Genesys Info Mart 8.1

Deployment Guide

About Genesys

Genesys is the world's leading provider of customer service and contact center software—with more than 4,000 customers in 80 countries. Drawing on its more than 20 years of customer service innovation and experience, Genesys is uniquely positioned to help companies bring their people, insights and customer channels together to effectively drive today's customer conversation. Genesys software directs more than 100 million interactions every day, maximizing the value of customer engagement and differentiating the experience by driving personalization and multi-channel customer service—and extending customer service across the enterprise to optimize processes and the performance of customer-facing employees. Go to www.genesys.com for more information.

Each product has its own documentation for online viewing at the Genesys Documentation 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 and the Genesys logo are registered trademarks of Genesys Telecommunications Laboratories, Inc. All other company names and logos may be trademarks or registered trademarks of their respective holders.

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.

Customer Care from Genesys

If you have purchased support directly from Genesys, please contact Genesys Customer Care. Before contacting Customer Care, please refer to the ***Genesys Care Program 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.genesys.com

Document Version: 81gim_dep_07-2014_v8.1.401.00

Table of Contents

List of Procedures

Preface

Welcome to the *Genesys Info Mart 8.1 Deployment Guide*. This document describes the procedures that you must complete in order to configure and install Genesys Info Mart 8.1 and the Genesys Info Mart Manager or Genesys Info Mart Administration Console. It also provides in-depth information about Genesys Info Mart functioning.

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

- Overview of Genesys Info Mart architecture and functionality
- Pre-installation considerations
- Installation instructions

This document is valid only for the 8.1 releases of this product.

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

This preface contains the following sections:

For information about related resources and about the conventions that are used in this document, see the supplementary material starting on page 469.

About Genesys Info Mart

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

Genesys Info Mart includes a server component, graphical user interface (GUI) for managing Genesys Info Mart, and database. The Genesys Info Mart Server

runs a set of predefined jobs that execute extract, transform, and load (ETL) processes to:

- Extract data that has been gathered by Interaction Concentrator (ICON) from data sources such as Configuration Server, T-Server, Interaction Server, and Outbound Contact Server. Genesys Info Mart stores this low-level interaction data, which is consolidated from Interaction Concentrator databases (Interaction Databases [IDBs]), in the Info Mart database.
- Transform the low-level interaction data and load it into a dimensional model (or star schemas) in the Info Mart database.

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

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

Intended Audience

This document is primarily intended for database administrators and system administrators. It has been written with the assumption that you have a basic understanding of:

- Relational database management systems (RDBMSs)
- Network design and operation
- Computer-telephony integration (CTI) concepts, processes, terminology, and applications
- Your own network and database configurations

Making Comments on This Document

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

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

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

Contacting Genesys Customer Care

If you have purchased support directly from Genesys, please contact Genesys Customer Care.

Before contacting Customer Care, please refer to the *Genesys Care Program Guide* for complete contact information and procedures.

Document Change History

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

New in Document Version v8.1.401.00

This document has been updated to support Genesys Info Mart release 8.1.4. The following topics have been added or changed since the previous release of this document:

Documentation Updates to Support New Features and Software Modifications

- A new subsection, "New in Release 8.1.4" on page 46, summarizes the feature enhancements that Genesys Info Mart 8.1.4 provides.
- The minimum Interaction Concentrator release requirement has been updated in Table 1 on page 64.
- In addition to the links provided for Genesys Info Mart Manager in "New in Release 8.1.4" on page 46, information throughout the document has been updated as required to include Genesys Info Mart Manager. In particular:
  - In Chapter 1 on page 31, the general descriptions of Genesys Info Mart components and architecture have been modified.
  - The Genesys Administrator Extension (GAX) 8.5 release requirement has been added to Table 1 on page 64.
  - A note in Task Summary: Enabling Access to Source and Target Databases, on page 209 confirms that Genesys Info Mart Manager does not require a DB Server and database access point (DAP) to access the Info Mart database directly.
  - "Installing Genesys Info Mart Manager" on page 330 provides full instructions for installing Genesys Info Mart Manager, and "Accessing Genesys Info Mart Manager" on page 340 provides instructions for accessing the GUI after it has been installed.
  - The generic term *management GUI* replaces previous references to the Genesys Info Mart Administration Console in contexts where the specific UI is not significant.

- The following updates relate to the new functionality described in "Support for Data Access Restrictions in GI2" on page 46:
  - The `cfg-annex` option has been added to the table of ICON configuration options that impact Genesys Info Mart, Table 9 on page 165.
  - Information about how the `GROUP_ANNEX` and `RESOURCE_ANNEX` tables are purged has been added to "Purging Info Mart Data" on page 405, and Table 27 on page 446, listing the `GROUP_ANNEX` and `RESOURCE_ANNEX` dimension tables, has been added to the appendix about Info Mart tables purged by the maintenance job.
  - In Appendix C on page 441, the `GC_ANNEX` table has been added to the list of IDB tables from which Genesys Info Mart extracts data.
- Descriptions of the following new Genesys Info Mart options have been added:
  - `days-to-keep-deleted-annex` on page 282
  - `ignored-reason-codes` on page 310

  The lists of data-related options in Table 14 on page 241 and operations-related options in Table 15 on page 248 have been updated to include the new options.
- The descriptions of the following options have been modified to reflect behavioral changes or to provide additional information in support of new features or enhancements:
  - `etl-start-date` on page 287
  - `extract-last-second` on page 290—The description has been modified to indicate that the option has been discontinued.
  - `max-parties-per-call` on page 293—The description has been modified to indicate that the option has been discontinued.
  - `max-time-deviation` on page 296—The minimum valid value has been changed from `0` to `1`.
- Updates related to the relaxed requirements for storage of Outbound Contact details (see "Outbound Contact Topologies" on page 48) include:
  - "Genesys Info Mart Requirements for ICON Details Storage" on page 70
  - The configuration requirements that Genesys Info Mart checks, as described in "General Deployment" on page 403
- In "Creating Genesys Info Mart Read-Only Tenant Views" on page 348, the instructions have been split into separate, RDBMS-specific procedures. The required modifications in the `make_gim_view_for_tenant.sql` script for PostgreSQL in release 8.1.4 and later have been updated (page 355).

  Related information about requirements for the Tenant User account for PostgreSQL has been updated in "Database Object Owners and User IDs" on page 93 and "Worksheet for PostgreSQL Databases" on page 426.

Documentation Enhancements and Corrections

- "Standby and Disaster Recovery" on page 72 has been revised to include an additional supported topology ("Active-Active Info Marts" on page 73).
- In the Notes on page 90, the requirement regarding general case insensitivity of the Info Mart database has been removed.
  - Replacement information about the case-insensitivity requirement for Microsoft SQL Server has been added in Task Summary: Preparing the Genesys Info Mart Database, on page 183.
  - Additional information about case sensitivity in the Oracle database has been added to Procedure: Tuning the Info Mart database on Oracle, on page 202.
  - Information about case-insensitive specification of schema names in PostgreSQL has been added in Task Summary: Preparing the Genesys Info Mart Database, on page 183 and in various places in Chapter 11 on page 207.
- The limitation for archive queues (page 101) now clarifies that the limitation applies to any interaction updates—for example, user-data updates after the interaction has been placed in an archive queue.
- The table of ICON configuration options that are important for Genesys Info Mart, Table 9 on page 165, has been updated to include `cseq-adjustment`.
- Table 11 on page 214 and Table 12 on page 220 are no longer divided into "JDBC and non-JDBC DAPs" and "Non-JDBC DAPs only" sections, since the `jdbc-*` options can be used in both JDBC and non-JDBC DAPs.
- The description of the `aggregate-schedule` option on page 315 has been enhanced.
- A warning against changing `-Duser` JVM parameters has been added to "Modifying JVM Startup Parameters" on page 360.
- In "Transaction Size in Transform" on page 388, the description of the chunk size for multimedia interactions has been corrected.
- On page 392, the recommendations for configuring chunk size have been enhanced.
- The notes on page 406 have been expanded to clarify that Genesys Info Mart does not purge GIDB configuration data.

New in Document Version v8.1.301.00

This document has been updated to support Genesys Info Mart release 8.1.3. The following topics have been added or changed since the previous release of this document:

Documentation Updates to Support New Features and Software Modifications

- A new subsection, "New in Release 8.1.3" on page 48, summarizes the feature enhancements that Genesys Info Mart 8.1.3 provides.
- Database-related information throughout the document has been updated as required to include PostgreSQL. In particular:
  - "RDBMS Considerations" on page 124
  - Procedure: Tuning the Info Mart database on PostgreSQL, on page 204
  - Recommended value for the `partitioning-interval-size-gim` option on page 298
  - "JDBC Driver for PostgreSQL" on page 326
  - Procedure: Creating Genesys Info Mart tenant views—Oracle, on page 351
  - "Additional Maintenance for PostgreSQL" on page 404 and "Scheduling the Maintenance Job" on page 409
  - "Worksheet for PostgreSQL Databases" on page 426
- Descriptions of the following new Genesys Info Mart options have been added:
  - `max-state-duration` on page 294
  - `canceled-queues` and `completed-queues` on page 309, with related information added in "Archive Queues" on page 100
  - `pipeline-timeout-in-hours` on page 312
  - `run-update-stats` on page 318
  - `update-stats-schedule` on page 319

  The lists of data-related options in Table 14 on page 241 and operations-related options in Table 15 on page 248 have been updated to include the new options.
- The descriptions of the following options have been modified to reflect behavioral changes or to provide additional information in support of new features or enhancements:
  - `delayed-data-threshold` on page 286. Related changes have been made in "Delayed Data for Extraction" on page 376, "Limitations" on page 377, "Delayed Data for Transformation" on page 389, and "Success Criteria for the Extraction Job" on page 404.
  - `max-session-duration-in-hours` on page 294—Changes have been made to the valid values and description.
  - `logging-level` on page 314—Valid values have changed.

Documentation Enhancements and Corrections

- The "New in Release 8.1.2" section has been updated to include the software enhancement that simplifies configuration management when a data source becomes unavailable (see page 53). Related updates include:
  - Additional criteria for data-source availability, as described in "Active Data Sources" on page 40
  - Additional steps in ICON and DAP deployment task summaries and procedures, to verify that the `Application` objects are enabled
  - Revisions to the description of the `delayed-data-threshold` configuration option on page 286, to replace the recommendation to temporarily disable unavailable data sources
  - In "Deployment Verification" on page 401, additional criteria for ICONs and DAPs to be considered part of the deployment and, in the subsection "Preventing Extraction Delays" on page 404, replacement of the recommendation to temporarily disable unavailable data sources
- Information about Java-related system requirements and host preparation in Chapters 2 and 15 has been updated to include the Server Java Runtime Environment (Server JRE).
- Table 1 on page 64 has been updated to include recommended minimum releases for:
  - DB Server for the Genesys Info Mart Administration Console in Microsoft SQL Server 2012 and PostgreSQL deployments.
  - Interaction Server and the Interaction Routing Designer (IRD) component of Routing, to support Genesys intelligent Workload Distribution (iWD).
- "Standby and Disaster Recovery" on page 72 has been revised to distinguish between *standby* and *disaster recovery* functionality.
- In the discussion of Interaction Concentrator topology requirements for Genesys Info Mart in Chapter 3, the requirement for each ICON to use only one database access point (DAP) to access its IDB has been spelled out (see "All ICON Details" on page 72).
- A note has been added on page 111 to clarify the format of skill combinations that Genesys Info Mart supports.
- In "Outbound Contact Solution" on page 113 and "Agent Desktop Applications" on page 113, misleading statements about how Genesys Info Mart uses the `GSW_CALL_ATTEMPT_GUID` KVP have been modified or removed.
- The valid values for the `max-backup-index` configuration option on page 314 have been corrected. Previously, the valid values incorrectly stated an upper limit of 99.
- The description of the `etl-end-time` configuration option on page 316 has been modified to provide additional guidelines for the ETL schedule and to remove an incorrect statement that setting `etl-start-time=etl-end-time` schedules the ETL cycle to run continuously.

- In "JDBC Driver for Microsoft SQL Server" on page 326, the following statement has been deleted as unnecessarily restrictive: "Although you can install the JDBC driver on a UNIX operating system, Genesys Info Mart supports accessing Microsoft SQL Server databases only from Windows operating systems."
- Since Genesys Info Mart does not use multimedia attached-data from the `GM_F_USERDATA` and `GM_L_USERDATA` tables in IDB in user-data processing, the sample ICON attached-data specification that is included in the Genesys Info Mart installation package (IP), `ccon_adata_spec_GIM_example.xml`, has been revised to remove misleading `<mcr-f>` and `<mcr-l>` elements. Appendix B, "Sample ICON Attached Data Specification," on page 437 has been revised accordingly.
- In Appendix C, `G_IS_LINK` and `GOX_CHAIN_CALL` have been added to the list of IDB tables from which Genesys Info Mart extracts (Table 25 on page 442), and a note has been added to clarify that Genesys Info Mart does not populate all of the GIDB tables in the Info Mart database schema.
- Additional minor corrections and clarifications have been made throughout the document.

New in Document Version v8.1.202.00

This document has been re-issued to support Genesys Info Mart release 8.1.2. The following topics have been added or changed since the previous release of this document:

- In "New in Release 8.1.2" on page 50, a new "Enhanced Support for User Data" item describes the enhancement that was introduced in release 8.1.201 to enable storage of custom user-data facts as date/time data types. Related updates include:
  - Modifications to statements in Chapters 6 and 10 and in "Worksheet for Mapping User Data" on page 428 about supported data types
  - Additional information about user-data processing (Step 4 on page 106)
  - Additional information about user-data configuration (Step 3 and Step 12 in Procedure: Customizing the user-data template script, on page 185)
- The description of the `user-event-data-timeout` configuration option (page 301) now correctly indicates that the valid values are `0` or any positive integer. (The previous 8.1.1 versions of this document incorrectly indicated that the maximum valid value had been changed to `3600`.)

