data in detail `DT_*` fact tables).
- Calculates only the date portion, dropping the time portion.

Figure 10 on page 95 shows a purging eligibility example for fact tables. When looking at the example, keep in mind the following:

- The `days-to-keep-gim-facts` for the tenant is configured as 380 days.
- The tenant's standard reporting time zone is three hours behind the Genesys Info Mart Server's time zone.
- `Job_MaintainGIM` is running at 02:00 on January 20, 2004.
- Year 2004 is a leap year.

`Job_MaintainGIM` determines the maximum standard tenant end date by each of the following actions:

- Converting 02:00 on January 20, 2004 (Genesys Info Mart Server's time zone) to 23:00 on January 19, 2004 (tenant's standard reporting time zone).
- Subtracting 380 days from 23:00 on January 19, 2004, which is 23:00 on January 5, 2003.
- Calculating only the date portion, which is January 5, 2003.

In Figure 10, `Job_MaintainGIM` purges interactions 1 and 2 and their segments, since all of the segments end on or before January 5, 2003. The job does not purge Interactions 3 and 4 and their segments, since some of their segments do

not end on or before January 5, 2003. Similarly, the job purges Resource States 1 and 2 and Resource Sessions 1 and 2, but does not purge Resource States 3 and 4 or Resource Sessions 3 and 4.

Note: Figure 10 does not depict all fact tables from which `Job_MaintainGIM` purges old data.

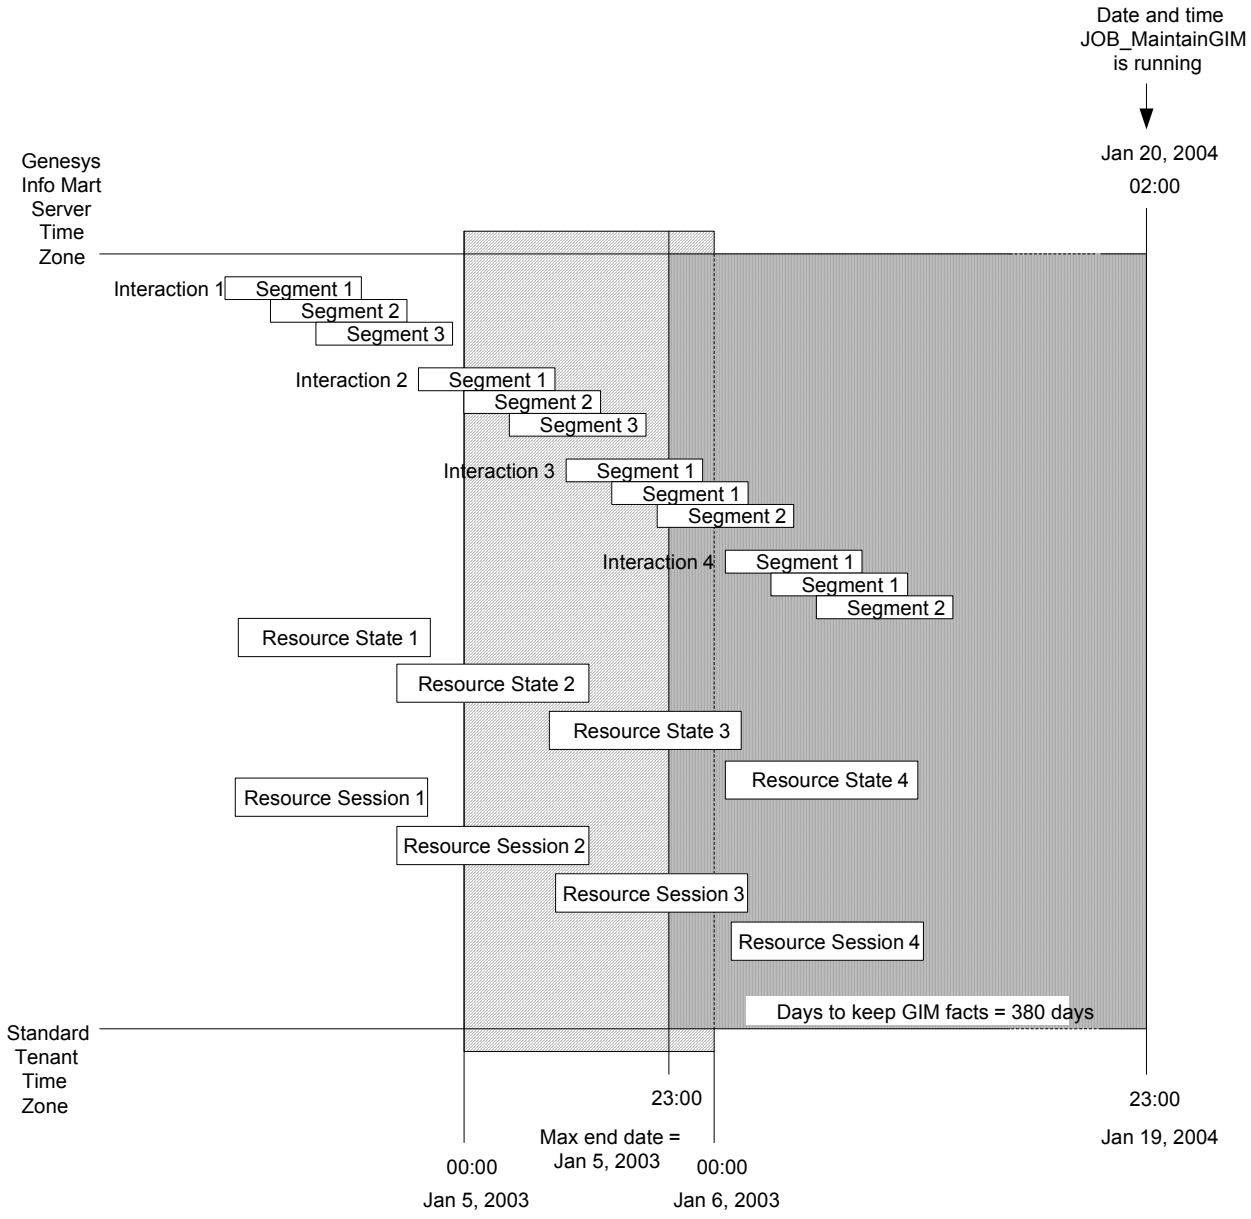

Figure 10: Purging Eligibility

Purging Data from Aggregate Tables

For aggregate tables, `Job_MaintainGIM` uses a procedure similar to the fact tables, to determine what data to purge. The following are the differences:

- As a date reference for determining purging eligibility, `Job_MaintainGIM` takes:
  - For Genesys Interactive Insights (GI2) out-of-box aggregates, the standard tenant date and the time that represent the start of the aggregation interval.
  - For CCPulse+ out-of-box aggregates, the date portion of the standard tenant date that represents the start of the interval.
- The data retention period is set by the following options in the [`gim-aggregates-tenant`] section:
  - `days-to-keep-hour-level-disposition-aggregates`
  - `days-to-keep-day-level-disposition-aggregates`
  - `days-to-keep-month-level-disposition-aggregates`
  - `days-to-keep-subhour-level-interval-aggregates`
  - `days-to-keep-hour-level-interval-aggregates`
  - `days-to-keep-day-level-interval-aggregates`

Staging Area Administrative Views

The Staging Area database provides several views to help database administrators maintain Genesys Info Mart and track ETL performance over time, and to facilitate technical support by Genesys. This section describes the following views:

- "ADMIN_ETL_STEP_HISTORY" on page 97
- "ADMIN_ETL_JOB_STATUS" on page 97
- "ADMIN_ETL_JOB_HISTORY" on page 98
- "ADMIN_EXTRACT_HISTORY" on page 99
- "ADMIN_LOAD_HISTORY" on page 100
- "ADMIN_AGGREGATE_HISTORY" on page 101
- "ADMIN_PURGE_HISTORY" on page 102
- "CV_MISSING_CFG_OBJECTS" on page 103

Note: The tables in the following subsections show the generic data types. The actual data types vary, depending on the RDBMS type.

ADMIN_ETL_STEP_HISTORY

This view provides information about the execution of each ETL job step. Rows are added to this view for completed ETL job steps only. As each ETL job completes, it adds rows for the completed steps of all currently running ETL jobs, including itself, that have not already been added to the view. Currently running ETL jobs may have steps that are in process or are waiting, and they do not yet appear in the view. Rows in this view are written once and are not updated.

Table 6 lists the `ADMIN_ETL_STEP_HISTORY` view columns and descriptions.