New in Document Version v8.1.201.00

This document has been updated to support Genesys Info Mart release 8.1.2. The following topics have been added or changed since the previous release of this document:

Documentation Updates to Support New Features and Software Modifications

- A new subsection, "New in Release 8.1.2" on page 50, summarizes the feature enhancements that Genesys Info Mart 8.1.2 provides.
- Minor updates in database-related material throughout the document indicate that partitioning support has been expanded to include certain Control tables.
- Information about support for partitioned IDBs has been added to "Genesys Info Mart Requirements for ICON Details Storage" on page 70 and "Source Data Retention and Purging" on page 92.
- "Standby and Disaster Recovery" on page 72 has been modified to include support for an Info Mart database for disaster recovery, by means of database replication with Oracle GoldenGate.
- "Genesys Info Mart and Attached Data" on page 103 has been slightly revised to match the structure of the User Data Assistant utility. Information about using the tool has also been included in "Enabling Storage of User Data" on page 148, "Customizing Your ICON Attached Data Specification File" on page 176, and "Customizing the User Data Template" on page 184.
- Information about support for user data in `MEDIATION_SEGMENT_FACT` (MSF) records has been included in Step 6 on page 106 in the section about processing user data; "Limitations for User Data in MSFs" on page 107; "Propagation Rules" on page 118; Task Summary: Enabling the Storage of User Data, on page 149; and Figure 8 on page 185. For links to additional related information, see the Document Change History item about `link-msf-userdata` on page 22.
- In various task summaries, steps have been modified to include new functionality about automatic upgrade of IDBs when required.
- The table of ICON-related `Switch` options that are important for Genesys Info Mart, Table 17 on page 256, has been updated to include `third-party-queue-in-divert`.
- In Procedure: Creating the Genesys Info Mart application, on page 235, a note has been added about configuring an optional connection to Configuration Server Proxy, instead of to Configuration Server, in deployments with distributed data centers.
- Descriptions of the following new Genesys Info Mart options have been added:
  - On the Genesys Info Mart `Application` object:
    - `etl-start-date` on page 287
    - `max-chain-processing-duration-in-hours` on page 292

    - max-parties-per-call on page 293
  - On DAP objects:
    - geo-location on pages 215 and 221. Related information has been included in the procedures to configure extraction and Info Mart DAPs that start on page 212; the description of the max-call-duration option on page 290; "Extracting Data in an HA Deployment" on page 381 (including a new subsection, "HA Data Extraction with Geolocation" on page 385); "Criteria for Best IDB" on page 413.
  - On DN and Script objects:
    - link-msf-userdata on pages 263 and 266, with related information in Step 7 on page 263 and Step 5 on page 266 in the respective procedures to configure the objects, and in Table 14 on page 241.
- The descriptions of the following options have been modified to reflect behavioral changes or to provide additional information in support of new features or enhancements:
  - days-to-keep-* options, starting on page 281—Dependencies have been introduced, and the timestamps from which purge thresholds are calculated have changed. Related changes have been made in "Purging Info Mart Data" on page 405 and Appendix D on page 445.
  - max-call-duration on page 290—The description includes the role of the option in voice deployments in which geolocation is a consideration. In addition, the range of valid values has been corrected to include 0 (zero).
  - max-thread-duration-after-inactive-in-days on page 295—Changes have been made to the default value, valid values, and dependencies.
  - partitioning-ahead-range on page 297—The starting point of the partitioning ahead range has changed.
- The lists of data-related options in Table 14 on page 241 and operations-related options in Table 15 on page 248 have been updated to include new options and changed or corrected default values.
- Procedure: Creating Genesys Info Mart tenant views—Oracle, on page 351 has been updated.
- The description of the "Configuration Check" on page 402 has been updated to include new checks for:
  - Purge option dependencies
  - The ICON role option

Documentation Enhancements and Corrections

- A new item, "Database, Database Schema, and Database Instance" on page 39, has been added to the terminology conventions.
- In "Software and Database Requirements" on page 62 and "Preparing the Genesys Info Mart Server Host" on page 323, references to Java 1.6 Java Development Kit (JDK) have been made generic and no longer refer to a

specific version. Consult the *Genesys Supported Operating Environment Reference Guide* for the Java versions that are supported in a particular Genesys Info Mart release.

- Information about not configuring auto-restart for the Genesys Info Mart application has been added to "Standby and Disaster Recovery" on page 72 and the step to configure "Start Info Tab" on page 237.
- Notes about case-insensitivity and database compression have been added on page 90.
- "Database Object Owners and User IDs" on page 93 has been revised to clarify the types of users for which you must provision user accounts, as well as required permissions. References to "owner IDs" or "owner accounts" have been removed from various database-related procedures.
- Information about multimedia-specific user data has been clarified in "eServices/Multimedia-Specific Attached Data" on page 112 and "Mapping Call-Based Attached Key-Value Pairs" on page 117.
- A new subsection, "Low-Cardinality vs. High-Cardinality User Data" on page 115, provides guidelines for deciding whether to store user data as facts or dimensions.
- The table of ICON configuration options that are important for Genesys Info Mart, Table 9 on page 165, has been updated to include `extended-route-result`.
- In the description of the ICON `observer-party` option on page 175, a note recommends setting this filter in Genesys Info Mart 8.1.1 and 8.1.0.
- A note has been added on page 194, recommending against the use of database links in Oracle RAC deployments.
- In Procedure: Tuning the Info Mart database on Microsoft SQL Server, on page 200, a new step (Step 5) provides information about configuring the `READ COMMITTED` isolation level.
- In Procedure: Tuning the Info Mart database on Oracle, on page 202:
  - The recommended minimum settings for the `processes`, `sessions`, and `open_cursors` initialization parameters have been increased to `1000`.
  - Information about changing the queue size for listener has been added.
- In the chapter about configuring DAPs, the descriptions of the `jdbc-url` option (for example, on page 216) have been enhanced to include examples for enabling Secure Socket Layer (SSL) over the connection and for using IPv6.
- The descriptions of the following Genesys Info Mart options have been corrected or enhanced:
  - `date-time-start-year` on page 273—A recommendation has been added about setting the start date.
  - `extract-data-cfg-facts-chunk-size` on page 287—The description clarifies that a value of `0` (zero) is considered to be the same as a negative value.

  - max-chunks-per-job on page 292—The range of valid values includes 100.
  - sm-resource-state-priority on page 301—The description clarifies that all four resource states must be specified in the option value.
  - user-event-data-timeout on page 301—The range of valid values includes 0 (zero), and the stated default value has been restored to 3600 seconds. (The previous 8.1.1 versions of this document incorrectly indicated that the default value had been changed to 300.) In addition, the description has been enhanced to note the role of the option in managing after-call work (ACW) dependencies between agent and interaction data.
- The diagrams illustrating HA data extraction for Voice details (Figure 23 on page 384 and Figure 24 on page 385) have been revised for accuracy and consistency.
- A note has been added on page 411 to clarify that support for HA means support for sets of redundant Interaction Concentrator instances, not HA of Genesys Info Mart itself.

New in Document Version v8.1.102.00

This document has been re-issued to support Genesys Info Mart release 8.1.1. The following topics have been added or changed since the previous release of this document:

- In "New in Release 8.1.1" on page 54, configuration requirements to support the "Enhanced Information in MSFs and IRFs" feature have been clarified.
- In the section about "Processing User Data" on page 105, an incorrect reference to the nonexistent STG_TRANSFORM_VIOLATIONS table has been corrected to STG_TRANSFORM_DISCARDS.

New in Document Version v8.1.101.00

This document has been updated to support Genesys Info Mart release 8.1.1. The following topics have been added or changed since the previous release of this document:

Documentation Updates to Support New Features and Software Modifications

- A new subsection, "New in Release 8.1.1" on page 54, summarizes the feature enhancements that Genesys Info Mart 8.1.1 provides.
- A new subsection, "Workbin Instance and Personal Workbin" on page 40, provides information about concepts that are related to the new Workbin dimension.
- "Security Features" on page 56 provides more information about the new security features. There are related additions to the following topics:
  - System requirements—"Security Requirements" on page 63 and Table 1 on page 64

  - Additional steps in various task tables and procedures, as well as a new task summary, "Enabling Secure Connection Features" on page 154
  - Additional supported startup parameters—page 359
- Information about compatibility with other Genesys software components (Table 1 on page 64) has been updated as required for the support of new features.
- Statements about Genesys Info Mart support for *disaster recovery* have been changed to refer to a *standby instance* only (for example, in "Standby and Disaster Recovery" on page 72), to avoid misleading customers about the extent of Genesys Info Mart 8.x functionality in this area.
- The new user-data propagation rule, `IRF_FIRST_UPDATE`, is described in "Propagation Rules" on page 118, and new scenarios have been added to Table 5 on page 119 to illustrate the effect of the new propagation rule.
- In Chapter 9 on page 159, the following ICON options have been added:
  - `cfg-auto-resync` on page 166
  - `route-res-vqid-hist-enabled` on page 169
  - `call-history` on page 174
- Descriptions of the following new Genesys Info Mart options have been added:
  - `fiscal-year-start` and `fiscal-year-week-pattern` on page 274, with related information about using the options to create a fiscal calendar included in the Genesys Info Mart configuration procedure, Step 6 on page 240
  - `error-policy-call-mergecall-missing` on page 276 and `error-policy-irf-exception-resumable` on page 277, with related information included in "Data Inconsistencies During Transformation" on page 397
  - `max-thread-duration-after-inactive-in-days` on page 295
  - `merge-failed-is-link-timeout` on page 297
  - `adjust-vq-time-by-strategy-time` on page 309
  - `ud-io-parallelism` on page 312
- The descriptions of the following options have been modified to reflect behavioral changes or to provide additional information in support of new features or enhancements:
  - `date-time-start-year` and `date-time-tz` in the `[date-time]` section starting on page 271.
  - `days-to-keep-active-facts` (including a change to the default value), `days-to-keep-gidb-facts`, and `days-to-keep-gim-facts`, starting on page 281, with related information updated in "Purging Info Mart Data" on page 405. For release 8.1.0 customers, a new appendix, Appendix E on page 451, provides information about purge functionality in previous 8.x releases.
  - `max-time-deviation` on page 296.

  - partitioning-interval-size-gim on page 298 (change to the default value only).
  - user-event-data-timeout on page 301 (including changes to the default value and valid values).
  - populate-mm-*-facts options starting on page 304.
  - populate-sm-resource-*-facts options starting on page 306.
  - populate-workbin-as-hold on page 307.
- The lists of data-related options in Table 14 on page 241 and operations-related options in Table 15 on page 248 have been updated to include new options or changed default values, as applicable. In addition, some of the options that were formerly listed under the "Performance tuning" section of Table 15 on page 248 have been moved into a new section, "Miscellaneous" on page 250.
- Information about Genesys Info Mart extraction behavior and related configuration requirements have been modified throughout the document. Significant changes are:
  - The definition of *active data sources* on page 40 and related changes throughout Chapter 18 on page 371 and in "Data-Source Availability" on page 403
  - Requiring the use of Advanced Disconnect Detection Protocol (ADDP) for ICON connections, as described in the task summaries in "Enabling Specific Functionality" on page 140 and Step 3 on page 164
  - Deletion of the section "Determining HA Mode" from Chapter 20 on page 411

Documentation Enhancements and Corrections

- To improve usability and to reduce duplication in the Genesys Info Mart documentation suite, some information that formerly appeared in the *Genesys Info Mart 8.1 User's Guide* has been incorporated into Chapter 18 on page 371 and Chapter 20 on page 411.
- Table 3 on page 95 consolidates and revises information that previously was provided in separate tables that listed required owner- and user-account database privileges. The names that describe the view-related database user accounts have been modified, to clarify the roles: *Info Mart Views* has been changed to *Tenant Admin,* and *Info Mart Tenant Views* has been changed to *Tenant User.* Related changes have been made in the section "Genesys Info Mart Database" on page 37; in Figure 20 on page 350; in the procedure to create the tenant views (page 351); and in the database worksheets starting on page 422.

- Genesys Info Mart support for the KVPs that Genesys Interactive Insights (GI2) uses for social media metrics is mentioned in "Special Requirements" on page 116 and reflected in the sample attached data specification file for ICON, which is reproduced in Appendix B on page 437.
- In Table 7 on page 122, a note has been added to the descriptions of the `Revenue` and `Satisfaction` KVPs.
- Because Genesys Info Mart does not require the `send_attribute` to be set in Outbound Contact–related `Field` objects, misleading information about configuring this option has been removed from "Mandatory Record Field Data" on page 128; Procedure: Configuring the mapping of OCS record fields, on page 268; and Task Summary: Enabling the Storage of Outbound Contact Details, on page 147. As a related change, the `Attached Data Key` column has been deleted from "Worksheet for Mapping OCS Record Fields" on page 432.
- There have been minor clarifications and additions in Task Summary: Enabling High Availability in a New or Existing Deployment, on page 152 and "Configuring for HA" on page 416. As a related change, a note that previously preceded the task summary on page 152 has been deleted, because it provided redundant information, which was confusing in the context.
- An item has been added to the Notes on page 163 about requirements for IDB data population.
- Information about using database links has been enhanced:
  - New subsections, "Creating Linked Servers—Microsoft SQL Server" on page 197 and "Verifying Use of Database Links" on page 200, have been added.
  - The SQL statement for Oracle that is provided in "Verifying Connection Information" on page 196 has been corrected, to specify `ALL_DB_LINKS` instead of `DB_LINKS`.
- To clarify configuration of custom user-data storage:
  - In the section about "Customizing the User Data Template" on page 184, a new diagram provides an overview of the user-data subject area in the Info Mart database. See Figure 8 on page 185.
  - The steps in Procedure: Customizing the user-data template script, on page 185 have been re-ordered and revised to match revisions that were made in the script.
- In the "schedule Section" on page 314, information about the aggregation-scheduling configuration options (`aggregate-duration`, `aggregate-schedule`, and `run-aggregates`) and the maintenance-scheduling configuration options (`maintain-start-time` and `run-maintain`) has been corrected so as not to show a dependency on the `run-scheduler` option.
- In Appendix B on page 437, a note about specifying a value of `all` for the `history` attribute for custom user data has been modified.

Genesys

Part

1 Overview and Planning

Part 1 of this document provides an overview of Genesys Info Mart functionality, as well as information about the implementation of Genesys Info Mart features that you must consider when you plan your deployment. This information appears in the following chapters:

Chapter

1 Genesys Info Mart Deployment Overview

This chapter describes the basic Genesys Info Mart architecture, the main Genesys Info Mart components and their functions, Genesys Info Mart 8.x features, and functionality that is new in release 8.1. It also includes a discussion of Genesys Info Mart database schemas.

This chapter contains the following sections:

Architecture

Genesys Info Mart 8.x extracts data from one or more Genesys Interaction Concentrator databases (Interaction Databases [IDBs]) and produces a data mart for contact center historical reporting.

Genesys Info Mart consists of a server component that extracts, transforms, and loads data into a data mart, based on a schedule that is configured in the Genesys Info Mart application. The Genesys Info Mart Administration Console or, starting with release 8.1.4, the Genesys Info Mart Manager provides a graphical user interface (GUI) for managing some of the Genesys Info Mart jobs. The Info Mart database stores low-level interaction data that is consolidated from any number of Interaction Databases (IDBs), as well as processed data that is suitable for end-user reports.

Figure 1 illustrates the Genesys Info Mart 8.x architecture and the primary data flow between the Genesys Info Mart components and other Genesys

components. (The diagram does not depict high-availability architecture for any components.)

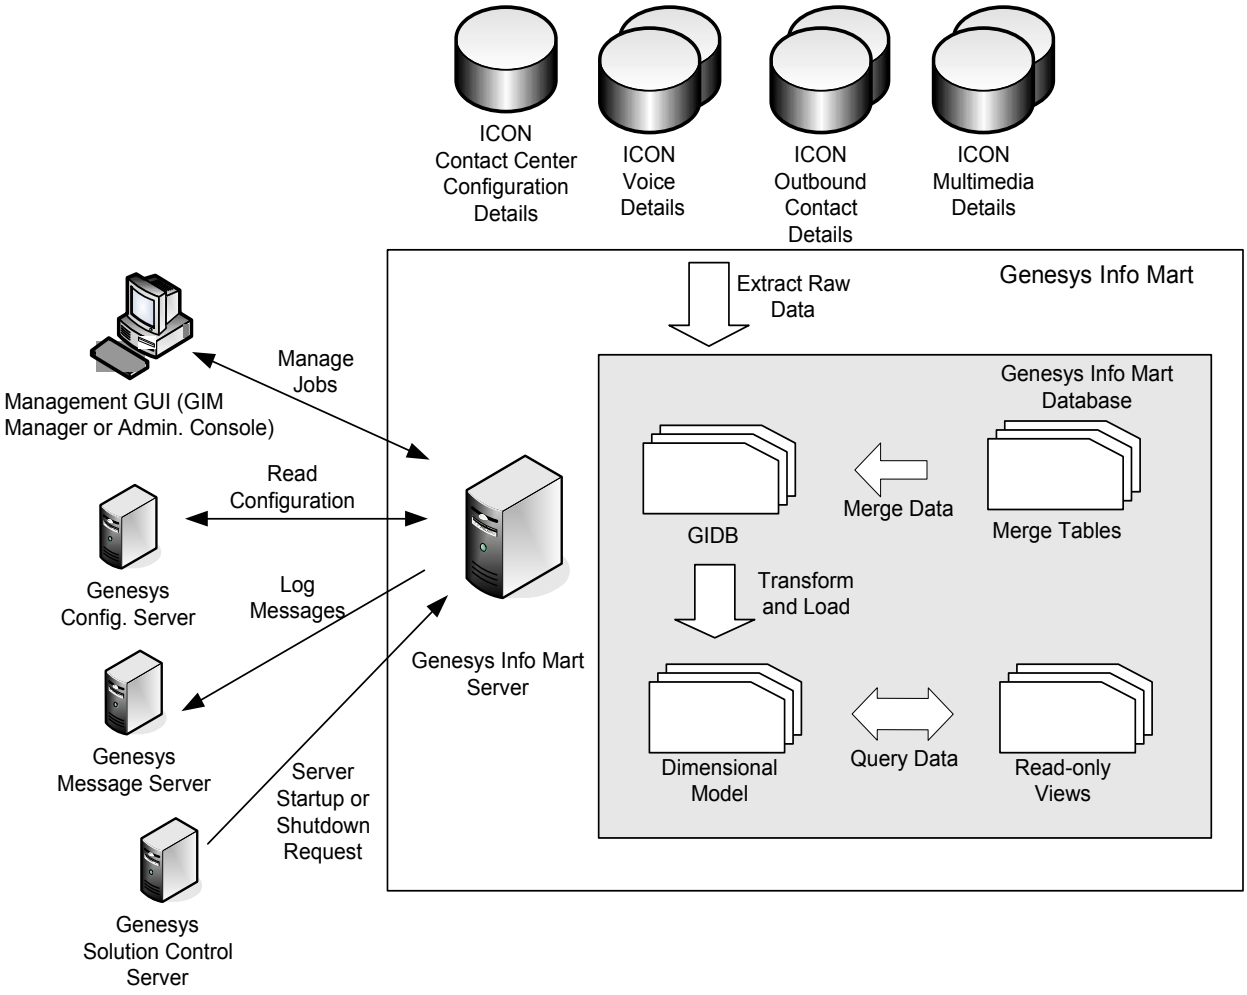

Figure 1: Genesys Info Mart Architecture and Data Flow Diagram

Types of Reporting Data

Depending on how it is configured, Genesys Info Mart stores the following types of data, which it extracts from one or more IDBs, which are populated by one or more Interaction Concentrator (ICON) applications:

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

  The ICON data source for Configuration details is Configuration Server.
- ICON Voice details, which include:
  - Voice interaction data

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

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

Note: In this document, the term *T-Server* is used generically to refer to all T-Server types (premise and network TDM Voice, SIP Server, IVR Server, and Virtual T-Server).

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

  The ICON data source for Multimedia details is Interaction Server.
- Outbound Contact details, which include:
  - History and results of campaigns, chains, and contact attempts
  - Associations between Outbound Contact objects (such as campaigns) and contact center objects (such as agent groups or place groups)
  - Precalculated Outbound Contact metrics

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

Default Genesys Info Mart Configuration

Configuration options enable you to customize almost every aspect of Genesys Info Mart functioning. However, the default configuration settings provide a Genesys Info Mart application that is suitable for a wide range of standard deployments.

The remaining sections in this chapter provide a high-level overview of Genesys Info Mart functioning. For advanced users who want a deeper understanding, Part 4 on page 369 provides more detailed information about how Genesys Info Mart works. If you decide that the default Genesys Info Mart application does not meet your requirements, Genesys recommends that you review the information in Part 4 during your deployment planning, to enable informed decision-making about application modifications.

Components and Functions

Genesys Info Mart consists of the following components:

- Genesys Info Mart Server
- Management GUI—Genesys Info Mart Administration Console or, starting with release 8.1.4, Genesys Info Mart Manager (see page 36)
- Info Mart database, which is created with Genesys-provided scripts (see page 37)

These components are described in the following subsections.

Genesys Info Mart Server

The Genesys Info Mart Server, a Java-based component, is the main executable process in Genesys Info Mart 8.x. Its main function is to run various functional jobs, including but not limited to ETL jobs. These jobs run according to the schedule that is configured in the Genesys Info Mart ETL application in Genesys Configuration Layer. The Genesys Info Mart Server also processes requests from applications outside of the Genesys Info Mart Server, such as requests from the management GUI to launch a specific job.

The Genesys Info Mart Server interfaces with:

- Solution Control Server (through Local Control Agent [LCA]), to control when the Genesys Info Mart Server starts and stops.
- Configuration Server, to read Genesys Info Mart application configuration options, as well as other configuration objects and options that affect Genesys Info Mart functionality.
- The management GUI (Genesys Info Mart Manager or Genesys Info Mart Administration Console), to start and stop jobs and to provide the status of ETL jobs.
- Message Server, to log messages to the Centralized Log Database.
- The log4j Java client, to log messages to the local log.

Genesys Info Mart Jobs

Genesys Info Mart jobs, which run under the Genesys Info Mart Server, do the following:

Initialize

- Set up your database and, starting with release 8.1.2, modify IDB for Info Mart use.

Extract

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

- Extract voice interaction, user data (including call-based attached data and UserEvent-based key-value pair (KVP) data), virtual queue, and agent activity details from one or more IDBs.
- Extract Outbound Contact details from one or more IDBs.
- Extract multimedia interaction, attached data, virtual queue, and agent activity details from one or more IDBs. For more information about the types of interactions that Genesys Info Mart considers to be *multimedia,* see "Multimedia Interactions" on page 40.
- In an HA configuration, evaluate all redundant IDBs by comparing ICON-provided session information for each set of redundant IDBs (that store Configuration, Voice, Outbound Contact, or Multimedia details), prior to extracting the data in a particular extraction cycle.

Merge

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

Transform

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

Maintain

- Purge old data from the Info Mart database.
- In partitioned databases, update partitioning as necessary.
- Maintain calendar data by prepopulating calendar dimensions for use in reports.

Migrate

- Run the necessary scripts to update your Info Mart database schema and, starting with release 8.1.2, IDB schema and also perform any other steps that are necessary to move from an earlier 8.x release to the current one.

Aggregate

In deployments that include Genesys Interactive Insights (GI2) or the separately installed Reporting and Analytics Aggregates (RAA) package, Genesys Info Mart also hosts the aggregation process. An aggregation job, which is implemented as a plug-in, runs the aggregation engine inside the Genesys Info Mart Server process, to calculate or recalculate the historical Aggregate tables in the Info Mart database. For a high-level description of how the aggregation job works in Genesys Info Mart release 8.x, see the description of `Job_AggregateGIM` in the chapter about understanding Genesys Info Mart jobs in the *Genesys Info Mart 8.1 Operations Guide*. For more information about the aggregation package and running the aggregation process, see the *Reporting and Analytics Aggregates 8.x Deployment Guide.*

Management GUI

Genesys Info Mart provides two alternative GUIs that enable monitoring and real-time administration of some aspects of the Genesys Info Mart processes: Genesys Info Mart Manager, which was introduced in Genesys Info Mart release 8.1.4, and the Genesys Info Mart Administration Console, which has been available since the first release of Genesys Info Mart 8.0. Both are included on the Genesys Info Mart CD as separate installation packages (IPs).

Functionality

The management GUIs provide the following functionality:

- Display the current job execution status.
- Display a history of job execution, including start time, stop time, duration, and final status.
- Filter the job execution history that is displayed, based on time and/or status.
- Execute a single job on an ad hoc basis.
- (Administration Console only) Schedule a single job on an ad hoc basis to run at a specified future time and date. Also, if necessary, cancel a job scheduled in this way.
- (Starting with release 8.1.4) Issue re-aggregation requests in deployments that include GI2 or the separately installed RAA package.
- Shut down a running job.

Interfaces

Genesys Info Mart Manager interfaces directly with the Genesys Info Mart Server to start and stop jobs on an ad hoc basis and to obtain job status and history from the Info Mart database.

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

- The Genesys Info Mart Server, to start, schedule, and stop jobs on an ad hoc basis.
- The Info Mart database, to query job status, job history, and job schedules from the Info Mart database. The Genesys Info Mart Administration Console connects to the Info Mart database through a DB Server and a database access point (DAP).

Relative Advantages of Genesys Info Mart Manager

Genesys Info Mart Manager has the following advantages over the Genesys Info Mart Administration Console:

- Genesys Info Mart Manager works with any operating system. The Administration Console is available for Windows only.
- Genesys Info Mart Manager can be localized. For links to information about installing language packs, see the Next Steps on page 332.
- Genesys Info Mart Manager is easier to deploy. You do not need to provide a DB Server or database access point (DAP) for Genesys Info Mart Manager to connect to the Info Mart database, and no installation or configuration is required on client hosts.

Genesys recommends that you use Genesys Info Mart Manager instead of the Administration Console in Genesys Info Mart 8.1.4 and later deployments, particularly if your deployment already includes Genesys Administrator Extension (GAX).

More Information

For more information about how to use the management GUIs to manage your Genesys Info Mart operations, see the chapter about working with jobs in the *Genesys Info Mart 8.1 Operations Guide*.

Note: Genesys Info Mart Manager displays timestamps in the time zone selected in the GAX preferences for locale.

All timestamps in the Genesys Info Mart Administration Console are in the local time zone of the host on which the Genesys Info Mart Administration Console is installed and running.

Info Mart Database

The Info Mart database contains all of the data that is populated by the ETL jobs, including the low-level interaction data that is consolidated from one or more IDBs as well as the processed data that is suitable for end-user reports.

Genesys Info Mart Database

Info Mart Tables

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

- Control tables—Store information that controls ETL execution, such as the status of running jobs, job schedules, execution history, audit logs, and similar bookkeeping information.
- Merge tables—Used for merge of voice interactions. After the merge, voice interactions are moved into Global Interaction Database (GIDB) tables.
- GIDB tables—Store the low-level interaction data that is consolidated from one or more IDBs.

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

  The low-level reporting data in GIDB supports the possibility of custom detailed reporting to drill down from the dimensional model.
- Staging tables—Store information that is required in the transformation process. Staging tables contain data that has not yet been transformed or cannot be transformed because of incompleteness or inconsistency of source data and other auxiliary data that are necessary for the ETL process.
- Temporary tables—Store data that is used only during the lifetime of one instance of a job.

- Fact and dimension tables (collectively referred to as *dimensional model)*—Contain transformed data that downstream reporting applications can query and combine in meaningful reports.

In deployments that use GI2 or RAA, the Info Mart database also includes aggregate tables and views that are used by GI2. For more information, see the *Reporting and Analytics Aggregates 8.x Reference Manual.*

Info Mart Views

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

- Predefined views—Read-only views of certain configuration dimensions and facts that are contained in GIDB tables. These views are created in the Info Mart database schema.
- Tenant-specific, read-only views of the dimensions, facts, and predefined views in the dimensional model of the Info Mart schema. These views are created in a Genesys Info Mart schema.
- Tenant-specific sets of views, which are created in a separate database schema for each tenant, so that each tenant user can access only its own data. In addition, these views shield business users from changes to the underlying database schema and prevent users from accidentally changing the contents of the underlying database. Customers should use these views to query Genesys Info Mart data.

  For an illustration of the relationship between the tenant views and the Info Mart database schema, see Figure 20 on page 350.

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

For more information about the database schemas, see the Genesys Info Mart overview section in the *Reference Manual* for your relational database management system (RDBMS):

- *Genesys Info Mart 8.1 Microsoft SQL Server Reference Manual*
- *Genesys Info Mart 8.1 Oracle Reference Manual*

RDBMS-Specific SQL Scripts

A set of SQL scripts is provided for each RDBMS type that is supported. The scripts perform a variety of tasks, including the following:

- Create the Info Mart database objects.
- Create the single-tenant and multi-tenant read-only views.
- Update source databases for efficient data extraction.

Terminology Conventions

This subsection describes the usage of terms that have specific meanings in the Genesys Info Mart documentation.

Database, Database Schema, and Database Instance

The word *database* has different meanings in the Genesys Info Mart documentation, depending on the context. It may refer to Genesys components, such as "Interaction Database" or "Info Mart database." It may also be used in reference to general RDBMS concepts and procedures, such as "database export" or "database replication." Where it is significant to refer to a particular organization of tables, views, indexes, and other database objects, the term *database schema* is used. Where it is significant to refer to the RDBMS that manages database files, the term *database instance* is used. This terminology might not necessarily match the terminology that is used by leading RDBMS vendors.

Database Area

As described in the section "Info Mart Database" on page 37, the Info Mart database consists of only one schema, which comprises several groupings of tables (GIDB tables, Merge tables, and so on). In this document, the groupings of tables might be referred to as a *database area*—for example, the Merge area.

Data Domains

The scope of Genesys Info Mart activity, in terms of the type of details that it processes, is defined by the configured role of the DAP(s) through which Genesys Info Mart accesses IDB. The *data domains* correspond to the type of details that each IDB stores—Configuration details, Voice details, Multimedia details, or Outbound Contact details. Genesys Info Mart processes each data domain separately.

For more information about the types of details, see "Types of Reporting Data" on page 32.

Data Source

The immediate source of data for Genesys Info Mart is IDB, which is populated by ICON. The source of data for ICON is Configuration Server, T-Server, Interaction Server, or OCS, depending on the configured role of the ICON application.

In this guide, the term *data source* refers to the upstream data provider—the source of data for ICON.

Available Data Sources

Starting with release 8.1.1, Genesys Info Mart extracts data from all the DAPs in its connections—that is, from all the data sources that populate the IDBs

from which Genesys Info Mart is configured to extract data. Starting with release 8.1.2, the ICONs and the extraction DAPs must be enabled in order for Genesys Info Mart to consider them and the associated data sources and IDBs to be part of the deployment. (*Enabled* means that the `State Enabled` check box on the `General` tab of the `Application` objects is selected.)

Active Data Sources

While Genesys Info Mart will extract data from all available data sources, Genesys Info Mart will wait for delayed data only from *active data sources*. In a Genesys Info Mart deployment, active data sources are data sources that are:

- Currently monitored by enabled ICONs that are connected to Genesys Info Mart
- Enabled—the `State Enabled` check box on the `General` tab of the T-Server, Interaction Server, or Outbound Contact Server `Application` object is selected

Voice and Multimedia Interactions

Genesys Info Mart supports reporting on both voice and multimedia interactions.

Voice Interactions

The term *voice interactions* refers to traditional telephony calls.

Multimedia Interactions

The term *multimedia interactions* refers collectively to all interactions that are processed through Genesys eServices/Multimedia solution, such as:

- *eServices/Multimedia interactions.* E-mail and chat are two of the Genesys-provided media types that Genesys Info Mart currently supports.
- *3rd Party Media interactions* (formerly referred to as *Open Media*). These are interactions of any custom media channel that is supported on top of Genesys eServices/Multimedia. The `Workitem` media type is an example of 3rd Party Media.

Genesys Info Mart processes data that is related to all multimedia interactions in a similar manner.

Workbin Instance and Personal Workbin

Workbin Instance

A workbin can be used to hold interactions for resources of a given type: `Agent`, `Place`, `AgentGroup`, or `PlaceGroup`. The `Script` object of type `Interaction Work Bin` in the Configuration Layer indicates the type of resource.

In the Genesys Info Mart documentation suite, the term *workbin instance* does not simply refer to an `Interaction Work Bin` object, but also to the resource that is indicated as the owner of the interaction in the workbin.

For example, if an `Interaction Work Bin` object that is named Drafts has been defined in the Configuration Layer for use by `Agent` resources, the expression "Agent1's Drafts workbin" refers to a workbin instance that represents the use of the Drafts workbin for interactions that are assigned to Agent1.

Personal Workbin

As a special case for workbins of type `Agent` or `Place`, *personal workbin* refers to the situation in which `Agent` or `Place` resources place interactions in their own workbin instances. The concept of a personal workbin does not apply to `AgentGroup` and `PlaceGroup` workbins. To extend the previous example, Agent1 placing an interaction into the Drafts workbin—with Agent1 specified as the owner of this workbin interaction—is an example of a personal workbin.