Table 6: ADMIN_ETL_STEP_HISTORY View Columns

Column Name	Data Type	Description
JOB_ID	VARCHAR	ID that uniquely identifies the execution instance of the job.
JOB_NAME	VARCHAR	The name of the job, such as `Job_ ExtractICON`.
JOB_VERSION	VARCHAR	The version of the job, such as 7.6.000.10.
START_TIME	DATETIME	The date and time the step started (Genesys Info Mart Server time zone).
END_TIME	DATETIME	The date and time the step ended (Genesys Info Mart Server time zone).
DURATION	INTEGER	The duration of the step in seconds.
DBCONNECTION	VARCHAR	The connection name of the data source or target.
OWNER	VARCHAR	The owner of the tables in the data source or target.
SOURCE_TYPE	VARCHAR	The primary data source or target type involved in the step, such as `GVP_VAR`, `ICON_CFG`, `ICON_CORE`, `ICON_OCS`, `ICON_MM`, `STG_AREA`, `STAT_SERVER`.
WORKFLOW_TYPE	VARCHAR	The type of the step, such as `CHECK`, `CLEANSE`, `EXT`, `INIT`, `AGGR`, `LOAD`, `PURGE`, `TRUNC`, `WINDOW`, `XFORM`, `CRITICAL`, `NONCRITICAL`.
STATUS	VARCHAR	The status of the step, such as `COMPLETE` or `FAILED`.

ADMIN_ETL_JOB_STATUS

This view provides information about the most recent execution of each ETL job. A row is added to this view after each ETL job starts, and is updated as the job status changes.

Table 7 lists the `ADMIN_ETL_JOB_STATUS` view columns and descriptions.

Table 7: ADMIN_ETL JOB_STATUS View Columns

Column Name	Data Type	Description
JOB_ID	VARCHAR	ID that uniquely identifies the execution instance of the job.
JOB_NAME	VARCHAR	The name of the job, such as `Job_ExtractICON`.
JOB_VERSION	VARCHAR	The version of the job, such as 7.6.000.10.
START_TIME	DATETIME	The date and time the step started (Genesys Info Mart Server time zone).
END_TIME	DATETIME	The date and time the step ended (Genesys Info Mart Server time zone).
DURATION	INTEGER	The duration of the step in seconds.
DBCONNECTION	VARCHAR	The connection name of the data source or target.
OWNER	VARCHAR	The owner of the tables in the data source or target.
ROLE	VARCHAR	The primary data source or target type involved in the step, such as `GVP_VAR`, `ICON_CFG`, `ICON_CORE`, `ICON_OCS`, `ICON_MM`, `STAT_SERVER`.
STATUS	VARCHAR	The status of the step, such as `INSTALLED`, `RUNNING`, `WAITING`, `SHUTDOWN`, `COMPLETE`, or `FAILED`.

ADMIN_ETL_JOB_HISTORY

This view provides information about the execution of each ETL job. A row is added to this view after each ETL job completes. Currently running ETL jobs do not appear in this view. Rows in this view are written once and are not updated.

Table 8 lists the `ADMIN_ETL_JOB_STATUS` view columns and descriptions.

Table 8: ADMIN_ETL JOB_HISTORY View Columns

Column Name	Data Type	Description
JOB_ID	VARCHAR	ID that uniquely identifies the execution instance of the job.
JOB_NAME	VARCHAR	The name of the job, such as `Job_ExtractICON`.
JOB_VERSION	VARCHAR	The version of the job, such as 7.6.000.10.
START_TIME	DATETIME	The date and time the step started (Genesys Info Mart Server time zone).

Table 8: ADMIN_ETL JOB_HISTORY View Columns (Continued)

Column Name	Data Type	Description
END_TIME	DATETIME	The date and time the step ended (Genesys Info Mart Server time zone).
DURATION	INTEGER	The duration of the step in seconds.
DBCONNECTION	VARCHAR	The connection name of the data source or target.
OWNER	VARCHAR	The owner of the tables in the data source or target.
ROLE	VARCHAR	The primary data source or target type involved in the step, such as `GVP_VAR`, `ICON_CFG`, `ICON_CORE`, `ICON_OCS`, `ICON_MM`, `STAT_SERVER`.
STATUS	VARCHAR	The status of the step, such as `COMPLETE` or `FAILED`.

ADMIN_EXTRACT_HISTORY

This view provides information about the data extracted from each source database table. A row is added to this view after `Job_ExtractICON`, `Job_ExtractSS`, and `Job_ExtractGVP` successfully complete extracting a source data table. Rows in this view are written once and are not updated.

Table 9 lists the `ADMIN_EXTRACT_HISTORY` view columns and descriptions.

Table 9: ADMIN_EXTRACT_HISTORY View Columns

Column Name	Data Type	Description
JOB_ID	VARCHAR	ID that uniquely identifies the execution instance of the job.
JOB_NAME	VARCHAR	The name of the job, such as `Job_ExtractICON`.
JOB_VERSION	VARCHAR	The version of the job, such as 7.6.000.10.
START_TIME	DATETIME	The date and time the step started (Genesys Info Mart Server time zone).
DURATION	INTEGER	The duration of the step in seconds.
DBCONNECTION	VARCHAR	The connection name of the data source.
OWNER	VARCHAR	The owner of the tables in the data source.
GSYS_DOMAIN	INTEGER	ID that uniquely identifies the Interaction Concentrator domain. Applies only to tables extracted by `Job_ExtractICON`.
GSYS_SYS_ID	INTEGER	ID that uniquely identifies the ICON application instance. Applies only to tables extracted by `Job_ExtractICON`.

Table 9: ADMIN_EXTRACT_HISTORY View Columns (Continued)

Column Name	Data Type	Description
TABLE_NAME	VARCHAR	The name of the table from which data is extracted.
LAST_INSERT_SEQUENCE	NUMBER (20)	The highest sequence number extracted from the table. For IDB tables this will be a value from the GSYS_SEQ column. For ICON Outbound Contact details in an HA data extraction deployment, this will be a UTC time value indicating the number of seconds since January 1, 1970 at midnight.
LAST_UPDATE_SEQUENCE	NUMBER (20)	The highest sequence number extracted from the table. For IDB tables this will be a value from the GSYS_USEQ column.
LAST_MERGE_SEQUENCE	NUMBER (20)	The highest merge sequence number extracted from the table. This applies only to ICON Voice Interaction details and will be a value from the GSYS_MSEQ column of the G_IR table.
LATEST_DATA_TIME	DATETIME	The most recent date and time recorded by the source system in the data extracted for this table (Genesys Info Mart Server time zone).
ROW_COUNT	INTEGER	The number of rows extracted.

ADMIN_LOAD_HISTORY

This view provides information about data loaded into each Info Mart database table. A row is added to this view after Job_LoadRecent or Job_LoadGIM successfully completes loading each fact table. Rows in this view are written once and are not updated.

Table 10 lists the ADMIN_LOAD_HISTORY view columns and descriptions.

Table 10: ADMIN_LOAD_HISTORY View Columns

Column Name	Data Type	Description
JOB_ID	VARCHAR	ID that uniquely identifies the execution instance of the job.
JOB_NAME	VARCHAR	The name of the job, such as Job_LoadGIM.
JOB_VERSION	VARCHAR	The version of the job, such as 7.6.000.10.
START_TIME	DATETIME	The date and time the step started (Genesys Info Mart Server time zone).
DURATION	INTEGER	The duration of the step in seconds.
DBCONNECTION	VARCHAR	The connection name of the data source or target.

Table 10: ADMIN_LOAD_HISTORY View Columns (Continued)

Column Name	Data Type	Description
OWNER	VARCHAR	The owner of the tables in the data source or target.
TABLE_NAME	VARCHAR	The name of the table into which data is loaded.
HIGH_WATER_MARK	INTEGER	The primary key of the last row loaded into the table.
LATEST_DATA_TIME	DATETIME	The most recent date and time recorded by the source system in the data loaded for this table (Genesys Info Mart Server time zone).
ROW_COUNT_ADD	INTEGER	The number of rows inserted.
ROW_COUNT_ UPDATE	INTEGER	The number of rows updated.

ADMIN_AGGREGATE_HISTORY

This view provides information about data loaded into each Info Mart database aggregate table. A row is added to this view after `Job_LoadRecent` or `Job_AggregateGIM` successfully completes loading each aggregate table. Rows in this view are written once and are not updated.

Table 11 lists the `ADMIN_AGGREGATE_HISTORY` view columns and descriptions.