Genesys Info Mart 8.x Features and Functionality

Genesys Info Mart 8.x architecture and implementation is different from Genesys Info Mart 7.6. However, Genesys Info Mart 8.x provides at least the same scope and granularity of data for reporting that was provided in Genesys Info Mart 7.6. This section highlights additional or changed functionality that was introduced in Genesys Info Mart 8.0. For more information about changes between Genesys Info Mart 7.6 and 8.0, see the appendix about functional changes in the *Genesys Info Mart 8.0 Deployment Guide*. See also the appendix about database schema changes in the *Genesys Info Mart 8.0 Reference Manual* for your RDBMS.

Depending on configuration, Genesys Info Mart 8.x provides the following features and functionality. See "New in This Release" on page 45 for additional, new functionality introduced in Genesys Info Mart 8.1.

Multimedia Data Processing

- Genesys Info Mart 8.x enhances support for extracting and transforming all interactions that are processed through Genesys eServices/Multimedia solution, including 3rd Party Media interactions. In particular, Genesys Info Mart 8.x:
  - Populates interaction resource facts (IRFs) and interaction resource state facts (IRSFs) for Genesys eServices/Multimedia interactions.
  - Supports reporting on Interaction Queue and Interaction Workbin activity in addition to previously supported mediation DNs.
  - Supports dynamically adding new media types that are encountered during transformation of multimedia interactions. For more information, including configuration considerations, see "Media Types" on page 97.
  - Supports dynamically adding new multimedia interaction subtypes that are encountered during transformation. For more information, including configuration considerations, see "Interaction Types and Subtypes" on page 99.

Configured Thresholds by Media Type

- Genesys Info Mart 8.x supports configuration of answer and abandon thresholds separately for each media type. You can configure the thresholds:
  - Per media type
  - Per tenant per media type

- Per tenant per media type per DN

Configuration sections and options in the Genesys Info Mart `Application` and in supporting objects enable you to customize media-specific thresholds, as required. For more information, see "gim-etl-media-<media type> Section" on page 302 and Table 14 on page 241.

Outbound Contact Data Processing

- Genesys Info Mart 8.x processes Outbound Contact data independently from interaction data.
- Support for Outbound Contact reporting includes support for agent group information, through the `RESOURCE_GROUP_COMBINATION_KEY` field in the `CONTACT_ATTEMPT_FACT` table. For details, see the description of Outbound Contact dimensions in the *Genesys Info Mart 8.1 User's Guide*.

User Data Handling

- Genesys Info Mart 8.x provides a unified mechanism for processing user data from both `EventUserEvents` and call-based `TEvents`. You can configure flexible data storage according to the number and types of user data that are captured in your contact center environment. A customizable database schema enables you to treat each key-value pair (KVP) field as either a fact or a dimension, and to store user-data KVPs in a configurable number of user-data dimensions and fact extension tables. Genesys Info Mart 8.x also processes the user data that arrives after call completion, and updates call records accordingly.

  For more information about user-data handling in Genesys Info Mart 8.x, see Chapter 6 on page 103.
- Genesys Info Mart provides predefined support for storage of KVPs for Revenue and Satisfaction. For descriptions of these and other KVPs, see Table 7 on page 122.

Aggregation

- Genesys Info Mart 8.x default functionality does not include aggregation. Instead, Genesys provides aggregation software that is bundled with the Genesys Interactive Insights (GI2) product. The aggregation software is also available as a separate installation package (IP), called Reporting and Analytics Aggregation (RAA), which is delivered with Genesys Info Mart 8.x. The aggregation layer is installed on top of Genesys Info Mart, which hosts the aggregation process.

High Availability

- Genesys Info Mart 8.x supports high availability (HA) of Configuration, Voice, Multimedia, and Outbound Contact details.
- Unlike Genesys Info Mart 7.6, Genesys Info Mart 8.x does not require special configuration to identify primary and secondary IDB data sources. Instead, Genesys Info Mart relies on session-control information that is provided by redundant Interaction Concentrator 8.x applications in IDBs. The extraction job evaluates all redundant IDBs and extracts data from the IDB that is the best source for Configuration, Voice, Multimedia, or Outbound Contact details for a particular timespan.

  For more information about HA in Genesys Info Mart 8.x, see Chapter 20 on page 411.

Multiple Calendars

- Genesys Info Mart 8.x supports multiple, customizable calendars with flexible week-numbering rules that can be configured to conform to the ISO 8601 standard for the representation of dates and times.
- Genesys Info Mart stores time facts in Coordinated Universal Time (UTC) time. Scalable support for multiple calendars means that Genesys Info Mart can be configured to express time data in any Java time zone format.
- Genesys Info Mart 8.x provides one default calendar (`DATE_TIME` dimension). The default configuration expresses UTC time in the GMT time zone and conforms to legacy Genesys Info Mart week-numbering rules, which are not the ISO 8601 standard.

Data Lineage

- Genesys Info Mart 8.x provides the capability to store the processing history of jobs and the extraction and transformation history of each piece of data.

  There are two aspects of data lineage:

  - Voice of Data—Special fields store service data that enables you to trace a particular reporting data item to its source system, as well as to trace data in the opposite direction (from source to target). Information that is stored as Voice of Data enables data tracking for the purpose of validating data and troubleshooting data-quality issues. For more information, see the *Genesys Info Mart 8.1 User's Guide.*
  - Voice of Process— Special fields store ETL processing history that enables you to trace which ETL process created or updated what piece of data. You can use this data in data-quality investigations—for example, when a review of a particular reporting item requires identification and review of other items that were processed by the same ETL job. For more information, see the *Genesys Info Mart 8.1 Operations Guide.*

Genesys Voice Platform Support

- Genesys Info Mart 8.x supports reporting on the interaction aspect of Genesys Voice Platform (GVP) 7.6 or 8.x activity. For GVP 8.x, it does not matter whether GVP has been configured for computer-telephony integration (CTI) through SIP Server or through IVR Server, provided that GVP has been configured either as a series of `Voice Treatment Port` DNs or as a `Trunk Group` DN. For information about how to configure GVP as a series of `Voice Treatment Port` DNs or as a `Trunk Group` DN, see the *Voice Platform Solution 8.1 Integration Guide.*

  Genesys Info Mart does not support configurations in which GVP 8.x is used simply to play treatments solely under the control of the GVP Media Control Platform. For example, if the call is on a Routing Point and GVP is used by the strategy to play a treatment, with no other IVR activity until the call is delivered to an agent, there will be no IRF record for the IVR (in other words, for GVP); no IVR-related metrics will be reported and, except for the route point duration, no mediation activity will be reported.

  Genesys Info Mart does not support detailed reporting on voice application usage (for example, subcallflows), based on data from GVP Voice

Application Reporter (VAR) 7.6 or the GVP 8.x Reporting Server. However, for GVP 7.6 deployments and for GVP 8.x deployments in which GVP is configured for CTI through IVR Server, Genesys Info Mart does report the limited voice application data that is available for any IVR through IVR Server. For information about how to obtain detailed reporting on voice application usage, see the GVP 7.6 or 8.x documentation, as applicable.

Simplified Database Model

- A simplified database model aligns the lowest level of Genesys Info Mart data details with the Interaction Concentrator model.

  Global Interaction Database (GIDB) within the Genesys Info Mart database schema represents a subset of Interaction Database (IDB) tables that consolidates data from one or more IDBs. GIDB, which is a replacement of the segment model that was used in previous releases of Genesys Info Mart, serves as a storage of low-level interaction details, which require a less resource-consuming transformation.

  For more information about the Genesys Info Mart database schema, see the *Genesys Info Mart 8.1 Reference Manual* for your particular RDBMS.

Support for Partitioning

- Genesys Info Mart 8.x supports partitioning of the Info Mart database in Oracle (range partitioning only), Microsoft SQL Server, and, starting with release 8.1.3, in PostgreSQL deployments. Genesys Info Mart provides SQL scripts to create a partitioned database schema out of the box, and Genesys Info Mart jobs automatically create and maintain the partitions.

ETL Cycle

- The ETL cycle consists of two main jobs:
  - The extraction job retrieves all data from available IDBs and, merging data for voice interactions as necessary, consolidates all low-level details data within a single GIDB.
  - The transformation job processes all extracted data, populating dimensions and loading data directly into the fact tables.

  When enabled, the aggregation job aggregates the transformed data in parallel with the ETL cycle, in environments that include GI2 or in which RAA is deployed as a separate package. Within scheduled daily intervals, which are configurable, `Job_AggregateGIM` runs continuously.

  For more information about the ETL jobs, see Chapter 18 on page 371.

Error Handling

- Configuration options enable you to control Genesys Info Mart behavior when Genesys Info Mart encounters errors during transformation. For more information about the error-handling configuration options, see "error-policy Section" on page 276.

Configuration Verification

- A configuration-checking process verifies the validity of environment and application settings for the Genesys Info Mart and ICON applications, as well as the validity of Genesys Info Mart–related options on supporting objects (such as `DNs` or `Media Type Business Attributes`) and the

availability of configured data sources. For more information about the configuration-checking process, see "Deployment Verification" on page 401.

Maintenance

- A maintenance job, `Job_MaintainGIM`, maintains the Info Mart database. In particular, it:
  - Purges eligible data from the Info Mart database.
  - In partitioned databases, maintains partitions.
  - Prepopulates customizable calendar tables.

  For more information about the maintenance job, see "Maintenance" on page 404 and "Purging Info Mart Data" on page 405.

  For information about how to schedule or run the maintenance job, see the chapter about working with Genesys Info Mart jobs in the *Genesys Info Mart 8.1 Operations Guide*.

Flexible DAP Configuration

- To simplify your deployment, you can reuse the non-JDBC DAPs in your deployment, and make these DAPs suitable for Genesys Info Mart to access the same databases. For more information, see "Reusing DAPs" on page 208.

Support for Management GUIs

- Genesys Info Mart 8.x supports the Genesys Info Mart Administration Console, a GUI for managing Genesys Info Mart jobs, on Windows platforms. In addition, starting with release 8.1.4, Genesys Info Mart supports Genesys Info Mart Manager, a web-based GUI for managing Genesys Info Mart jobs on any platform.

New in This Release

This section describes new or changed functionality that was introduced in Genesys Info Mart 8.1.x releases:

New in Release 8.1.4

This section describes the new or changed functionality that was introduced in release 8.1.4.

Note: Genesys Info Mart 8.1.4 requires Interaction Concentrator 8.1.000.24 or higher. For additional information, see Table 1 on page 64.

Genesys Info Mart Manager

- Genesys Info Mart Manager is a new, web-based GUI to manage Genesys Info Mart operations. Genesys Info Mart Manager provides the same functionality as the Genesys Info Mart Administration Console, except for the ability to schedule a job on an ad-hoc basis to run at a future time. Genesys Info Mart release 8.1.4 supports using either Genesys Info Mart Manager or the Administration Console, or both, to manage day-to-day operations.

  For information about deploying Genesys Info Mart Manager, which is implemented as a plug-in for Genesys Administrator Extension (GAX) 8.5.0 and later, see "Installing Genesys Info Mart Manager" on page 330; see also the Document Change History item about Genesys Info Mart Manager on page 15. For information about using Genesys Info Mart Manager, see the *Genesys Info Mart 8.1 Operations Guide* or the Help in Genesys Info Mart Manager itself.

Improved Aggregation Management

- In deployments that include GI2 or the separately installed RAA package:
  - You can now issue re-aggregation requests from Genesys Info Mart Manager or the Genesys Info Mart Administration Console. Previously, re-aggregation requests could be issued only from the command line.
  - RAA log messages now use the Genesys Info Mart log format, and RAA messages are now available in Solution Control Interface (SCI). For more information, see the descriptions of the Genesys Info Mart log events in the *Framework 8.x Combined Log Events Help*.

Support for Data Access Restrictions in GI2

- Genesys Info Mart supports GI2 functionality to control visibility of GI2 data and reports. New dimension tables in the Info Mart database—`GROUP_ANNEX` and `RESOURCE_ANNEX`, which are populated from data in a new GIDB table, `GIDB_GC_ANNEX`—store configuration data that enables GI2 to control visibility of certain data and reports, based on attributes such as a user's location or line of business.

  A new configuration option, `days-to-keep-deleted-annex` in the `[gim-etl]` section, enables you to specify a data-retention policy for deleted `*_ANNEX` dimension data. For more information, see the option description on page 282.

  For more information about the new Info Mart database tables, see the *Genesys Info Mart 8.1 Reference Manual*.

The configuration data is available in the Configuration details IDB only if the `callconcentrator.cfg-annex` option is configured in the `cfg` ICON application. This option was introduced in ICON release 8.1.4. For information about configuring ICON to store the required data, see the *Interaction Concentrator 8.1 User's Guide* for release 8.1.4.

For information about configuring Genesys person and group configuration objects to enable role-based data access control in GI2, see the *Genesys Interactive Insights 8.1 User's Guide* for release 8.1.4.

Multimedia Transformation Improvements

- Various internal enhancements, mostly arising from a more incremental approach to transformation processing, improve the performance and data quality of multimedia transformation, particularly for very long-living interactions or interactions with inappropriate routing strategies that result in very large numbers of parties or virtual queues ("runaway strategies").

  To prevent performance and out-of-memory issues with runaway strategies, it is no longer necessary to artificially limit how much of the data that is associated with a single interaction will be selected for transformation in a single chunk. Therefore:

  - The `max-parties-per-call` configuration option (see page 293) has been discontinued.
  - Log message numbers 55-20120, 55-20121, and 55-20122 are no longer generated.

Agent Transformation Improvements

- Changes in agent-activity processing improve the performance and data quality of agent transformation. The most significant observable change is a new configuration option, `ignored-reason-codes` in the `[gim-transformation]` section, which enables you to filter out agent-state reason codes that are not useful for reporting. For more information, see the option description on page 310.

New Technical Descriptors

- The following new technical descriptor combinations have been introduced:

  - `InConference/Redirected/PulledBack`, which can occur in multimedia deployments with ICON release 8.1.200.18 or later, when an agent leaves a chat interaction abnormally. For more information about the scenarios for which Genesys Info Mart reports the new technical descriptor combination, see the description of `Redirected/Pullback` in the table of technical results and technical result reasons in the Technical Descriptors chapter in the *Genesys Info Mart 8.1 User's Guide*.

    As a related change, the role reason and technical result reason of `PulledBackTimeout` have been renamed to `PulledBack`, because Interaction Server might have pulled the interaction back for reasons other than a timeout.

- `RoutedTo/Redirected/PulledBack`, which can occur in multimedia deployments when an interaction is routed to an agent and is then pulled back from the agent into an interaction queue (for example, after a timeout).

get_gim_data

- `get_gim_data`, a utility to extract Info Mart data for troubleshooting purposes, is now included in the Genesys Info Mart IP. Previously, `get_gim_data` was obtainable only from Genesys Customer Care.

  There are basic instructions for using `get_gim_data` in the `readme.txt` file. Genesys expects that you will use `get_gim_data` only under the direction of Genesys Customer Care, who will provide full instructions.

Outbound Contact Topologies

- Starting with release 8.1.301.07, Genesys Info Mart no longer requires you to use separate Interaction Concentrator instances or HA sets to store Voice and Outbound Contact details. Previously, Genesys Info Mart enforced this requirement. However, Genesys continues to recommend separate Interaction Concentrator instances or HA sets for Voice and Outbound Contact data. The information in the remainder of the Genesys Info Mart documentation suite assumes that you observe this recommendation.

OS and RDBMS Changes

- For information about operating system (OS) and RDBMS changes in this release, see the *Genesys Supported Operating Environment Reference Guide*.

New in Release 8.1.3

This section describes the new or changed functionality that was introduced in release 8.1.3 or carried over from enhancements in interim Genesys Info Mart 8.1.2 releases.

Note: Genesys Info Mart 8.1.3 requires Interaction Concentrator 8.1.000.19 or higher. For additional information, see Table 1 on page 64.

Support for Business Edition Premise (formerly Genesys One)

- Starting with release 8.1.202.01, Genesys Info Mart provides historical reporting in Business Edition Premise (formerly Genesys One), the all-in-one on-premise Genesys contact center solution.

  When deployed with Business Edition Premise, Genesys Info Mart provides all functionality except support for custom user data. In Business Edition Premise deployments, Genesys Info Mart limits user-data transformation to user data that populates the `INTERACTION_DESCRIPTOR` and `IRF_USER_DATA_GEN_1` tables (in other words, the KVPs that are identified with a single asterisk in Table 7 on page 122).

Security Enhancements

- Transport Layer Security (TLS) features have been enhanced to include support for:
  - Mutual TLS
  - Compliance with Federal Information Processing Standards (FIPS)

Support for PostgreSQL

- Genesys Info Mart supports using PostgreSQL 9.2 for the Info Mart database. The Info Mart database on PostgreSQL must be partitioned.

  A new job, `Job_UpdateStats`, performs supplementary maintenance on the Info Mart database on PostgreSQL. New configuration options in the `[schedule]` section, `run-update-stats` (see page 318) and `update-stats-schedule` (see page 319), control the job schedule.

  A new document, the *Genesys Info Mart 8.1 PostgreSQL Reference Manual,* provides information about the Info Mart database schema on PostgreSQL.

Other OS and RDBMS Changes

- For information about additional operating system (OS) and RDBMS changes in this release, see the *Genesys Supported Operating Environment Reference Guide*.

Improved Reporting for iWD

- For intelligent Workload Distribution (iWD) and other scenarios in which multimedia interactions are placed into "archiving" Interaction Queues, instead of being terminated immediately after processing, special handling enables Genesys Info Mart to provide:
  - More accurate reporting data—Genesys Info Mart can now indicate that an agent completed work on an interaction when the agent moves the interaction into an archive queue. Previously, Genesys Info Mart indicated that the agent transferred the interaction.
  - Improved performance—Genesys Info Mart now terminates the multimedia interaction fact when the interaction is moved to an archive queue. Previously, the interaction remained active in the Info Mart database until an iWD strategy terminated the archived interaction or Genesys Info Mart artificially terminated the interaction in accordance with the `days-to-keep-active-facts` setting (default 30 days). Such long-living interaction facts caused performance issues for the ETL and for RAA aggregation.

  To support the improved reporting behavior and results:
  - Two new configuration options in the `[gim-transformation]` section, `completed-queues` and `canceled-queues`, enable you to specify which Interaction Queues are archive queues. For more information, see the descriptions of the options starting on page 309.
  - Two new technical result reasons, `ARCHIVED` and `CANCELED`, are used in combination with the technical result of `COMPLETED` to identify when a handling resource places an interaction into one of the archive queues.

Agent Transformation Improvements

- Changes in agent-activity processing improve the performance and data quality of agent transformation. The most significant observable changes are:
  - A new configuration option, `max-state-duration` in the `[gim-etl]` section, which enables you to limit the duration of agent states. For more information, see the description of the option on page 294.
  - Additional changes related to the implementation of `max-session-duration-in-hours` (see page 294).

Miscellaneous

- Minor schema changes improve database performance and enhance data quality. For more information, see the "New in Release 8.1.3" section in the *Genesys Info Mart 8.1 Reference Manual* for your RDBMS.
- Various enhancements improve Genesys Info Mart reliability, performance, and data quality. Software improvements include the following:
  - A new configuration option, `pipeline-timeout-in-hours` in the `[gim-transformation]` section, enables you to control how long Genesys Info Mart allows for the execution of a single transformation pipeline. If the timeout is exceeded, Genesys Info Mart tries to abort the pipeline, and the transformation job fails. For more information, see the description of the option on page 312.
  - Starting with release 8.1.202.01, a new startup parameter, `cfg.protocol.timeout`, enables you to customize the Configuration Server response timeout. Increasing the timeout helps to improve resilience to low-bandwidth or very busy network conditions when Genesys Info Mart reads configuration information. For more information, see "Modifying JVM Startup Parameters" on page 360.
  - Additional enhancements optimize the loading of configuration data during Genesys Info Mart startup, as well as processing of configuration updates.
  - Enhanced logging about application connections.

New in Release 8.1.2

This section describes the new or changed functionality that was introduced in release 8.1.2 or carried over from enhancements in interim Genesys Info Mart 8.1.1 releases.

Note: Genesys Info Mart 8.1.2 requires Interaction Concentrator 8.1.000.19 or higher. For additional information, see Table 1 on page 64.

Support for the SIP Cluster Solution

- Genesys Info Mart can be deployed in the usual way to provide historical reporting in the SIP Cluster solution. There are no special configuration requirements and no special features that are unique to SIP Cluster deployments. However, there are some internal operational and data differences in the way that Genesys Info Mart functions in a SIP Cluster deployment. For more information, see the "Historical Reporting" section in the *SIP Cluster Solution Guide,* which is provided as wiki pages.

Note: Genesys Info Mart supports a restricted release of SIP Cluster. For information on Genesys SIP Cluster technology, contact your Genesys representative.

Enhanced Support for Distributed Data Centers

- When high availability (HA) of Interaction Concentrator is implemented across multiple data centers, a new database access point (DAP) configuration option, `geo-location` in the `[gim-etl]` section, enables the Genesys Info Mart extraction job to give preference to the IDBs that are local to the Info Mart database, provided that data quality in a remote IDB is not superior. For more information about the `geo-location` option, see Table 11 on page 214 and Table 12 on page 220. For more information about how geolocation affects Genesys Info Mart extraction behavior, see "HA Data Extraction with Geolocation" on page 385.

Support for Business Continuity

- Genesys Info Mart 8.1.x supports the use of Oracle GoldenGate for database replication, to provide Disaster Recovery for reporting data when a site failure occurs.

  For full information about setting up and using Oracle GoldenGate, see the *Genesys Info Mart Business Continuity Deployment Guide*, which is provided as wiki pages. Refer to the *Genesys Info Mart 8.1 Release Notes* for details about the location of the wiki documentation.

Enhanced Support for User Data

- Starting with release 8.1.201, Genesys Info Mart supports storing custom user-data facts as date/time data types. By default, Genesys Info Mart converts date/time KVP values to date/time using `yyyy-mm-ddThh24:mi:ss.ff` format. A new column, `CONVERT_EXPRESSION`, in the `CTL_UD_TO_UDE_MAPPING` table enables you to customize the conversion expression to convert from a different date format. For more information, see Procedure: Customizing the user-data template script, on page 185.
- Genesys Info Mart now optionally stores the user data for interactions that are in mediation, including user data for interactions that are not distributed to a handling resource after the mediation. A new `DN` and `Script` configuration option, `link-msf-userdata` in the `[gim-etl]` section, enables this functionality. For links to more information about user data in Mediation Segment Fact (MSF) records, see the Document Change History on page 21. For more information about how MSF user data is stored in the Info Mart database, see the "New in Release 8.1.2" section in the *Genesys Info Mart 8.1 Reference Manual* for your RDBMS.
- Starting with release 8.1.101.07, a new startup parameter, `filterUserData`, enables you to control whether Genesys Info Mart will filter the call-based attached data in IDB to extract only from the `UserData` attribute of data-source events, or whether it will also extract from the `Reasons` and `Extensions` attributes. For more information, see "Source Attributes in Events" on page 107. For information about changing the parameter value, see "Modifying JVM Startup Parameters" on page 360.

Enhanced Support for Tenant Reporting

- Genesys Info Mart now supports creating tenant views for a group of tenants or for all tenants, as well as for a single tenant. This functionality simplifies deployment of the reporting solution, as described in "Creating Genesys Info Mart Read-Only Tenant Views" on page 348.

Deployment Simplification

- The Genesys Info Mart IP includes a new utility called User Data Assistant, a Microsoft Excel 2007 file, to help you configure ICON and Genesys Info Mart to capture and store custom user data. For more information, see "Enabling Storage of User Data" on page 148.
- `Job_InitializeGIM` and `Job_MigrateGIM` now *automatically* execute the scripts to modify IDB schemas for use with Genesys Info Mart when required. For more information, see "Preparing IDBs" on page 178 and "IDB Schema Compatibility" on page 403.

Enhanced Partitioning Support

- Support for database partitioning has been extended to include certain Control tables in the Info Mart database. In particular, audit and history tables can now be partitioned.

  For information about the new fields that have been added to enable partitioning of these tables, as well as information about other enhancements to improve database performance, see the "New in Release 8.1.2" section in the *Genesys Info Mart 8.1 Reference Manual* for your RDBMS.

Enhanced Multimedia Processing

- Processing of multimedia interactions that involve very large numbers of parties or virtual queues has been enhanced, resulting in improved data quality and transformation behavior when Genesys Info Mart abbreviates the representation of repeatedly unsuccessful routing attempts. For information about related changes to the calculation of mediation durations that are reported in the IRFs and MSFs, see the *Genesys Info Mart 8.1 Reference Manual* for your RDBMS.

  Starting with release 8.1.103.03, a new configuration option, `max-parties-per-call` in the `[gim-etl]` section, enables you to control how much of the data that is associated with a single interaction will be selected for transformation. Be aware that, by default, the number of parties and virtual queues that triggers the abbreviated representation of "runaway strategy" scenarios is smaller than it was in earlier 8.1 releases. For more information, see the description of the `max-parties-per-call` option on page 293.

Improved Purge Processing

- Improvements in purge processing reduce job and data problems that were sometimes encountered in release 8.1.1, arising primarily from the different characteristics of voice and multimedia interactions. The most significant observable changes are:
  - The values of the `days-to-keep-active-facts` and `days-to-keep-gidb-facts` configuration options must always be smaller than `days-to-keep-gim-facts`, and the configuration-check process enforces this dependency. For more information about the options, see the descriptions starting on page 281.
  - The usual point from which Genesys Info Mart calculates the `days-to-keep-*-facts` purge thresholds is the transformation high-water mark (HWM). In release 8.1.1, it was the extraction HWM.

Miscellaneous

- This release simplifies the changes in configuration that are required when a data source that is monitored by a dedicated ICON temporarily becomes unavailable. To avoid delays in data processing, you can exclude the data source from data extraction by disabling the configuration objects for the relevant ICON and DAP `Applications`. The previous method that required removal of the relevant ICON and DAP `Applications` from the Genesys Info Mart `Application` connections is still supported.
- In Outbound Contact deployments with a partitioned Info Mart database, a new configuration option, `max-chain-processing-duration-in-hours` in the `[gim-etl]` section, enables you to customize the expected maximum durations of chain processing. For more information, see the description of the option on page 292.
- Starting with release 8.1.103.07, a new configuration option, `etl-start-date` in the `[gim-etl]` section, enables you to control the date from which Genesys Info Mart starts extracting data. For more information, see the description of the option on page 287.
- Starting with release 8.1.103.03, a new startup parameter, `queryParallelism`, enables you to control the degree of query parallelism that the Oracle RDBMS will use, to improve performance. For more information, see “Modifying JVM Startup Parameters” on page 360.
- Enhanced notifications for RAA improve aggregation performance.
- When an option value is changed to an invalid value, Genesys Info Mart now uses the default value of the option. In releases earlier than 8.1.2, Genesys Info Mart ignored the change and continued to use the old option value.
- Genesys Info Mart supports IDB partitioning, which was introduced in ICON release 8.1.1. There are no additional configuration requirements or operational considerations on the Genesys Info Mart side.

  For information about partitioning on IDBs, see the *Interaction Concentrator 8.1 Deployment Guide* and the *Interaction Concentrator 8.1 User's Guide*. For important considerations about purging IDB, see also the chapter about managing data sources in the *Genesys Info Mart 8.1 Operations Guide*.
- Genesys Info Mart supports Internet Protocol version 6 (IPv6). For information about configuring and using IPv6 in your Genesys Info Mart deployment, see the *Framework 8.1 Deployment Guide*. For an example of the JDBC URL information you must provide in the DAPs that Genesys Info Mart uses to access the Info Mart database and IDBs, see page 217.

New in Release 8.1.1

This section describes the new or changed functionality that was introduced in release 8.1.1 or carried over from enhancements in interim Genesys Info Mart 8.1.0 releases.

Note: Genesys Info Mart 8.1.1 requires Interaction Concentrator 8.1 or higher. For additional information, see Table 1 on page 64.

Enhanced Information in MSFs and IRFs

- New fields in the `MEDIATION_SEGMENT_FACT` (MSF) and `INTERACTION_RESOURCE_FACT` (IRF) tables provide additional information about relationships between MSF activity and IRFs, regardless of whether attempts to reach the handling resource were successful or were abandoned or cleared from virtual queues. The new fields also enhance capabilities for reporting on queue activity and transfer details.

  To enable this functionality, you must use Universal Routing Server (URS) release 8.1.100.08 or later, and you must set the ICON `route-res-vqid-hist-enabled` configuration option to `true`, so that ICON populates the `G_ROUTE_RES_VQ_HIST` table in IDB.

  For more information about the new fields, see the "New in Release 8.1.1" section in the *Genesys Info Mart 8.1 Reference Manual* for your RDBMS. For information about the minimum releases of ICON and URS that are required to support certain functionality, see Table 1 on page 64.

Workbin Dimension

- A new dimension table, `WORKBIN`, stores information about all the workbin instances that have been associated with multimedia interactions in the environment. In cases in which placement of a multimedia interaction in a workbin is considered to be mediation, the MSF references this dimension table, to identify the type of resource that is associated with the workbin and the specific resource that is associated with the workbin mediation.

  For the meaning of *workbin instance* in Genesys Info Mart, see "Workbin Instance and Personal Workbin" on page 40.

Stop Action Enhancements for Multimedia

- Use of the `STOP_ACTION` flag in the IRF table has been extended to multimedia interactions, for which the flag indicates whether the interaction was stopped by an IRF resource or by an outside entity that was not a party to the interaction (for example, Interaction Server or a Media Server).

  For more information, see the description of the `STOP_ACTION` column in the IRF table in the *Genesys Info Mart 8.1 Reference Manual* for your RDBMS.

- A new technical result, `AbnormalStop`, and associated technical result reasons, `AbnormalStopWhileRinging` and `AbnormalStopWhileQueued`, have been introduced.

  For more information, see the table of technical results and technical result reasons in the chapter about technical descriptors in the *Genesys Info*

Mart 8.1 User's Guide. For the available technical descriptor combinations that include the new technical result, see the appendix in the *User's Guide.*

Support for Thread Metrics

- Earlier releases of Genesys Info Mart 8.x provided information in interaction fact (IF) records to identify interaction facts that belong to the same interaction thread (in other words, to identify whether a new Interaction is a follow-on to a prior interaction). Genesys Info Mart now provides additional data in IRF records that enable downstream reporting applications to identify whether the IRF represents:
  - This agent's first participation in the interaction (where "this agent" refers to the agent resource that is the subject of the IRF).
  - This agent's first participation in a reply within the interaction (which might or might not be the agent's first participation in the interaction).
  - This agent's first participation in the interaction thread.
  - This agent's first participation in a reply within the interaction thread (which might or might not be the agent's first participation in the thread).
  - The first participation by any handling resource in the interaction thread. For example, if a strategy that sends an `AutoResponse` is the first handling resource to participate in the interaction thread, the IRF for the strategy will indicate this.

  To implement this functionality:
  - A new configuration option in the `[gim-etl]` section, `max-thread-duration-after-inactive-in-days`, enables you to control how long Genesys Info Mart keeps track of inactive threads (in other words, threads in which all interactions in the thread have terminated). For information about the option, see page 295.
  - The rules that govern population of the `MEDIA_SERVER_ROOT_IXN_ID` and `MEDIA_SERVER_ROOT_IXN_GUID` fields in the IF record have changed; these fields now depend on the tracking of inactive threads. Genesys Info Mart reliably identifies interactions that belong to the same thread, provided that the thread has not timed out.
  - Schema changes have been made, to include internal staging tables and a new dimension table, `ANCHOR_FLAGS`, which relates to a new IRF column, `ANCHOR_FLAGS_KEY`. For more information about the schema changes, see the *Genesys Info Mart 8.1 Reference Manual* for your RDBMS.