Table 11: ADMIN_AGGREGATE_HISTORY View Columns

Column Name	Data Type	Description
BATCH_ID	INTEGER	Numeric ID that uniquely identifies the execution instance of the job. This ID also appears in every row of the aggregate tables.
JOB_ID	VARCHAR	ID that uniquely identifies the execution instance of the job.
JOB_NAME	VARCHAR	The name of the job, such as `Job_AggregateGIM`.
JOB_VERSION	VARCHAR	The version of the job, such as 7.6.000.10.
DBCONNECTION	VARCHAR	The connection name of the data source or target.
OWNER	VARCHAR	The owner of the tables in the data source or target.
TENANT_KEY	INTEGER	The tenant to which the data belongs.
TABLE_NAME	VARCHAR	The name of the table into which the data is loaded.
START_TIME	DATETIME	The date and time the aggregation started (Genesys Info Mart Server time zone).

Table 11: ADMIN_AGGREGATE_HISTORY View Columns (Continued)

Column Name	Data Type	Description
DURATION	INTEGER	The duration of the aggregation in seconds.
ROW_COUNT_ADD	INTEGER	The number of new rows added.
ROW_COUNT_ DELETE	INTEGER	The number of rows that were deleted.

ADMIN_PURGE_HISTORY

This view provides information about data purged from each Info Mart database table. A row is added to this view after `Job_MaintainGIM` successfully completes purging each fact table. Rows in this view are written once and are not updated.

Table 12 lists the `ADMIN_PURGE_HISTORY` view columns and descriptions.

Note: `ADMIN_PURGE_HISTORY` reports the actions that are performed by `Job_MaintainGIM` only. When `Job_MaintainGIM` flags rows that are eligible for purging, `ADMIN_PURGE_HISTORY` does not indicate when the data is actually purged.

Table 12: ADMIN_PURGE_HISTORY View Columns

Column Name	Data Type	Description
JOB_ID	VARCHAR	ID that uniquely identifies the execution instance of the job.
JOB_NAME	VARCHAR	The name of the job, such as `Job_MaintainGIM`.
JOB_VERSION	VARCHAR	The version of the job, such as 7.6.000.10.
START_TIME	DATETIME	The date and time the step started (Genesys Info Mart Server time zone).
DURATION	INTEGER	The duration of the step in seconds.
DBCONNECTION	VARCHAR	The connection name of the data source or target.
OWNER	VARCHAR	The owner of the tables in the data source or target.
TENANT_ID	INTEGER	The tenant for which data is purged. The value will match `TENANT_KEY` from the `TENANT` table.

Table 12: ADMIN_PURGE_HISTORY View Columns (Continued)

Column Name	Data Type	Description
DAYS_TO_KEEP_FACTS	INTEGER	The value of the configuration option that specifies the number of days that the fact or aggregate should be retained in the Info Mart database.
MARKED_ONLY	INTEGER	Indicates whether data was purged or whether it was flagged as eligible to be purged. If the `purge-action-is-delete` configuration option is set to `TRUE`, then `MARKED_ONLY` is set to `0`. If the `purge-action-is-delete` configuration option is set to `FALSE`, then `MARKED_ONLY` is set to `1`.
CALCULATED_PURGE_DATE	DATETIME	The latest start date and time that facts or aggregates are considered eligible to be purged, using the configuration option that specifies the number of days that the fact or aggregate should be retained in the Info Mart database.
TABLE_NAME	VARCHAR	The name of the table from which data is purged.
ROW_COUNT	INTEGER	The number of rows purged (or flagged as eligible to be purged).

CV_MISSING_CFG_OBJECTS

This view provides information about configuration objects that are missing from the Staging Area database. A row is added to this view when an unresolved reference is identified in the extracted data from any data source. Rows in this view are written once and are not updated. You can use this information to determine which configuration objects are missing in the source IDB from which Genesys Info Mart extracts configuration details.

Table 13 lists the `CV_MISSING_CFG_OBJECTS` view columns and descriptions.

Table 13: CV_MISSING_CFG_OBJECTS View Columns

Column Name	Data Type	Description
ERROR_TYPE	VARCHAR	The type of an unresolved configuration object, such as `Missing DN (DN DBID).`
MISSING_DBID	INTEGER	The database identifier (DBID) of the missing object, as stored in the Configuration Database.

Chapter

4 Managing Data Sources

Purging old data from data source databases should be performed periodically to prevent database size from growing too large. For more information about purging the Genesys Info Mart database, see "Purging the Genesys Info Mart Database" on page 94.

This chapter provides some recommendations for migrating and purging databases that Genesys Info Mart uses as data sources. It also provides recommendations for stopping and restarting a Multimedia ICON to avoid data quality issues within Genesys Info Mart.

This chapter provides information on the following:

Migrating Data Sources

There are certain special requirements to consider when you upgrade Interaction Concentrator.

Upgrading the ICON Application

If you are upgrading Interaction Concentrator, and Genesys Info Mart has already been extracting data from the Interaction Database (IDB) into which the existing Interaction Concentrator (ICON) application stores data, do not create a new ICON `Application` object in the framework configuration. Instead, use the existing application in the framework configuration when you install the Interaction Concentrator upgrade. Upgrading the existing ICON `Application` object ensures consistent data in the IDB, if the upgraded ICON application subsequently updates data that was inserted by the previous

ICON application. If you do not use the same ICON `Application` object, Genesys Info Mart will not be able to extract IDB data that was inserted by the previous ICON application and updated by the upgraded ICON application.

Upgrading an Outbound Contact Application

The following applies if you are upgrading Outbound Contact (for example—during migration from a previous release, or when installing a newer version of the same release) and your Genesys Info Mart deployment extracts Outbound Contact details in an HA data extraction topology. If Genesys Info Mart has already extracted Outbound Contact details from the HA pair of IDBs based on events ICON received from the existing Outbound Contact application, do not create a new Outbound Contact `Application` object in the Configuration Layer. Instead, use the existing Outbound Contact `Application` object in the Configuration Layer when you install the Outbound Contact upgrade. If for any reason, you must create a new Outbound Contact `Application` object in the Configuration Layer, then you must also create a new pair of ICON `Application` objects, and a new HA pair of Outbound Contact details IDBs for Genesys Info Mart to extract. See the "Configuring HA Data Extraction of Outbound Contact Details in a New Deployment" section of the *Genesys Info Mart 7.6 Deployment Guide* for information about how to configure HA data extraction for Outbound Contact details.

Upgrading the Merge Staging Area Database Schema

In Genesys Info Mart deployments that require multi-IDB merge, you must upgrade the Merge Staging Area database schema if the IDB upgrade involved changes to the `gsysIRMerge` stored procedure or to the tables that the `gsysIRMerge` procedure needs. To ensure that you do not inadvertently omit performing a required upgrade, Genesys recommends that you always upgrade the Merge Staging Area database schema if, as part of the Interaction Concentrator upgrade, you had to upgrade the IDB that stores voice details.

To upgrade the Merge Staging Area database schema:

1. In Configuration Manager, set the `run-scheduler` option to `FALSE` in the `[schedule]` section of the Genesys Info Mart application.
2. Wait until there are no ETL jobs running. Verify this in the `Status` tab window in the Genesys Info Mart Administration Console.
3. Upgrade the Merge Staging Area database schema by using the SQL scripts that are provided with the Interaction Concentrator software.

   For more information about using the scripts to perform the upgrade:
   - To apply an Interaction Concentrator hot fix, see the hot fix Deployment Procedure.

- To upgrade Interaction Concentrator from an earlier release, see the Interaction Concentrator chapters in the *Genesys Migration Guide*.

4. In Configuration Manager, reset the `run-scheduler` option to `TRUE` in the `[schedule]` section of the Genesys Info Mart application.

   The Genesys Info Mart Server will resume normal ETL scheduling.

Moving Data Sources

Genesys Info Mart identifies a data source by the database connection information in the Database Access Point (DAP). Generally, Genesys Info Mart does not allow source databases to be moved when this move results in a change in the database connection information. Moving a data source can also result in data issues, such as job failures and stranded or duplicate data. However, there are circumstances where moving a data source may be needed, such as disaster recovery or system maintenance. With the introduction of High Availability in Genesys Info Mart release 7.6, it is now possible to move the data source and to change the database connection information for the following deployment topologies:

- "HA of Configuration Details"
- "Non-HA of Configuration Details"
- "HA of Outbound Contact Details"

It is also possible to move the data source and change the database connection information for the following deployment topologies:

- Voice Details (HA and non-HA, single IDB or multiple IDBs)
- Multimedia Details (non-HA)
- Outbound Contact Details (non-HA)

When an IDB moves to a different host, or there is a change in `DBCONNECTION` or `OWNER`, Genesys Info Mart sees this IDB as a new source and attempts to extract all available data from that source starting from the beginning.

To prevent this re-extraction from the beginning, Genesys Info Mart recognizes that the database connection information represented by `DBCONNECTION` and `OWNER` has changed, logs message `55-21003`, and does not run `Job_ExtractICON`.

Note: Genesys recommends setting an alarm condition for this log message—should this situation occur, you would be notified.

In order for Genesys Info Mart to correctly recognize the new location of the IDB as an existing data source rather than a new one, updates to metadata in the Staging Area database need to be made.

Please contact Genesys Technical Support for the correct procedure to move the ICON data sources listed above.

Warning! Genesys Info Mart is not able to detect when a move has occurred for the Stat Server or GVP VAR databases. If you plan on moving either of these data sources, please contact Genesys Technical Support *BEFORE* you start.

Configuration Details Data Source

The following sections examine things to consider and the steps you need to follow to successfully move your Configuration details data source (the target of the `ICON_CFG` role) in both an HA and Non-HA deployment.

Planning Considerations

Before you attempt to move your data source, consider the following:

- Be sure to carefully review all of the following steps before performing any of them.
- Plan a time when you can complete the installation in one continuous session. Several steps can take significant time to complete, and this potentially affects the availability of Info Mart data.
- Some steps, such as configuring new Interaction Concentrator applications and DAPs or creating new IDBs, can be done at any time; while others, such as stopping and starting Interaction Concentrator or Genesys Info Mart Server should be done during a planned maintenance window to prevent a loss of data and minimize the size of the backlog that the ETL must process when it is restarted.
- The following procedures only apply to installations where the DAP for Configuration details contains the role `ICON_CFG` and no others.

Moving a Configuration Details Data Source

HA of Configuration Details

In a current deployment which includes HA of Configuration details, to replace one side of the HA pair with a new `ICON_CFG` data source, perform the following:

1. Create a new Interaction Concentrator application recording only Configuration details. For more information, see the "Configuring ICON application to capture configuration details" section of the *Genesys Info Mart 7.6 Deployment Guide*.
2. Create a new IDB in the desired location and associate it with the new Interaction Concentrator application. For more information, see the "preparing IDBs" section of the *Genesys Info Mart 7.6 Deployment Guide*.

3. Start Interaction Concentrator to populate the new IDB with Configuration details.
4. Create a new data extraction DAP. Configure the `ICON_CFG role` in the `[gim-etl]` section to represent the new data source of Configuration details. For more information, see the "Configuring DAPs for ICON Configuration details HA" section of the *Genesys Info Mart 7.6 Deployment Guide*. Make sure the values for the `ha-pair-id` and `ha-pair-primary` options match the values of the old DAP for `ICON_CFG` that you are replacing.
5. Complete a final intraday ETL cycle (all extract jobs, `Job_TransformGIM`, and `Job_LoadRecent`).
6. When all jobs have successfully completed, shut down Genesys Info Mart Server.
7. Add the new DAP to the `Connections` tab of the Genesys Info Mart `Application` object.
8. Remove the old DAP for the `ICON_CFG` role, from the `Connections` tab of the Genesys Info Mart `Application` object.
9. Change the `run-scheduler` option to `FALSE`, in the `[schedule]` section of the Genesys Info mart `Application` object.
10. Start Genesys Info Mart Server.
11. From Genesys Info Mart Administration Console run `Job_ExtractICON` for `ALL SOURCES`.
12. When the extract job completes, the data from the new IDB should be reconciled with the existing data in the Staging Area database.
13. Change the `run-scheduler` option in the `[schedule]` section to `TRUE`.
14. Resume normal operations.

Non-HA of Configuration Details

In a current deployment which does not include HA of Configuration details, to move an `ICON_CFG` data source, perform the following:

1. Create a new Interaction Concentrator application recording only Configuration details. For more information, see the "Configuring ICON application to capture configuration details" section of the *Genesys Info Mart 7.6 Deployment Guide*.
2. Create a new IDB in the desired location and associate it with the new Interaction Concentrator application. For more information, see the "preparing IDBs" section of the *Genesys Info Mart 7.6 Deployment Guide*.
3. Start the new Interaction Concentrator application to populate the new IDB with Configuration details.
4. Create a new data extraction DAP. Configure the `ICON_CFG role` in the `[gim-etl]` section to represent the new data source of Configuration details. For more information, see the "Configuring DAPs for ICON

Configuration details HA" section of the *Genesys Info Mart 7.6 Deployment Guide*. Configure the `ha-pair-id` option and set the `ha-pair-primary` option to `TRUE`.

5. Complete a final intraday ETL cycle (all extract jobs, `Job_TransformGIM`, and `Job_LoadRecent`).
6. When all jobs have successfully completed, shut down Genesys Info Mart Server.
7. Add the new DAP to the `Connections` tab of the Genesys Info Mart `Application` object.
8. In the old DAP with `ICON_CFG` role, set the `ha-pair-id` option to the same value you configured in the new DAP and set the `ha-pair-primary` option to `FALSE`.
9. Change the Genesys Info Mart `Application` option `ha-cfg-all-connections-required` from `TRUE` to `FALSE`.
10. Change the `run-scheduler` option to `FALSE`, in the `[schedule]` section of the Genesys Info mart `Application` object.
11. Start Genesys Info Mart Server.
12. From Genesys Info Mart Administration Console run `Job_ExtractICON` for `ALL SOURCES`.
13. When the extract jobs complete, the data from the new IDB should be reconciled with the existing data in the Staging Area database.
14. Stop Genesys Info Mart Server.
15. Remove the old DAP for the `ICON_CFG` role, from the `Connections` tab of the Genesys Info Mart `Application` object.
16. Remove the `ha-pair-id` and `ha-pair-primary` options from the new DAP with the `ICON_CFG` role.
17. Change the `run-scheduler` option in the `[schedule]` section to `TRUE`.
18. Start Genesys Info Mart Server.
19. Resume normal operations.

Outbound Contact Details Data Source

The following sections examine things to consider and the steps you need to follow to successfully move your Outbound Contact details data source in an HA of Outbound Contact details deployment.

Planning Considerations

Before you attempt to move your data source, consider the following:

- Be sure to carefully review all of the following steps before performing any of them.

- Plan a time when you can complete the installation in one continuous session. Several steps can take significant time to complete, and this potentially affects the availability of Info Mart data.
- Some steps, such as configuring new Interaction Concentrator applications and DAPs or creating new IDBs, can be done at any time; while others, such as stopping and starting ICON, Genesys Info Mart Server or Outbound Contact campaigns, should be done during a planned maintenance window to prevent a loss of data and minimize the size of the backlog that the ETL must process when it is restarted.

Moving an Outbound Contact Details Data Source

HA of Outbound Contact Details

In a current deployment which includes HA of Outbound Contact details, to replace one side of the HA pair with a new data source, perform the following:

1. Create a new Interaction Concentrator application recording only Outbound Contact details. For more information, see the "Configuring ICON application to capture Outbound Contact details" section of the *Genesys Info Mart 7.6 Deployment Guide*.
2. Create a new IDB in the desired location and associate it with the new Interaction Concentrator application. For more information, see the "preparing IDBs" section of the *Genesys Info Mart 7.6 Deployment Guide*.
3. Stop all of your Outbound Contact campaigns.
4. Start Interaction Concentrator to populate the new IDB with Outbound Contact details.
5. Restart your Outbound Contact campaigns.
6. Create a new data extraction DAP. Configure the `ICON_OCS role` in the `[gim-etl]` section to represent the new data source of Outbound Contact details. For more information, see the "Configuring DAPs for ICON Outbound Contact details HA" section in the *Genesys Info Mart 7.6 Deployment Guide*. Make sure the values for the `ha-pair-id` and `ha-pair-primary` options match the values of the old DAP for `ICON_OCS` that you are replacing.
7. Complete a final intraday ETL cycle (all extract jobs, `Job_TransformGIM`, and `Job_LoadRecent`).
8. When all jobs have successfully completed, shut down Genesys Info Mart Server.
9. Add the new DAP to the `Connections` tab of the Genesys Info Mart `Application` object.
10. Remove the old DAP for the `ICON_OCS` role, from the `Connections` tab of the Genesys Info Mart `Application` object.
11. Change the `run-scheduler` option to `FALSE`, in the `[schedule]` section of the Genesys Info mart `Application` object.

12. Start Genesys Info Mart Server.
13. From Genesys Info Mart Administration Console run `Job_ExtractICON` for `ALL SOURCES`.
14. Change the `run-scheduler` option in the `[schedule]` section to `TRUE`.
15. Resume normal operations.