  For information about how interaction threads are identified, see the subsection about interaction threads, in the section about populating interaction data, in the *Genesys Info Mart 8.1 User's Guide*.

Support for Fiscal Calendars

- You can now configure either the default or custom calendars in Genesys Info Mart to provide calendar dimensions for fiscal years. Two new configuration options in the `[date-time]` section, `fiscal-year-start` and `fiscal-year-week-pattern`, enable this functionality.

For more information about the calendar-related configuration options, see their respective descriptions in "date-time Section" on page 271. For links to more information about how to create and maintain calendars, see "Maintaining Calendars" on page 405.

Enhancement of MSF Time Metrics

- Genesys Info Mart no longer uses URS timestamps to identify the start or end times of MSF records. This prevents gaps or overlaps in MSF time metrics when URS and Interaction Server are not synchronized.

  Enhancements in Interaction Concentrator 8.1 enable Genesys Info Mart to increase the accuracy of MSF time reporting for virtual queues, provided that you configure ICON to populate the `G_ROUTE_RES_VQ_HIST` table (see the description of the `route-res-vqid-hist-enabled` option on page 169).
- In general, the mediation duration of MSFs for ACD Queues, Interaction Queues, and Interaction Workbins does not include strategy time, except for scenarios that are described in connection with the `populate-mm-ixnqueue-facts` and `populate-mm-workbin-facts` options starting on page 304. In previous releases, mediation duration included all the time that the interaction spent in the last strategy before the interaction was distributed to a handling resource. Starting with release 8.1.1, Genesys Info Mart reports only the actual time that the interaction spent in these mediation resources.

  A new configuration option in the `[gim-transformation]` section, `adjust-vq-time-by-strategy-time`, enables you to control whether the time that interactions spend in strategies but not in associated virtual queues will be included in the MSF records for virtual queues. Setting the option to exclude the strategy time that is outside the virtual queue (the default setting) means that there might be gaps in time between MSFs for virtual queues and other MSFs or IRFs. For information about the option, see page 309.

Artificial Termination of Long-Living Interactions

- As an enhancement to existing purge functionality, the `days-to-keep-active-facts` configuration option sets the maximum duration for long-living multimedia interactions, after which Genesys Info Mart artificially terminates them. The enhancement increases flexibility in the management of long-living interactions. For more information about the option, see page 281.

Security Features

- Genesys Info Mart supports secured connections between Genesys Info Mart Server and Configuration Server and between Genesys Info Mart Server and Message Server, by means of:
  - Transport Layer Security (TLS)—You can configure the applications to use the TLS protocol for the connections between them. With features such as authentication and data encryption, TLS enables you to secure data and service integrity.

- Client-side port definition—In the Genesys Info Mart configuration, you can predefine the connection parameters that Genesys Info Mart will use to connect to the Configuration Server and Message Server applications. This feature enables you to secure access for Genesys Info Mart across a firewall and to maintain service availability.

These optional security features operate independently of each other; thus, depending on your requirements, you can configure either, both, or neither.

For general information about these Genesys security features, see the *Genesys 8.x Security Deployment Guide.* For more information about Genesys Info Mart system requirements for these features, see "Security Requirements" on page 63. For more information about how to enable these security features in Genesys Info Mart, see "Enabling Secure Connection Features" on page 154.

Note: Genesys Info Mart does not support secured connections between Genesys Info Mart Server and the Genesys Info Mart Administration Console or Genesys Info Mart Manager.

- Genesys Info Mart supports secured connections between Genesys Info Mart Server and the IDB and Info Mart databases on Oracle and Microsoft SQL Server, by using the Secure Socket Layer (SSL) protocol on the JDBC connection.

  To set up SSL for the JDBC connection, configure the `jdbc-url` option in the applicable DAPs. For more information about how to configure the `jdbc-url` option, see Procedure: Configuring JDBC DAPs, on page 225 or Procedure: Configuring non-JDBC DAPs, on page 228. For information about the specific parameters that you must specify in the `jdbc-url` option, see your RDBMS vendor's documentation.
- Genesys Info Mart supports the Transparent Database Encryption (TDE) features that Oracle and Microsoft SQL Server provide, to enable secure storage of data in the Info Mart database.

  For more information about how to enable and use database encryption, see your RDBMS vendor's documentation.

Performance Enhancements

- A new configuration option in the `[gim-transformation]` section, `ud-io-parallelism`, enables you to enhance performance of the transformation job in deployments that involve large amounts of user data in historical reporting. For information about the option, see page 312.
- Internal changes to the processing of agent activity improve the performance of `SM_*` transformation. As a result, it is no longer necessary to control population of the `SM_RES_*_FACT` tables, and hence the `populate-sm-resource-*-facts` options have been discontinued. Genesys Info Mart always populates the `SM_RES_*_FACT` tables.

Miscellaneous

- A new column in the RESOURCE_ dimension table, SWITCH_DBID, enables downstream reporting applications to reliably identify the switch that is associated with switch resources. The field is populated from the GC_ENDPOINT table in the Configuration details IDB.
- The behavior, default value, and range of valid values of the user-event-data-timeout option (see page 301) have changed. The transformation job now calculates the timeout for UserEvent-based KVP data based on the end time of the call.
- A new user-data propagation rule, IRF_FIRST_UPDATE, enhances the flexibility of user-data reporting, particularly in two-step transfer, consultation, or conference scenarios and in route on no answer (RONA) scenarios. For more information, see "Propagation Rules" on page 118.
- A new configuration option in the [gim-etl] section, merge-failed-is-link-timeout (see page 297), specifies the time interval, in seconds, for which the merge of failed IS-Links will be delayed to enable Genesys Info Mart to receive both sides of the links. In earlier Genesys Info Mart 8.x releases, the merge of failed links was performed immediately, without waiting for the other side of the failed link, thus resulting in dangling links.

New in Release 8.1.0

This section describes the new or changed functionality that was introduced in the initial 8.1 release of Genesys Info Mart or carried over from enhancements in Genesys Info Mart 8.0.x releases.

Numeric User Data

- Genesys Info Mart supports numeric data types, in addition to character data types, for KVP values in user-data fact tables. For more information, see the description of user-data transformation on pages 104 and 106.

Identifying Who Released the Call

- For voice interactions from those T-Servers that support this feature, and provided that ICON has been configured to store the required information, Genesys Info Mart provides indication of whether a particular handling resource (for example, agent or self-service IVR) initiated release of the call.

  A new field in the IRF table—STOP_ACTION—acts as a flag. For more information about the STOP_ACTION field, see the *Genesys Info Mart 8.1 Reference Manual*.

  For more information about how to configure ICON to support this functionality, see store-releasing-party on page 172.

Multimedia Enhancements

- Genesys Info Mart supports reporting on placement of a multimedia interaction into an Interaction Workbin as hold.
  - populate-workbin-as-hold, a new option in the [gim-etl-populate] section, enables you to specify whether placement of an interaction into an Interaction Workbin by an Agent or Place resource is considered as hold or as mediation. For more information about this option, see page 307.

- For workbin activity that is considered to be hold, various hold metrics are reported in the IRF table. For more information, see the *Genesys Info Mart 8.1 Reference Manual*.

- Genesys Info Mart reports Initiated Conference metrics for multimedia interactions. In the IRF table, the `CONF_INIT_TALK_COUNT` and `CONF_INIT_TALK_DURATION` fields are now populated for multimedia interactions as well as for voice interactions. For more information, see the *Genesys Info Mart 8.1 Reference Manual*.

Configuration Check Enhancements

- As part of the configuration check, Genesys Info Mart verifies the following additional items:
  - The options that control ETL frequency and the chunk size of extracted data are internally consistent.
  - Based on the roles that are configured for the extraction DAPs, Genesys Info Mart is set up to extract Voice, Multimedia, and Outbound Contact details from separate Interaction Concentrator instances. This check minimizes the possibility of errors that arise from deployment topologies that are not supported.

  For more information, see "Configuration Check" on page 402.

Additional Outbound Contact Metrics

- New fields in the `CONTACT_ATTEMPT_FACT` table report additional metrics for the following:
  - Dialing—`CPD_DIAL_COUNT, CPD_DIAL_DURATION_MS`
  - Call Progress Detection (CPD)—`CPD_COUNT, CPD_DURATION_MS`
  - Transfers—`CPD_TRANSFER_COUNT, CPD_TRANSFER_DURATION_MS`

  For more information about the new fields, see the *Genesys Info Mart 8.1 Reference Manual*.

Error Handling

- Modifications to the error-policy configuration options and to the Info Mart database schema enhance the flexibility of Genesys Info Mart behavior when the transformation job encounters data inconsistencies.
  - A new field, `STATUS`, in the `INTERACTION_FACT` table serves as an indicator of data quality. Depending on configuration, the `STATUS` field reports information about data inconsistencies that are encountered.
  - An option in the `[error-policy]` section, `error-policy-campaign-group-missing`, enables you to control Genesys Info Mart behavior when an Outbound Contact campaign record refers to a campaign group for which there is no group record. The option was introduced in an 8.0.1 release. For more information, see page 276.
  - The default value of most interaction-level `error-policy-*` options is `resume`.

Supported Topologies

- Starting with release 8.0.100.21, Genesys Info Mart supports topologies in which multiple Interaction Servers are monitored by a single Interaction Concentrator instance. For more information about the topologies that Genesys Info Mart supports, see "Interaction Concentrator Topologies" on page 70.

Route-thru-Queue Target

- An option in the `[gim-transformation]` section, `msf-target-route-thru-queue`, enables you to specify which party is recorded as the target of mediation in the MSF table in scenarios where a voice call is routed from a Routing Point through an ACD queue to an agent. The option was introduced in an 8.0.1 release. For more information, see page 311.

Increased Precision of Voice Agent States Reporting

- The precision with which voice agent states are reported has been increased. Internal calculations of summarized agent states and state reasons are now made in milliseconds for the `SM_RES_SESSION_FACT`, `SM_RES_STATE_FACT`, and `SM_RES_STATE_REASON_FACT` tables. This improvement results in more accurate reporting of states and state reasons that occur within the same second for agents handling voice interactions.

Miscellaneous

- Various internal enhancements improve Genesys Info Mart reliability, performance, and data quality. Software improvements that might result in observable behavioral changes include the following:
  - If you set the ICON `dest-busy-processing` configuration option (see page 172) to `true`, Genesys Info Mart now reports a technical result of `DESTINATIONBUSY` in the IRF table when an agent dials a busy destination. Previously in these scenarios, Genesys Info Mart reported a technical result of `CUSTOMERABANDONED`.
  - In scenarios that involve a queueing resource that interacts with another mediation resource while an interaction is queued, `Job_TransformGIM` now identifies the parallel processing and represents it as a single IRF. Previously, this type of parallel processing was represented by multiple IRF records.

    For more information, see “Transforming Calls in Parallel Queues” on page 390.
  - Additional transformation logic protects Genesys Info Mart from being overwhelmed by strategies that cause excessive numbers of parties and virtual queues to be involved in a single interaction. For example, if a strategy keeps retrying busy agents every second, 1000 parties could be associated with the interaction within less than 10 minutes. When the transformation job detects these scenarios, it abbreviates the representation of unsuccessful routing attempts.

    Log events (message numbers 55-20120, 55-20121, or 55-20122) identify when excessive numbers of Party, Party History, or Virtual Queue records, respectively, trigger this behavior. You can set an alarm on the log messages, to prompt you to investigate strategies that might be inappropriate for your deployment.

Chapter

2 System Requirements

This chapter describes the system requirements for Genesys Info Mart. It contains the following sections:

Supported Operating Systems and Databases

For information about operating systems and the relational database management systems (RDBMSs) that Genesys Info Mart and Genesys Info Mart Administration Console support, see the *Genesys Supported Operating Environment Reference Guide* document, which is available on the Genesys Documentation website.

Interoperability Requirements

Genesys Info Mart can operate only with the Genesys components that are listed in Table 1 on page 64. Other Genesys software components that you might have in an environment with Genesys Info Mart must be compatible with Interaction Concentrator 8.x.

Note: For Genesys Info Mart to provide accurate and reliable data, the system clocks on all hosts on which Genesys applications are running (for example, T-Servers, Interaction Servers, and Universal Routing Servers) must be synchronized.

If you plan to use the Genesys Info Mart Administration Console with Genesys Info Mart, Genesys recommends that you install both the Administration Console and Genesys Configuration Manager 8.x on the machine on which an administrator will administer the Genesys Info Mart jobs.

Note: There is *no* requirement that the Genesys Info Mart Server and the Genesys Info Mart Administration Console reside on the same machine. The Genesys Info Mart Administration Console displays timestamps in the local time of the host on which it is installed.

For specific interoperability requirements, see the *Genesys Interoperability Guide*.

Software and Database Requirements

The following software must be installed on the Genesys Info Mart Server host to support Genesys Info Mart 8.x:

- Java Development Kit (JDK) or Server Java Runtime Environment (Server JRE) (see Java)
- Java Database Connectivity (JDBC) driver (see page 63)

System Resources

The Genesys Info Mart Administration Console requires approximately 10 MB of hard-disk space to accommodate the installed program. It does not require any additional RAM beyond what Genesys Configuration Manager and Wizard Framework require.

Genesys Info Mart Manager, the web-based management GUI that is available starting with release 8.1.4, requires negligible hard-disk space on the Genesys Administrator Extension (GAX) host. Genesys Info Mart Manager requires no additional installations on client hosts and no additional RAM.

The Genesys Info Mart Server requires approximately 20 MB of hard disk space and a sufficient amount of disk space for local log files. It requires a minimum of 1 GB of additional RAM, depending on the configuration options that are set to define the data chunk size. For more information about these options, see “Configuration Recommendations” on page 391.

Java

Genesys Info Mart uses the Server Java Virtual Machine (JVM) that ships with the JDK or Server JRE packagings of Java. (The non-Server JRE packaging of Java does not include the Server JVM.) You must install the JDK or Server JRE on the server on which you plan to install the Genesys Info Mart Server.

Several Genesys Info Mart software components use Java. Note that Genesys Info Mart operates with 32-bit or 64-bit versions of Java. For the Java versions that Genesys Info Mart supports, see the *Genesys Supported Operating Environment Reference Guide*.

Note: Periodically check that the daylight saving time (DST) definitions and other time zone information for your Java version are current and correct.

You must also modify your `PATH` and `JAVA_HOME` environment variables so that Genesys Info Mart can locate the Server JVM. The `PATH` and `JAVA_HOME` environment variables that you modify depend on the operating system and user account under which the Genesys Info Mart Server runs.

For specific information about installing the JDK or Server JRE and modifying the environment variables, see "Preparing the Genesys Info Mart Server Host" on page 323.

JDBC Drivers

Genesys Info Mart Server and the ETL jobs use JDBC to access all databases. For specific information about installing the appropriate JDBC driver for your environment, see "Preparing the Genesys Info Mart Server Host" on page 323.

Security Requirements

Starting with release 8.1.1, Genesys Info Mart supports Transport Layer Security (TLS) connections and client-side port definition, to provide secured connections to the Configuration Layer and Management Layer—specifically, to Configuration Server and Message Server. Genesys supports these security features on all the operating systems that Genesys Info Mart supports.

Starting with release 8.1.3, Genesys Info Mart also supports mutual TLS and compliance with Federal Information Processing Standards (FIPS).

On Windows platforms, support for TLS is integrated into the operating system, and there are no additional requirements to enable Genesys Info Mart to support it. On UNIX-based platforms, you must install the Genesys Security Pack on the Genesys Info Mart host (see page 67).