Purging Data Sources

Interaction Database (IDB) size is one of the most significant factors that affect ICON operational performance. For this reason, ICON provides several stored procedures that purge old IDB data:

ICON Purge Procedures

- Available starting with ICON release 7.6.1:
  - `gsysPurge76`—Purges voice, Multimedia, Open Media, attached data, and agent login sessions. This is the preferred ICON purge procedure for these types of data, because it outperforms the purge procedures from previous ICON releases (see below—`gsysPurgeIR`, `gsysPurgeUDH`, and `gsysPurgeLS`).
- Available in ICON releases 7.6.1, 7.6.0, and 7.5:
  - `gsysPurgeIR`—Purges IR voice data. This procedure is required only for IDBs from which Genesys Info Mart extracts voice details. This procedure does not purge Multimedia data.
  - `gsysPurgeUDH`—Purges `UserData` history for voice interactions. This procedure is required only for IDBs from which Genesys Info Mart extracts voice details. This procedure does not purge Multimedia data.
  - `gsysPurgeLS`—Purges agent login session data. This procedure is required only for IDBs from which Genesys Info Mart extracts voice or Multimedia details.
  - `gsysPurgeOS`—Purges Outbound-specific data. This procedure is required only for IDBs from which Genesys Info Mart extracts Outbound Contact details.

For more information about these stored procedures, see the *Interaction Concentrator 7.6 User's Guide.*

Input Parameters

Each ICON stored procedure requires an input parameter that specifies a retention period (the number of days preceding the current date, for which data is not to be purged). Data that is older than the current date minus the retention period parameter is purged.

`gsysPurge76` accepts a second input parameter, which controls whether ICON deletes all records older than the retention period or only terminated records older than the retention period. When you use `gsysPurge76`, Genesys recommends that you delete only terminated records.

Genesys Info Mart and ICON Purge Procedures

Genesys Info Mart does not provide automated purging of old IDB data. You must write your own program or stored procedure to invoke the ICON mechanisms 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.

This section provides recommendations for the execution of the ICON stored procedures to prevent accidental purging of data before Genesys Info Mart has an opportunity to extract it. In particular, this section provides guidance for selecting the smallest safe value that can be used for the retention period parameter.

IDB Data Retention

The amount of historical data you are able to retain in IDB and still preserve reasonable performance depends on the database server's hardware resources, such as memory and disk space, and disk subsystem performance. You should retain as much IDB data as you require to allow ongoing interactions to complete and as a hedge against operational problems. You should not retain more IDB data than you require for these purposes, because excess data slows system efficiency, increases backup requirements, and complicates troubleshooting.

In an environment where Genesys Info Mart is maintaining a regular schedule of ETL cycles that run every sixty minutes or less, the recommended value for the retention period parameter is between `7` and `30` days. Retaining fewer than `7` days could result in a loss of data.

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 `limit-extract-data` configuration option) while it works through a backlog of data. The IDB data retention period should be increased 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, the IDB data retention period should be increased 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) you should 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.

IDB Purge Frequency

Genesys Info Mart recommends that you run the ICON stored procedures 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 jobs, `Job_LoadGIM`, `Job_AggregateGIM`, and `Job_MaintainGIM`. ICON 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 intraday ETL cycle.

Determining the Retention Period

The retention period parameter for the `gsysPurge76` stored procedure is `<number of days>`. The retention period parameter for the other ICON purge procedures is `<count days>`. For a brief description of the available purge procedures, see "ICON Purge Procedures" on page 112.

The procedure documented below gives an estimate for the `<number of days>` or `<count days>` input parameter. This procedure takes into account the last time Genesys Info Mart successfully extracted all of the data for an ETL cycle and whether Genesys Info Mart is limiting the amount of data extracted while processing a backlog of data. It also includes a safety buffer of seven days.

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 seven days. If you are able to store more than seven 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.

The value that is returned from the procedure should be input for each of the ICON stored procedures listed on page 112.

Run the following RDBMS-specific SQL statements against your Staging Area database to return the minimum value for the `<number of days>` or `<count days>` input parameter. To ensure an accurate calculation, issue the SQL

statements prior to running each ICON stored procedure. Also make sure that you log in using the Staging Area database Owner account before issuing the statements.

Oracle

Issue the following SQL query against an Oracle-based Staging Area database to determine the value for the <number of days> or <count days> parameter:

```
with z as (
/* number of days since 1st extract preceding last completed transformation */
/* (and following the prior completed transformation) or 999, if none       */
   select coalesce((
      select round(current_date - min(start_time)) + 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 from dual
   union
   /* number of days since max target extract date if limiting extract data */
   select round(current_date -
               coalesce((TO_DATE('19700101000000','YYYYMMDDHH24MISS') +
                         NUMTODSINTERVAL(
                            (select maxutc from stg_extract_throttle),
                            'SECOND')
                        ), current_date) ) + 7 as x from dual
)
select max(x) as MINIMUM_DAYS from z ;
```

Microsoft SQL Server

Issue the following SQL query against a Microsoft SQL Server–based Staging Area database to determine the value for the <number of days> or <count days> parameter:

```
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
   union
/* number of days since maximum target extract date if limiting extract data */
   select coalesce(round(datediff(second, dateadd(second, maxutc,
                  '01/01/1970'), CURRENT_TIMESTAMP)/86400.0, 0) +7, 7)
      as x from stg_extract_throttle
) z;
```

IBM DB2

Issue the following SQL query against a DB2-based Staging Area database to determine the value for the <number of days> or <count days> parameter:

```
with z as (
/* number of days since 1st extract preceding last completed transformation */
/* (and following the prior completed transformation) or 999, if none */
   select coalesce((
      select round(days(current_date) - days(min(start_time)) +
                  float(midnight_seconds(current_time)
                  - midnight_seconds(min(start_time))) / 86400, 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 from admin_etl_job_history
   union
/* number of days since maximum target extract date if limiting extract data */
   select coalesce(round(float(
                      (days(current_date) - days(DATE('1970-01-01')))*86400
                      - maxutc) / 86400, 0) + 7, 7)
      as x from stg_extract_throttle
)
select decimal(max(x)) as MINIMUM_DAYS from z ;
```

Note: For information about purging Stat Server databases in legacy Reporting environments, see "Purging the Stat Server Database" on page 136.

Restarting a Multimedia ICON

The following sections examine the considerations and the steps that you need to follow to successfully stop and restart a Multimedia ICON to avoid data quality issues within Genesys Info Mart.

Special Considerations

If your environment contains an Interaction Concentrator that stores Multimedia details, the following section outlines which Interaction Concentrator options Genesys recommends that you configure, and what effects not configuring these options may have on your Genesys Info Mart data. It also discusses the consequences of missing or inaccurate information due to incorrect configuration, data collection during different ICON sessions, or data quality issues brought on during the restart of a Multimedia ICON, which need to be considered.

Note: The term *ICON session* is used to refer to the time span when ICON is connected to Interaction Server. For example—an ICON session starts when ICON is started and connects to Interaction Server, and ends when ICON is stopped, and disconnects from Interaction Server.

Recommended Configuration for a Multimedia ICON

Genesys Info Mart recommends that you configure the following options for your Multimedia Interaction Concentrator:

- `calls-in-the-past=true`

- `om-force-adata=true`

For more information about properly configuring these options, see the *Genesys Info Mart 7.6 Deployment Guide*.

If you do not configure `calls-in-the-past` as recommended, ICON will not record any data about multimedia interactions that started outside of the current ICON session. As a result:

- If the Interaction started during one ICON session (for example—before ICON was stopped) and was still active in a subsequent ICON session (for example—after ICON was restarted), ICON's IDB and Info Mart's intraday fact tables store reporting data for that Interaction only up to the time the first ICON session ended. The Interaction remains active in ICON's IDB and Info Mart's intraday fact tables, and will not move to Info Mart's historical fact tables.
- If the Interaction started outside of an ICON session (for example—while ICON was stopped) and was still active in a subsequent ICON session (for example—after ICON was restarted), ICON's IDB and Info Mart will not store any reporting data for that Interaction.

If you do not configure `om-force-adata` as recommended, ICON will not record initial values of multimedia interaction attributes in `GM_F_USERDATA` for any interaction that started outside of the current ICON session. As a consequence, the corresponding Interaction in the Info Mart database will have missing or incorrect data, most notably the Interaction Type dimension.

If you configure `calls-in-the-past` as recommended, ICON will record data about multimedia interactions that started outside of the current ICON session, beginning with the first party added to the interaction during the next ICON session. As a result:

- If the Interaction started during one ICON session (for example—before ICON was stopped) and was still active in a subsequent ICON session (for example—after ICON was restarted), ICON's IDB and Info Mart's intraday fact tables store reporting data up to the time the first ICON session ended, and resume storing reporting data beginning with the first interaction party added to the interaction during that subsequent ICON session. There may be missing data for the portion of the interaction that occurred outside of any ICON session, and data inaccuracies for any portion of the interaction that depended on that missing data. For example—some Interaction Segment Facts might be missing, an Interaction Segment Fact's Technical Descriptor might be incorrect, or an Interaction Segment Fact's end times or state counts and durations might be incorrect.
- If the Interaction started outside of an ICON session (for example—while ICON was stopped) and was still active in a subsequent ICON session (for example—after ICON was restarted), ICON's IDB and Info Mart's intraday fact tables store reporting data beginning with the first interaction party added to the interaction during that subsequent ICON session. There may be missing data for the portion of the interaction that occurred outside of any ICON session, and data inaccuracies for any portion of the interaction

that depended on that missing data. For example—Interaction Segment Facts might be missing, an Interaction Segment Fact's Technical Descriptor might be incorrect, or an Interaction Segment Fact's end times or state counts and durations might be incorrect.

If you configure `om-force-adata` as recommended, ICON will record initial values of multimedia interaction attributes in `GM_F_USERDATA`, for any interaction that started outside of the current ICON session. As a result, the corresponding Interaction in the Info Mart database will have correct data, most notably the Interaction Type dimension.

Note: Even if you configure `calls-in-the-past` and `om-force-adata` as recommended, the following consequences still can occur:

- If the Interaction started during one ICON session (for example—before ICON was stopped) and completed outside of an ICON session (for example—while ICON was stopped), ICON's IDB and Info Mart's intraday fact tables store reporting data up to the time the ICON session ended. The Interaction remains active in ICON's IDB and Info Mart's intraday fact tables, and will not move to Info Mart's historical fact tables.
- If the Interaction started and completed outside of an ICON session (for example—while ICON was down), ICON's IDB and Info Mart will have no reporting data for that Interaction.

Most of the data quality issues described in the "Recommended Configuration for a Multimedia ICON" section can be avoided for planned ICON outages by configuring ICON as recommended and following the procedure described in the next section, "Recommended Procedure for Restarting a Multimedia ICON". However, any unplanned ICON outage, such as a network connectivity issue between ICON and Interaction Server or an ICON crash, that occurs during the entire lifetime of a multimedia interaction can still result in the data quality issues described in the "Recommended Configuration for a Multimedia ICON" section.

Recommended Procedure for Restarting a Multimedia ICON

As noted above, some data quality issues regarding Multimedia interactions may occur when the interaction data is recorded during different ICON sessions—some information may be lost between these sessions. However, sometimes it is necessary to stop an ICON to apply upgrades or to allow for configuration changes. Most of these data quality issues can be avoided by

following the steps outlined in the procedure below for a planned ICON outage.

Note: It is best to perform the following 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.

Restarting a Multimedia ICON Procedure

Each of the steps in this procedure can be started when the previous step has successfully completed:

1. Stop Interaction Server prior to stopping the Multimedia ICON. (eServices recommends using SCI to start and stop Interaction Server.)
2. After allowing sufficient time for the Multimedia ICON to record all final queue and agent activity (for example—after Interaction Server is fully stopped), stop the Multimedia ICON.
3. Restart the Multimedia ICON.
4. Restart Interaction Server. (eServices recommends using SCI to start and stop Interaction Server.)

By first stopping Interaction Server, interactions are moved to a *home* state—returning them to Interaction Queues or Workbins, and no interaction activity can occur during this time. Since ICON is restarted first, it is able to record subsequent interaction activity.

Virtual Queue Activity

Following Steps 1 to 4 above to restart a Multimedia ICON eliminates most data quality issues that may occur when a Multimedia interaction spans more than one ICON session. However, it is possible that some Virtual Queue activity that occurs during this procedure could be lost.

- Interactions in Virtual Queues when Interaction Server stopped are returned to Interaction Queues. The events sent by Universal Routing Server indicating that the interactions were cleared from Virtual Queues may not reach Interaction Server before it stops its connection to Universal Routing Server and the data may not show that the interaction left the Virtual Queue. As a result, the `MEDIATION_SEGMENT_FACT` in Genesys Info Mart associated with a Virtual Queue that was not seen to complete may have incomplete information and will have to be artificially terminated. Since no agent can be logged in to Multimedia at this time, no interactions can be routed from the Virtual Queue to an agent, so any loss of information should not be significant.
- When Interaction Server is restarted, some Virtual Queue or Routing Strategy activity may occur before ICON successfully reestablishes its connection to Interaction Server. But, when following the procedure above,

there will be more than enough time for ICON to establish its connection to Interaction Server before agents begin logging in, so no agent related interaction data is lost, not even for agents who receive interactions because of this Routing activity.

If there was Virtual Queue activity that was not recorded in the IDB during this startup period, then there would also be no `MEDIATION_SEGMENT_FACT` data in Genesys Info Mart for the associated Virtual Queue activity. If interactions that were placed in these Virtual Queues were simply cleared from the Virtual Queue because no agents are logged in and available, then any resulting loss of data should be insignificant. If the interactions remain in these Virtual Queues until the agents do become available, then there may be some loss of mediation information in Genesys Info Mart, but there will be no loss of the interaction segment data capturing the handling of these interactions by agents, or any subsequent data.

Similarly, if there is Routing Strategy activity that is not recorded in the IDB during this startup period, then there will be no Interaction Segment Fact capturing this Routing Strategy activity, and there will be no loss of the Interaction Segment data capturing the handling of these interactions by agents, or any subsequent data.

Chapter

5 Troubleshooting ETL Jobs

The ETL 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 Interdependencies" on page 87, 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.

No matter which method you use to schedule or execute jobs, recovering from any failed job requires using the Genesys Info Mart Administration Console to manually restart the failed job or jobs. Use the information in this section to help you determine the cause of a job failure and how to recover from it.

This chapter also provides recommendations for situations in which the 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 provides information on the following:

Types of Errors

Several types of errors can cause an ETL job to fail:

Configuration Errors

- The Genesys Info Mart Server could not communicate with the Genesys Configuration Server to read its configuration parameters. The Configuration Server is not running or the network connection between the Genesys Info Mart Server and the Configuration Server is down.
- The Genesys Info Mart Server or an ETL job encountered errors in the Genesys Info Mart configuration, such as missing parameters or invalid parameter values. See also "Unresolved References to Configuration Objects" on page 125.

Database Connection Errors

- The ETL 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 ETL 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 ETL 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, the ETL job `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.
- In a large-scale deployment using high levels of parallel processing, the ETL job or Genesys Info Mart Server might not be able to connect to a source or target database because the RDBMS server allows an insufficient number of concurrent database connections.

When you encounter error messages about insufficient available database connections, use the tools provided by your RDBMS to clean up inactive sessions and to review, and potentially increase, the number of concurrent database connections that the RDBMS server will allow.

To test connectivity for each database access point (DAP), you should always run the `gim_etl_config_check` command. It will provide additional detailed information about any possible connection errors.

SQL Error

- The ETL 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.
- 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. Running the `gim_etl_config_check` command will log the versions of all JDBC drivers that have previously been installed.

ETL Job Error

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

Unresolved References to Configuration Objects

- `Job_TransformGIM` encountered an unresolved reference to a configuration object in the extracted source data. In this situation, `Job_TransformGIM` logs one of the following error messages, and shuts down:
  - `GIM_ETL_MISSING_REQ_CFG_OBJ (20069)`
  - `GIM_ETL_MISSING_REQ_CFG_OBJN (20070)`

  In order for you to be notified when this situation occurs, Genesys strongly recommends that you use Solution Control Interface (SCI) to set alarms for these error messages. The generated log message identifies the missing configuration object that caused the job failure. You will have to reset the alarm manually once the problem is resolved.

  When you encounter these error messages, follow the recovery procedure on page 130.

Unrecognized Call Flow Errors

- Records extracted from ICON that do not pass a certain criterion (`IR_CALL_FLOW_UNRECOGNIZED`) are written to the `STG_ICON_DATA_DISCARDS` table.

- Due to deployment-specific factors, such as voice interaction topologies, data source configuration, or failures in the collection of source data, Job_TransformGIM might not be able to transform some extracted voice interactions. When this occurs, the job:
  - Logs message 34204 - GIM_TRANSFORM_BADDATA.
  - Records some information about the voice interaction in Staging Area table STG_ICON_DATA_DISCARDS.
  - Discards the voice interaction.
- Information in STG_ICON_DATA_DISCARDS can be used to investigate why some voice interactions were discarded:
  - TABLE_NAME - indicates the IDB table that had problematic data.
  - ROOTIRID - indicates the IRID of the root G_IR row that represents the voice interaction.
  - REASON - indicates the reason the voice interaction was discarded. See Table 14 for a description of some of the reasons.
- Data is retained in STG_ICON_DATA_DISCARDS for seven days.