Compatibility with Genesys Software

For general requirements on interoperability with Genesys Configuration Layer, see the *Genesys Interoperability Guide*.

Table 1 lists the Genesys software components with which Genesys Info Mart operates directly. For requirements for various other Genesys software

components that you may have in an environment with Genesys Info Mart, refer to Interaction Concentrator 8.x documentation. The two exceptions are:

- Universal Routing Server (URS) 8.0 or higher.
- Outbound Contact Server (OCS) release 7.6, which Genesys Info Mart 8.1 requires in order to populate Outbound Contact–related data correctly.

For each component, Table 1 provides the *minimum* release number with which Genesys Info Mart release 8.1 is compatible.

Note: Genesys recommends that you install the most recent generally available release of Genesys products. Refer to the corresponding product *Release Notes* for information about new or improved functionality.

Refer to the *Genesys Info Mart 8.1 Release Notes* and *Release Advisory* for any updates to the release requirements for the various components.

Table 1: Genesys Info Mart Compatibility (Minimum Releases)

Area of Functionality	Component/Product Release	Comments
Configuration Layer	• Configuration Server release 7.6 • DB Server release 7.6	Configuration Server release 7.6 or higher provides improved support for the configuration history log. Genesys Info Mart 8.x supports Advanced Disconnect Detection Protocol (ADDP) for the connection from Genesys Info Mart to Configuration Server. Genesys Info Mart 8.1.1 and later supports Transport Layer Security (TLS) protocol for the connection from Genesys Info Mart to Configuration Server. Genesys Info Mart 8.1.3 and later supports mutual TLS and compliance with Federal Information Processing Standards (FIPS). The Genesys Info Mart Administration Console requires DB Server 8.1.3 or later to connect to an Info Mart database on Microsoft SQL Server 2012 or PostgreSQL.

Table 1: Genesys Info Mart Compatibility (Minimum Releases) (Continued)

Area of Functionality	Component/Product Release	Comments
Management Layer	• Local Control Agent (LCA) release 7.6 • Message Server release 7.6 • Solution Control Interface (SCI) release 7.6 • Solution Control Server (SCS) release 7.6	Genesys Info Mart 8.1.1 and later supports TLS protocol for the connection from Genesys Info Mart to Message Server. Genesys Info Mart 8.1.3 and later supports mutual TLS and compliance with FIPS.
User Interaction Layer	Genesys Administrator Extension (GAX) release 8.5.0	Genesys Info Mart requires this component only if you want to use the Genesys Info Mart Manager management GUI, which is supported in Genesys Info Mart 8.1.4 and later.
Interaction Concentrator	• Interaction Concentrator release 8.1 • (For Genesys Info Mart release 8.1.4) Interaction Concentrator release 8.1.000.24	For all Genesys Info Mart 8.1.x releases, Genesys recommends that you use Interaction Concentrator release 8.1.000.24 or later, to enable functionality that improves the reliability of Configuration details and virtual-queue reporting. There are additional minimum requirements for specific functionality. For recommendations for each Genesys Info Mart release, see the *Genesys Info Mart 8.1 Release Note.*
T-Server	T-Server release 7.6	There are additional requirements for the feature to identify the party that released a call. For information about the T-Servers that support this functionality and the minimum release of T-Server that is required, see the *Interaction Concentrator 8.0.x Release Note.*

Table 1: Genesys Info Mart Compatibility (Minimum Releases) (Continued)

Area of Functionality	Component/Product Release	Comments
Interaction Server	eServices/Multimedia Interaction Server release 7.5 or higher	For Genesys intelligent Workload Distribution (iWD) and other deployments that have routing strategies that repeatedly move a multimedia interaction into and out of the same interaction queue or workbin, Genesys recommends that you use Interaction Server release 8.1.200.21 or higher and that you configure your deployment to suppress Interaction Server reporting events and notifications for this kind of routing strategy activity. For more information, see the *Interaction Server 8.1.x Release Note* and the *Universal Routing 8.1 Business Process User's Guide.*
Routing	• For deployments that use virtual queues, Universal Routing Server (URS) release 8.0 • For deployments that do not use virtual queues, URS release 7.2 • To support the Genesys Info Mart feature that links MSF records to the parent IRF, URS release 8.1.100.08 or later	• URS 8.0 is required for processing missing virtual-queue configuration objects. For more information about why Genesys Info Mart requires URS 8.0, see "Missing Virtual Queue Resource Example" on page 397. For information about the minimum release of URS that is required to support this functionality, see the *Interaction Concentrator 8.0.x Release Note.* • For iWD and other deployments that have routing strategies that repeatedly move a multimedia interaction into and out of the same interaction queue or workbin, Genesys recommends that you use Interaction Routing Designer (IRD) release 8.1.3 or higher, to enable you to mark strategies for which redundant activity should not be distributed by Interaction Server. For more information, see the *Universal Routing 8.1 Business Process User's Guide.*

Table 1: Genesys Info Mart Compatibility (Minimum Releases) (Continued)

Area of Functionality	Component/Product Release	Comments
Outbound Contact	Outbound Contact Server (OCS) release 7.6	For information about how to configure Genesys Info Mart and other contact center objects to store reporting data for Outbound Contact campaign activity, see "Enabling Storage of Outbound Contact Details" on page 146.
Genesys Voice Platform	Genesys Voice Platform (GVP) release 7.6	For more information, see "Genesys Voice Platform Support" on page 43.
Security	Genesys Security Pack 8.0	This component is required for UNIX platforms only, for deployments in which you want to use the TLS protocol to secure the connections among Genesys Info Mart, Configuration Server, and Message Server. For more information, see the chapter about Security Pack installation in the *Genesys 8.x Security Deployment Guide*.

Chapter

3 Supported Topologies

This chapter describes and illustrates the various data-source and Genesys Info Mart topologies that Genesys Info Mart 8.1 supports.

This chapter contains the following sections:

Data Domains

The Interaction Concentrator server (ICON) monitors data sources and stores data about data-source activity in the Interaction Database (IDB). Genesys Info Mart extracts data from one or more IDBs, according to the configuration. Genesys Info Mart extracts each of the following data domains separately:

- *ICON Configuration details*—Contains data about contact center configuration objects and object relationships. The data source for Configuration details is Configuration Server.
- *ICON Voice details*—Contains interaction, user data, virtual queue, and agent activity for voice media. The data source for Voice details is T-Server.

> **Note:** In this document, the term *T-Server* is used generically to refer to all T-Server types (premise and network TDM Voice, SIP Server, IVR Server, and Virtual T-Server).

- *ICON Multimedia details*—Contains interaction, attached data, virtual-queue and agent activity for Genesys eServices/Multimedia e-mail, chat, and 3rd Party Media. The data source for Multimedia details is Interaction Server.
- *ICON Outbound Contact details*—Contains data that relates to Outbound Contact activity, such as the history and results of campaigns, chains, and contact attempts; associations between Outbound Contact objects and contact center objects; and precalculated Outbound Contact metrics. The data source for Outbound Contact details is Outbound Contact Server (OCS).

Genesys Info Mart has specific minimum requirements for the types of ICON details that must be included in the deployment. For more information, see "Genesys Info Mart Requirements for ICON Details Storage" on page 70.

Topology Considerations

Genesys Info Mart supports a variety of data-source topologies. The topology that you choose for each data source depends on several deployment-specific factors, including the number of sites, the data network capacity between sites, the interaction volume, and the required level of data-source redundancy or high availability (HA).

Review the topologies on the following pages to determine which ones meet your contact center's needs for performance and HA.

Interaction Concentrator Topologies

In a contact center that has a large Genesys configuration environment or that processes high call volumes—possibly, with large amounts of attached data and/or UserEvent-based key-value pair (KVP) data—you can improve performance of both ICON and Genesys Info Mart by deploying multiple ICON instances to collect data for a particular data domain. When data is stored in multiple IDBs, Genesys Info Mart extracts data from these IDBs in parallel, thus decreasing the extraction time.

Genesys Info Mart Requirements for ICON Details Storage

The Interaction Concentrator topologies that Genesys Info Mart 8.1 supports are similar for all types of ICON details, except for the following special requirements:

- Your deployment must include only one IDB (or one HA set of redundant IDBs) that stores Configuration details.
- Your deployment must include at least one IDB (or one HA set of redundant IDBs) that stores either Voice or Multimedia details.

- You can mix partitioned and non-partitioned IDBs for Voice details or Outbound Contact details within the same deployment.

The following general requirements apply to all supported topologies:

- Each ICON application must populate its own IDB. In other words, consider each ICON-IDB pair (Interaction Concentrator instance) a unit.
- Genesys recommends that each Interaction Concentrator instance (or HA set) process and store data for only one data domain.
  - You must use separate Interaction Concentrator instances (or HA sets) for Voice details and Multimedia details.
  - Starting with release 8.1.301.07, you can combine storage of Outbound Contact details with Voice details. However, Genesys recommends that you use separate Interaction Concentrator instances (or HA sets) to store Voice and Outbound Contact details.
  - You can combine storage of Configuration details with any of the other types of ICON details. However, Genesys recommends that you use a separate Interaction Concentrator instance to store Configuration details.

    To minimize the possibility of missing configuration data, Genesys further recommends that you co-locate the Configuration details IDB on the same host as the Configuration Database (see "Recommendations on Hosting" on page 83).
- You can have one instance or multiple Interaction Concentrator instances (or HA sets) that store Voice, Multimedia, or Outbound Contact details.
- For Voice details, Multimedia details, or Outbound Contact details, each Interaction Concentrator instance can store data from one or multiple instances (or HA pairs) of T-Server, Interaction Server, or OCS, as applicable. In other words, the relationship between the data source(s) and Interaction Concentrator can be one-to-one, many-to-one, or many-to-many.

ICON Roles

The `role` option in the ICON application specifies the type of data that each ICON instance processes; similarly, the `role` option in the Interaction Concentrator database access point (DAP) specifies the type of data that the ICON instance stores in IDB. For a thorough discussion of ICON roles, see the *Interaction Concentrator 8.x Deployment Guide*. For more information about setting the roles that are required for Genesys Info Mart, see Procedure: Configuring the ICON application, on page 163.

When you plan your deployment, consider the following requirements:

ICON Configuration Details

- In all Genesys Info Mart topologies for ICON Configuration details, you cannot have more than one ICON instance (or HA set) monitoring the same Configuration Server (or HA pair) and storing configuration data in the same IDB. To store Configuration details, the `role` option of the ICON

application must contain the value `cfg`. Be aware that the default value of the ICON `role` option is `all`. If you have more than one ICON application in your deployment, ensure that you specifically exclude `cfg` from the value of the `role` option in the ICON applications that will not be storing Configuration details.

ICON Voice or Multimedia Details

- In all Genesys Info Mart topologies for ICON Voice or Multimedia details:
  - The ICON application must be configured to store interaction activity, attached data, virtual queue, resource login, and agent state and work mode details. Optionally, the ICON application may also be configured to store UserEvent-based key-value pair (KVP) data. In other words, the `role` option of the ICON application must contain the values `gcc`, `gud`, and `gls`.
  - A single ICON application (or a single HA set of redundant ICON applications) must record *all* activity for a particular agent. If, for example, a particular agent in your contact center logs in to two switches, the same ICON application (or the same HA set of redundant ICON applications) must monitor both switches.

ICON Outbound Contact Details

- In all Genesys Info Mart topologies for ICON Outbound Contact details, the `role` option of the ICON application must contain only the value `gos`.

All ICON Details

- In all Genesys Info Mart topologies, each ICON must use only one DAP to access its IDB. For example, to store data for a Voice or Multimedia details ICON, do not use separate DAPs for the `gcc`, `gud`, and `gls` ICON providers. Ensure that the `role` configured in the `callconcentrator` section of the DAP matches the ICON role.

Standby and Disaster Recovery

Genesys Info Mart relies on sets of redundant Interaction Concentrator instances to prevent data loss. If the Genesys Info Mart Server is down for a period of time and then becomes available again, it will just resume processing data from where it left off. If the Info Mart database is down for a period of time and then becomes available again, there will be a delay in data available for reporting, but no data loss.

Genesys Info Mart Server Redundancy

To protect against the Genesys Info Mart Server being unavailable for an extended period of time or to enable quick substitution (for example, for server maintenance), you can deploy a second instance of Genesys Info Mart Server to act as a standby. (*Standby* refers to an instance that can be brought online quickly and easily.)

For a description of the topology where there are two instances of the Genesys Info Mart Server with a single instance of the Info Mart database, see “Standby Genesys Info Mart Server Instance” on page 73.

Info Mart Database Redundancy

To protect against the Info Mart database being lost, you can also deploy a second instance of the Info Mart database. There are two types of architecture:

- Two full instances of Genesys Info Mart (Genesys Info Mart Server and Info Mart database) in an *active-active* configuration, as described in "Active-Active Info Marts" on page 73.
- Two full instances of Genesys Info Mart (Genesys Info Mart Server and Info Mart database) in an *active-standby* configuration, as described in "Active-Standby Info Marts" on page 74.

Standby Genesys Info Mart Server Instance

If you deploy a second instance of the Genesys Info Mart Server as a standby for operation against the same Info Mart database, the second instance must be configured in exactly the same way as the first instance, with one exception: To ensure that only one instance of Genesys Info Mart Server accesses the Info Mart database at the same time, do not add a connection to the DAP that enables access to the Info Mart database (the Info Mart DAP) on the `Connections` tab of the second Genesys Info Mart `Application` object. If the first instance of Genesys Info Mart Server is unavailable, you must move the connection to the Info Mart DAP from the first Genesys Info Mart Server instance to the second instance, then start the second instance.

Note: In Oracle deployments, you can configure the connection to the Info Mart DAP in both instances of the Genesys Info Mart Server. A safeguard mechanism prevents more than one instance of Genesys Info Mart Server from accidentally accessing a single Info Mart database.

To further ensure that only one instance of Genesys Info Mart Server is active at any one time, do not select the `Auto-Restart` check box on the `Start Info` tab of both Genesys Info Mart `Application` objects.

Active-Active Info Marts

You can deploy two full instances of Genesys Info Mart to operate in parallel. Both instances of the Genesys Info Mart Server extract data from the same IDBs and independently populate their respective Info Mart databases.

This topology can be used to provide both disaster-recovery protection and protection against the Info Mart database being unavailable for an extended period of time because of, for example, network connectivity problems between the other Genesys Info Mart Server and Info Mart database hosts; for disaster recovery, the second Genesys Info Mart, at a different site, protects against the Info Mart database being lost because of a site failure.

For details about the architecture, deployment, and operation of Genesys Info Mart in the active-active configuration, see the 8.1.4 version of the *Genesys*

Info Mart Business Continuity Deployment Guide, which is provided as wiki pages.

Note: Although the *Genesys Info Mart Business Continuity Deployment Guide* provides information in the context of Business Continuity and disaster recovery, the active-active configuration is suitable for any deployment in which you want to provide database redundancy (for example, to simplify maintenance or migration), not only for disaster recovery. Furthermore, although the information is not available in versions of the document earlier than 8.1.4, any 8.1.x release of Genesys Info Mart supports the active-active configuration.

Active-Standby Info Marts

With one instance of Genesys Info Mart active, you can deploy a second instance of the Genesys Info Mart Server as well as a second instance of the Info Mart database, to which you replicate data.

In a Business Continuity environment, you can use Oracle GoldenGate to replicate the active Info Mart database (the source) to the second Info Mart database (the target) in a separate data center, to enable disaster recovery in the event of a site failure.

The active-standby configuration with Oracle GoldenGate database replication requires Genesys Info Mart 8.1.2 or later. For full information about the minimum version of Oracle GoldenGate that is required, the Info Mart tables that you should exclude from replication, and the procedures for setting up replication and for disaster recovery, see the *Genesys Info Mart Business Continuity Deployment Guide*.