Table 14: Sample STG_ICON_DATA_DISCARDS Table Entries

Table	Reason	Details
G_PARTY_HISTORY	MISSING_CREATED_DATA	The created state is missing from the G_PARTY_HISTORY data. This is important because it is used to determine the resource role for the G_PARTY row.
G_PARTY_HISTORY	MISSING_TERMINATED_DATA	The terminated state is missing from the G_PARTY_HISTORY data. This is important because it is used to determine the technical result for the G_PARTY row.
G_PARTY	MISSING_CREATED_DATA	The createdTs in a G_PARTY row is null.
G_PARTY	MISSING_TERMINATED_DATA	The terminatedTs in G_PARTY row is null.
G_PARTY	NO_PARTY_DATA	A G_CALL row extracted from ICON has no related G_PARTY rows.
G_CALL	NO_CALL_DATA	A G_IR row was extracted that had no related G_CALL rows.
G_PARTY_STAT	MISSING_PARTY_STAT_DATA	A G_PARTY row was extracted that had no matching G_PARTY_STAT data.

Table 14: Sample STG_ICON_DATA_DISCARDS Table Entries (Continued)

Table	Reason	Details
G_IR	ERROR_PROCESSING_ICON_DATA	An error was encountered by the transformation process for this interaction. This is typically a Java exception that is caught and handled by the transformation logic. The exception is logged and processing continues.
G_IS_LINK_HISTORY	MISSING_CALL_DATA_FOR_IS_LINK	IS_LINK data was encountered that did not have a CALLID associated with the link.
Multiple Tables	IR_CALL_FLOW_UNRECOGNIZED	Genesys Info Mart has a count of G_CALL rows that should be processed for an interaction. If the code that generates the interaction segment facts from the related G_CALL/G_PARTY data does not process every call in that count, then the interaction is discarded (the data could not be associated correctly). The following are the main reasons interactions are discarded: • G_CALL data Genesys Info Mart is expecting to see via IS_LINK data is missing. This occurs in deployments where ICON is not monitoring all T-Servers and calls are transferred or routed between a monitored and unmonitored T-Server and back. In this type of deployment both sides of an IS_LINK pair (callid from each T-Server) is not available to Genesys Info Mart and the attempt to follow the call data is disrupted. This is particularly true when the call flow bounces from monitored to unmonitored to monitored T-Servers.

Table 14: Sample STG_ICON_DATA_DISCARDS Table Entries (Continued)

Table	Reason	Details
Multiple Tables (continued)	IR_CALL_FLOW_UNRECOGNIZED (continued)	• The ICON merge procedure associated an incorrect rootirid to the calls in the interaction. Genesys Info Mart cannot traverse the related calls correctly as the transformation logic starts in the wrong call and the data that indicates to the transformation logic to associate the correct calls together is not as expected. This can occur when two calls have the same createdTs, which is typical in Network deployments where calls exist in the network sub-second and are routed to a premise T-Server. This can also occur when the T-Server boxes in a deployment are not time synchronized, because the call with the smallest createdTs will be pegged by ICON as the root call even when another call was the actual root call. • If for some reason an interaction is created without segment ordinal 1, the interaction is discarded.

Resources to Consult for Additional Information

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

- *Genesys Info Mart 7.6 Deployment Guide*—Contains information to help you tune performance parameters for your Staging Area and Info Mart databases and correct errors in Genesys Info Mart configuration parameters. This guide also contains information that you can use to configure the ICON, GVP VAR, and Stat Server (for legacy environments only) applications that populate the databases from which you want to extract data.
- *Genesys Info Mart 7.6 Operations Guide* (this document)—Contains information that you can use to correct errors in the ETL 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 ETL jobs. The logs indicate configuration errors, when an ETL job begins, when an ETL 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 ETL jobs. When a job fails, view these logs to obtain detailed information about the failure.
- Staging Area Administrative Views—Contain a history of ETL job execution and a list of unresolved configuration object references. You can use this information to track historical ETL performance and diagnose ETL problems. For views descriptions, see "Staging Area Administrative Views" on page 96.
- 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 ETL jobs fail:

General Recommendations

When the Genesys Central Logger or Genesys Info Mart Administration Console indicates that an ETL 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 ETL job.

Genesys Info Mart Server will not launch a scheduled ETL job until the jobs it depends on complete successfully. In the event of an ETL job failure, the

Genesys Info Mart Server will not launch any other scheduled job until the failed job is restarted and completes successfully. Use the Genesys Info Mart Administration Console to restart the failed ETL job (see Chapter 3, "Executing and Scheduling ETL Jobs," on page 76).

When a job fails, do *not* stop the Genesys Info Mart Server. If you experience a job failure, it is very important for Genesys Info Mart Server to keep performing intra-IDB merge for voice details. After a job fails, make sure that:

1. Genesys Info Mart Server continues running.
2. The `ir-merge-interval` option in the `[gim-etl]` section of the Genesys Info Mart application is set to a non-zero value. (For example—keep the default of 5 minutes.)

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 130 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 Unresolved References to Configuration Objects

When Genesys Info Mart reports an unresolved reference to a configuration object in the extracted source data (see the error messages on page 125), use the following procedure to recover from this situation:

Note: The following procedure assumes that you have not configured the Genesys Info Mart Server to ignore missing configuration objects, and either you have not configured automatic job retry, or you have configured it, but all retries have failed. For more information about configuring Genesys Info Mart to ignore missing configuration objects, and the deployments for which this type of configuration might be suitable, see "Ignoring Configuration Errors" on page 43.

1. The Staging Area database provides a database view, `CV_MISSING_CFG_OBJECTS` (see page 103), to identify the unresolved references in all the extracted source data. Use this information to

determine which configuration objects are missing in the source IDB from which Genesys Info Mart extracts configuration details. The view returns each missing object's type and configuration database identifier (DBID).

2. The source IDB might be missing some configuration details because of an application or network connection outage. Check to make sure that the following applications and network connections are operational:
   - The Interaction Concentrator (ICON) application that stores the configuration details in the IDB
   - The network connection between ICON and Configuration Server
   - The DB Server application that ICON uses to store the configuration details in the IDB
   - The network connection between the DB Server and the IDB
3. If you found and corrected an issue with any of the previously listed applications or network connections, some or all of the missing configuration details might now be stored in the IDB. Make sure to allow enough time for ICON to store any information that it had previously buffered into its IDB, before you continue.
4. Use the Genesys Info Mart Administration Console to rerun the following ETL jobs in the indicated order:
   a. `Job_ExtractICON`, specifying the appropriate connection. If your deployment extracts IDB data from more than one DAP, specify the `<All Sources>` connection.

   Note: When you run `Job_ExtractICON` for `All SOURCES`, you may see a popup box that indicates that some extraction jobs are not eligible to run until after `Job_TransformGIM` has completed.This popup box does not indicate an error, but is a reminder that a subset of extraction jobs has already ran after the last time `Job_TransformGIM` completed, and only the remaining extraction jobs will now run.

   b. `Job_TransformGIM`.
5. If `Job_TransformGIM` still logs an error message and shuts down because of an unresolved reference to a configuration object, the source IDB might be missing some configuration details, because ICON (or the network connection between Configuration Server and ICON) was down for too long. In this case, follow the instructions in the *Interaction Concentrator 7.6 User's Guide* to resynchronize the IDB with the contents of the Configuration Server database. Based on the information in that document, make sure that the resynchronization has completed before you continue.
6. Use the Genesys Info Mart Administration Console to rerun the following ETL jobs in the indicated order:
   a. `Job_ExtractICON`, specifying the appropriate connection. If your deployment extracts IDB data from more than one DAP, specify the `<All Sources>` connection.

> **Note:** When you run Job_ExtractICON for All SOURCES, you may see a popup box that indicates that some extraction jobs are not eligible to run until after Job_TransformGIM has completed. This popup box does not indicate an error, but is a reminder that a subset of extraction jobs has already ran after the last time Job_TransformGIM completed, and only the remaining extraction jobs will now run.

b. Job_TransformGIM.

7. If Job_TransformGIM still logs an error message and shuts down because of an unresolved reference to a configuration object, the missing configuration objects cannot be recovered. In this case, you must force the ETL to process the extracted data.

   Depending on the type of data that is being processed, the ETL will either discard the data that contained unresolved references or transform and load that data, providing references to the appropriate "unknown" dimension. To force the ETL to process the data, use the Genesys Info Mart Administration Console to rerun Job_TransformGIM, specifying -ignoreMissingConfigObjs as the parameter (see page 42).

   After Job_TransformGIM runs successfully, the Genesys Info Mart Server will resume normal ETL scheduling.

8. You should note that there might be additional unresolved references to configuration objects in source data that the ETL has not yet extracted. In this case, you must repeat Step 7 until the ETL has processed all the source data that contained unresolvable references to configuration objects.

High Availability Recommendations

> **Note:** The following Genesys recommendations apply only to high availability of Configuration and Voice details. These recommendations are not needed for high availability of Outbound Contact details.

By default, Job_ExtractICON (with the role ICON_CORE or ICON_CFG) fails when it cannot connect to one of the IDBs from an HA pair in an HA configuration. In this situation, consider allowing the ETL temporarily to extract data from the remaining IDB.

> **Warning!** Extracting data from a single IDB in an HA configuration has a negative impact on data quality. You should do this only in a situation in which having no reporting data is worse than having reporting data with suspect data quality (that is, when it is known that the unavailable IDB will remain inaccessible for a long time).

If you decide to extract data from one data source out of the HA pair, follow these instructions:

1. In Configuration Manager, set the following options to `FALSE` in the `[gim-etl]` section of the Genesys Info Mart application:
   - `ha-cfg-all-connections-required`, if the unavailable IDB stores ICON configuration details
   - `ha-agent-all-connections-required`, if the unavailable IDB stores ICON voice details, including agent activity
2. In Configuration Manager, verify that the following options in the `[gim-etl]` section of the Genesys Info Mart application are configured appropriately for your requirements:
   - `ha-ir-extract-comparison-timeout`—Defines the timeout interval, in minutes, for which Genesys Info Mart will wait for voice details from one IDB, after it has read the data from the other IDB, for deduplication purposes. The default is 60 minutes. In the situation in which you are extracting data from only one IDB, consider reducing the value of this option (`0` means no timeout), so that voice data can be extracted faster into the Staging Area database.
   - `days-to-keep-stg-ha-ir-ids`—Defines the period, in days, for which Genesys Info Mart will remember data that it extracted from the single, available IDB. The default is 3 days.
     - If the unavailable IDB becomes available within this time period, Genesys Info Mart will compare voice detail records in the newly available IDB against the voice details it has already extracted, and will not extract duplicate voice details from the newly available IDB.
     - If the unavailable IDB becomes available after the expiry of this time period, Genesys Info Mart will extract all voice details from the newly available IDB. If the newly available IDB is not purged before the extraction, this will result in duplicated voice detail records.

     In the situation in which you are extracting data from only one IDB, consider increasing this option to a value that exceeds the number of days that you expect the other IDB to be unavailable.
3. When the unavailable IDB becomes available again, consider whether you need to clean up the data before Genesys Info Mart extracts it. In particular, if the IDB was unavailable for a period longer than the value configured for the `days-to-keep-stg-ha-ir-ids` option (see Step 2), Genesys recommends that you purge the voice details records from the unavailable IDB, to ensure that you do not duplicate data.
4. In the Genesys Info Mart Administration Console, manually run the `Job_ExtractICON`, specifying the `<All Sources>` connection.

   The Genesys Info Mart Server will now extract the data from the available IDB in the HA pair, and normal ETL job scheduling will resume.

5. When both IDBs in the HA pair are available, reset Genesys Info Mart `Application` options in Configuration Manager. In the `[gim-etl]` section:
   a. Set the following options to `TRUE`:
      - `ha-cfg-all-connections-required`, if the HA pair of IDBs stores ICON configuration details
      - `ha-agent-all-connections-required`, if the HA pair of IDBs stores ICON voice details, including agent activity

      The ETL will connect to both IDBs during the next run of the `Job_ExtractICON`.
   b. Set the `ha-ir-extract-comparison-timeout` option back to its normal setting.

Appendix

Using Stat Server in Legacy Environments

For customers whose legacy reports continue to use Stat Server data, this appendix provides information for extracting and purging Stat Server data as described in the following sections:

Overview

Starting with release 7.6, Genesys Info Mart extracts voice agent state and reason details from Interaction Concentrator. For the benefit of the customers who have invested in Stat Server–based reports with a previous release, Genesys Info Mart continues to provide data extraction of voice agent state and reason details from Stat Server database.

As a courtesy to those customers whose legacy reports continue to use Stat Server data, this appendix provides information for extracting and purging data from a Stat Server database. New Genesys Info Mart deployments must extract data related to voice agent activity from Interaction Concentrator 7.5 or later.

Job_ExtractSS

This job extracts new and changed data from one Stat Server database, for backward compatibility with legacy reporting environments that are based on Stat Server details. The job stores the extracted data in the Staging Area database. The Genesys Info Mart Server launches this job for each Stat Sever database during each intraday ETL cycle. After all the extraction jobs have completed successfully, the Genesys Info Mart Server launches the job that transforms all the extracted IDB, GVP VAR, and Stat Server data.

Note: Extracting voice resource state and reasons details from Stat Server is supported for legacy reporting environments only. If you do not wish to extract these details, do not configure a database access point (DAP) for the `STAT_SERVER` role. If your legacy reports continue to rely on Stat Server data, the Genesys Info Mart Server will continue to run `Job_ExtractSS` for each DAP with this role.

Planning for Latency

When planning the intraday ETL schedule, make allowances for data latency. You want to ensure that the Stat Server database has had sufficient time to record all the previous day's data before the final extraction begins.

Controlling the Rate of Data Extraction

You can configure Genesys Info Mart to control the rate of the data extracted from a Stat Server database at any one time. You need to configure the following dependent configuration options in section [`gim-etl`] to set the rate of extraction:

- `limit-extract-data`
- `extract-data-after-date`
- `extract-data-time-range-limit`
- `extract-data-time-range-units`

When configured, the rate of data extraction applies equally to all Genesys databases in the system; you cannot apply this function individually to any one Genesys Info Mart data source. Controlling the rate of data extraction occurs once per ETL cycle. To set these configuration options to control the rate of data extraction, see the *Genesys Info Mart 7.6 Deployment Guide*.

Purging the Stat Server Database

To purge Stat Server data, manually run the following SQL queries against the Stat Server database. These queries use the value of `<count days>` to specify a retention period (the number of days preceding the current date, for which data is not to be purged). Data older than the current date minus `<count days>` is purged.

Note: The following statements are valid only when they are executed against a Stat Server database that has been modified for use as a data source for Info Mart.

Oracle Purge Commands

To purge data from the LOGIN_TABLE, QINFO_TABLE, and/or STATUS_TABLE tables in an Oracle-based Stat Server database, issue the following command:

```
DELETE FROM <table_name> WHERE
     TO_CHAR(ROUND(CURRENT_DATE - <count_days>), 'YYYYMMDD') > PARTITION_KEY ;
```

To purge data from the VOICE_REASONS table, issue the following command:

```
DELETE FROM VOICE_REASONS WHERE
     ROUND((CURRENT_DATE - <count_days>) -
          TO_DATE('19700101000000','YYYYMMDDHH24MISS')) * 86400 > END_TIME;
```

Microsoft SQL Server Purge Commands

To purge data from the LOGIN_TABLE table in an Microsoft SQL Server–based Stat Server database, issue the following command:

```
DELETE FROM LOGIN_TABLE WHERE
     DATEDIFF(SECOND, '01/01/1970', CURRENT_TIMESTAMP - <count_days>) > TIME
```

To purge data from the QINFO_TABLE and/or STATUS_TABLE tables, issue the following command:

```
DELETE FROM <table_name> WHERE
     DATEDIFF(SECOND, '01/01/1970', CURRENT_TIMESTAMP - <count_days>) > ENDTIME
```

To purge data from the VOICE_REASONS table, issue the following command:

```
DELETE FROM VOICE_REASONS WHERE
     DATEDIFF(SECOND, '01/01/1970', CURRENT_TIMESTAMP - <count_days>)
          > END_TIME
```

DB2 Purge Commands

To purge data from the LOGIN_TABLE, QINFO_TABLE, and/or STATUS_TABLE tables in a DB2-based Stat Server database, issue the following command:

```
DELETE FROM <table_name> WHERE
     (SUBSTR(CHAR((CURRENT_DATE - <count_days> DAYS),ISO),1,4) ||
      SUBSTR(CHAR((CURRENT_DATE - <count_days> DAYS),ISO),6,2) ||
      SUBSTR(CHAR((CURRENT_DATE - <count_days> DAYS),ISO),9,2) )
           > PARTITION_KEY;
```

To purge data from the VOICE_REASONS table, issue the following command:

```
DELETE FROM VOICE_REASONS WHERE
     (DAYS(CURRENT_DATE - <count_days> DAYS) -
          DAYS(DATE('1970-01-01'))) * 86400 > END_TIME;
```

Index

Symbols

A

C

D

E

F

G

H

I

J

L

M

N

O

P

R

S

T

U

V

W