Info Mart Database Backup

Whether or not you implement database redundancy, Genesys recommends that you observe usual best practices for database backup and recovery, as applicable for your RDBMS.

Topology Diagrams

Provided that the requirements for storage of ICON details are observed (see page 70), the Genesys Info Mart architecture is flexible and scalable. For example, supported topologies can include:

- Single- or multi-site deployments, with one T-Server per site or several T-Servers per site.

- For the deployment as a whole or on each site, a single media type or data domain (for example, voice only) or a combination (for example, voice and multimedia; voice and Outbound Contact; or voice and Outbound Contact and multimedia).
- Common components located at one of the data-source sites or at some other, central location. Common components include the Configuration Server and the Genesys Info Mart application and Info Mart database.

This section provides diagrams that illustrate the data-source topologies that Genesys Info Mart supports. The supported topologies implement the rules that are described in "Genesys Info Mart Requirements for ICON Details Storage" on page 70.

Given the range of potential combinations, the diagrams in this section are not intended to represent specific deployment architectures. Instead, the diagrams illustrate generic building blocks for topologies that Genesys Info Mart supports. This section includes the following diagrams:

- Conceptual Architecture—Illustrates the data-source architecture at the highest level. See Figure 2 on page 76 for the basic non-HA architecture and Figure 3 on page 77 for the basic HA architecture.
- One Data Source per ICON—Illustrates topology building blocks based on a one-to-one relationship between the data source and ICON. See Figure 4 on page 79.
- Multiple Data Sources per ICON—Illustrates topology building blocks based on a many-to-one relationship between the data sources for a particular data domain and ICON. See Figure 5 on page 81.
- Mixed Example—Multi-Site, All Details—Depicts a specific multi-site topology for all details, as an example of how you can mix and match architectural approaches. See Figure 6 on page 83.

For related information about distributing the applications among hosts, see "Recommendations on Hosting" on page 83.

Diagram Conventions

The topology diagrams use the following conventions:

- The diagrams do not show the DB Server and DAP that each ICON requires in order to access the IDB it populates. In all cases, the role that is configured for the DAP matches the role that is configured for the ICON application.

> **Note:** You can reuse the Interaction Concentrator DAP to enable Genesys Info Mart to access the same IDB. For more information, see "Reusing DAPs" on page 208.

- Square brackets ([ ]) indicate optional, additional data sources and Interaction Concentrator instances that you can include to scale your deployment.
- The diagrams show only one instance of the Genesys Info Mart application and Info Mart database. However, you can deploy a second, parallel instance of the Genesys Info Mart Server to act as a standby. For more information, see "Standby and Disaster Recovery" on page 72.

Conceptual Architecture

Figure 2 and Figure 3 illustrate the data-source architecture at the highest level. You can extrapolate from the basic architectural concepts to scale or customize your deployment topology as required.

Basic Architecture

Figure 2 illustrates the basic architecture for each data domain.

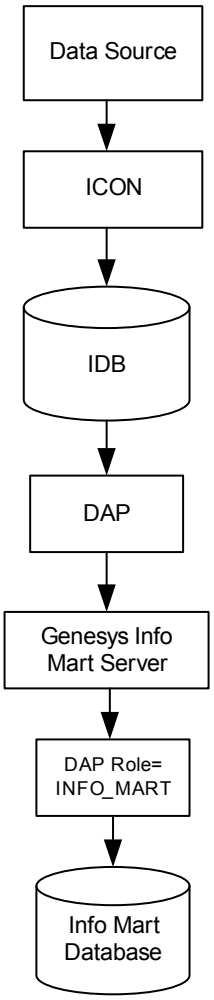

Figure 2: Conceptual Data-source Architecture

Basic HA Architecture

Figure 3 illustrates the basic architecture for high availability (HA).

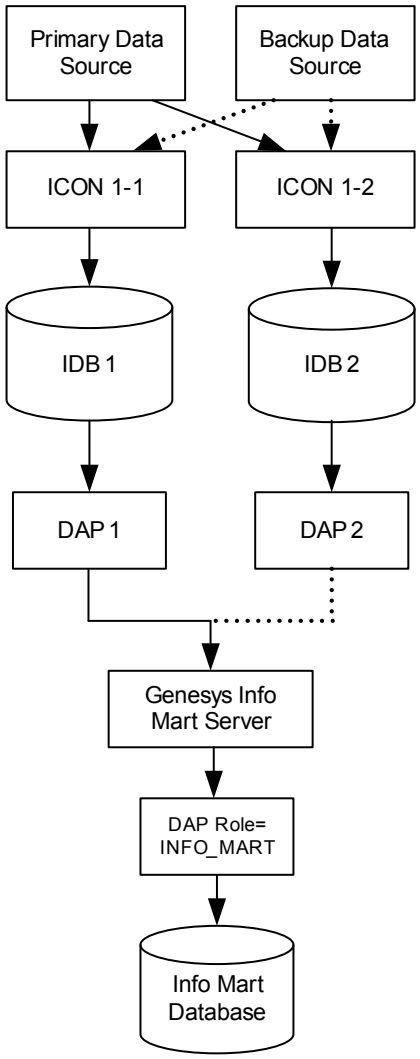

Figure 3: Conceptual HA Architecture

Note: Genesys Info Mart provides HA of reporting data by supporting HA at the Interaction Concentrator level. Genesys Info Mart itself cannot be configured to operate in HA mode.

In each HA set of redundant ICONs and IDBs, the ICON applications must be configured to perform the same role, and they must have configured connections to all the primary data sources in the HA set.

The Genesys Info Mart extraction job uses session information in the IDBs to identify which instance of IDB from the HA set of IDBs contains the most

complete and accurate set of data from each data source. For more information about Genesys Info Mart support for HA, see Chapter 20 on page 411.

One Data Source per ICON

Figure 4 on page 79 illustrates a topology in which each ICON monitors a single data source. This topology "building block" is supported for all data domains.

In this topology, the deployment consists of:

- A single Interaction Concentrator instance to store Configuration details.

  You can co-locate the Configuration details IDB in the same RDBMS instance with Configuration Database (see "Recommendations on Hosting" on page 83).
- At least one Interaction Concentrator instance to store data from either a T-Server (for Voice details) or an Interaction Server (for Multimedia details).
- Any number of additional Interaction Concentrator instances for the Voice, Multimedia, or Outbound Contact data domains, with each ICON storing data from a single data source.

Note: When this topology is deployed in a multi-site environment, temporary outages in network connectivity are less likely to result in a loss of data as long as the ICON application resides on the same site as its data source.

For HA support, provide redundant sets of Interaction Concentrator instances for each Interaction Concentrator in the non-HA deployment, as shown in Figure 3 on page 77.

Legend In Figure 4 on page 79:

- Data Source 1 represents either a T-Server (for Voice details) or an Interaction Server (for Multimedia details).
- Data Source *N* represents any number of additional, optional data sources (T-Server, Interaction Server, or OCS).
- DAP roles for the IDBs depend on the data domain:
  - DAP Role 1 is either `ICON_CORE` or `ICON_MM`.
  - DAP Role *N* is `ICON_CORE`, `ICON_MM`, or `ICON_OCS`.

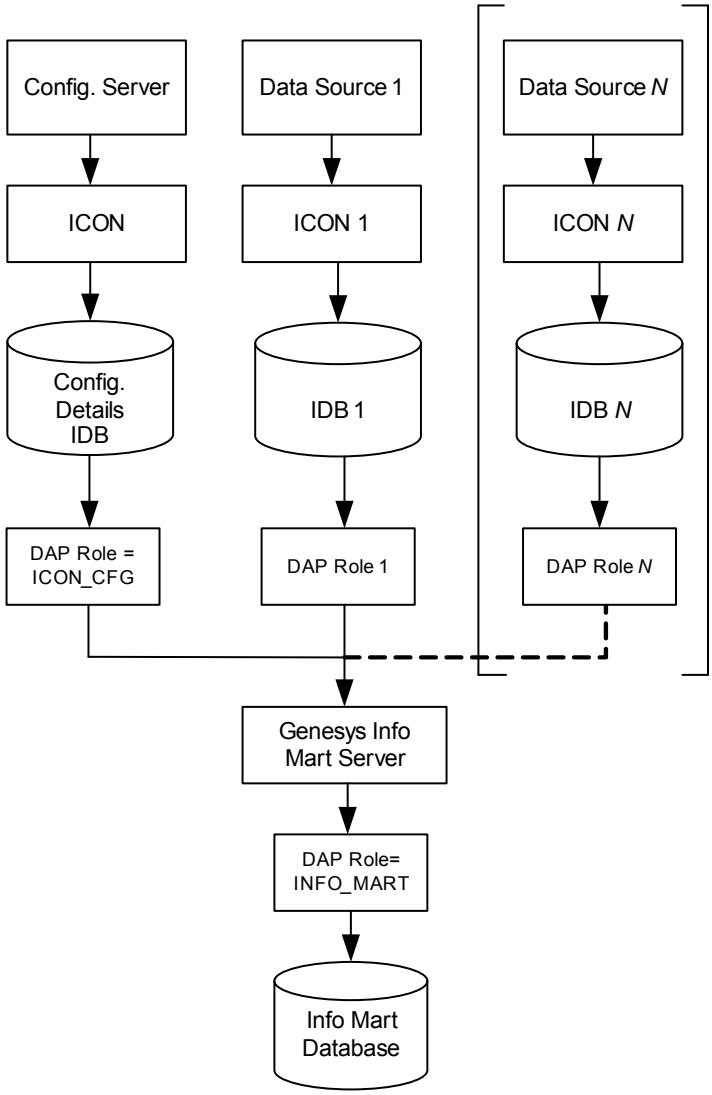

Figure 4: One Data Source per ICON

Multiple Data Sources per ICON

Figure 5 on page 81 illustrates a topology in which each ICON monitors multiple data sources for a data domain. This topology "building block" is supported for the Voice, Multimedia, and Outbound Contact data domains.

In this topology, the deployment consists of:

- A single Interaction Concentrator instance to store Configuration details.

  You can co-locate the Configuration details IDB in the same RDBMS instance with Configuration Database (see "Recommendations on Hosting" on page 83).
- At least one Interaction Concentrator instance to store data from multiple T-Servers (for Voice details) or Interaction Servers (for Multimedia details).

- Any number of additional Interaction Concentrator instances for the Voice, Multimedia, or Outbound Contact data domains, with each ICON storing data from multiple data sources of the same type.

Note: In a multi-site environment, this topology is susceptible to data delays when temporary outages in network connectivity affect connections between sites. If the ICON instance is not on the same site as one or more of the data sources, this topology might also be susceptible to data loss in a non-HA deployment, if temporary outages in network connectivity affect the connection between a data source and ICON.

For more information about identifying and mitigating data availability issues, see "Data-Source Availability" on page 403.

For HA support, provide redundant sets of Interaction Concentrator instances for each Interaction Concentrator in the non-HA deployment, as shown in Figure 3 on page 77.

Legend In Figure 5 on page 81:

- Data Sources 1-1 through 1-*N* represent a set of *N* T-Servers (for Voice details) or Interaction Servers (for Multimedia details).
- Data Sources N-1 through N-*N* represent any number of additional, optional sets of data sources, provided that each set consists of data sources of the same type (T-Server, Interaction Server, or OCS).
- DAP roles for the IDBs depend on the data domain:
  - DAP Role 1 is either `ICON_CORE` or `ICON_MM`.
  - DAP Role N is `ICON_CORE`, `ICON_MM`, or `ICON_OCS`.

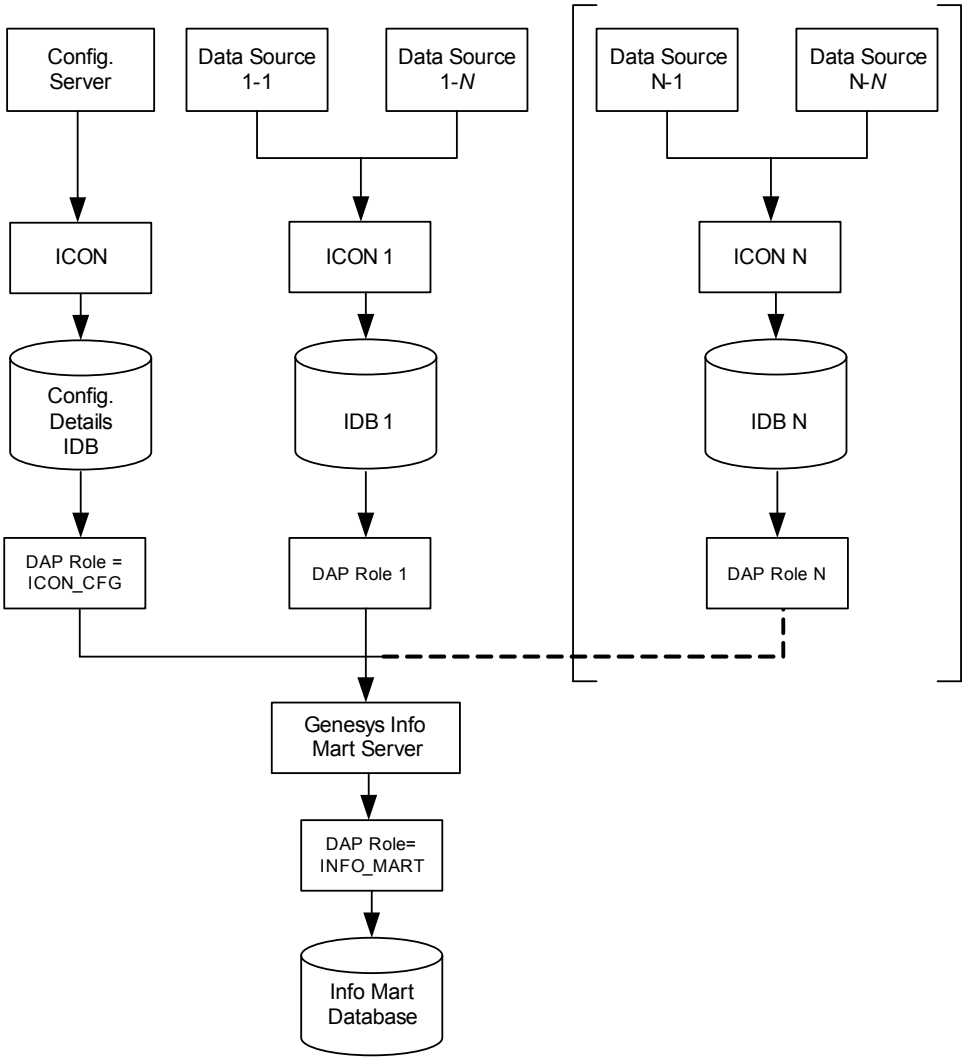

Figure 5: Multiple Data Sources per ICON

Mixed Example—Multi-Site, All Details

Figure 6 on page 83 illustrates a specific multi-site deployment, as an example of how you can combine the various approaches that are described in the preceding diagrams. In this example:

- HA is provided for Configuration details and Voice details, but not for the other data domains.
- The T-Server on Site 1 and the T-Server on Site 2 are monitored by a single Interaction Concentrator instance (on Site 1).
- Two Interaction Servers, both on Site 2, are each monitored by a separate Interaction Concentrator instance (on Site 2).

- Two OCS data sources, one on Site 1 and one on Site 2, are each monitored by a separate Interaction Concentrator instance.

> **Note:** There is nothing special about the Multimedia or Outbound Contact topology requirements. Figure 6 shows one Interaction Server and one OCS per ICON simply to illustrate the possible topology "building blocks" in a multi-site context. However, as with T-Server, multiple Interaction Server or OCS data sources can be monitored by a single ICON.

The example illustrates the following features:

- The deployment complies with the minimum requirements for supported topologies.
- Both sites host a combination of media types.
- In some cases, multiple data sources of the same type populate a single IDB; in other cases, multiple data sources of the same type each populate their own IDB.

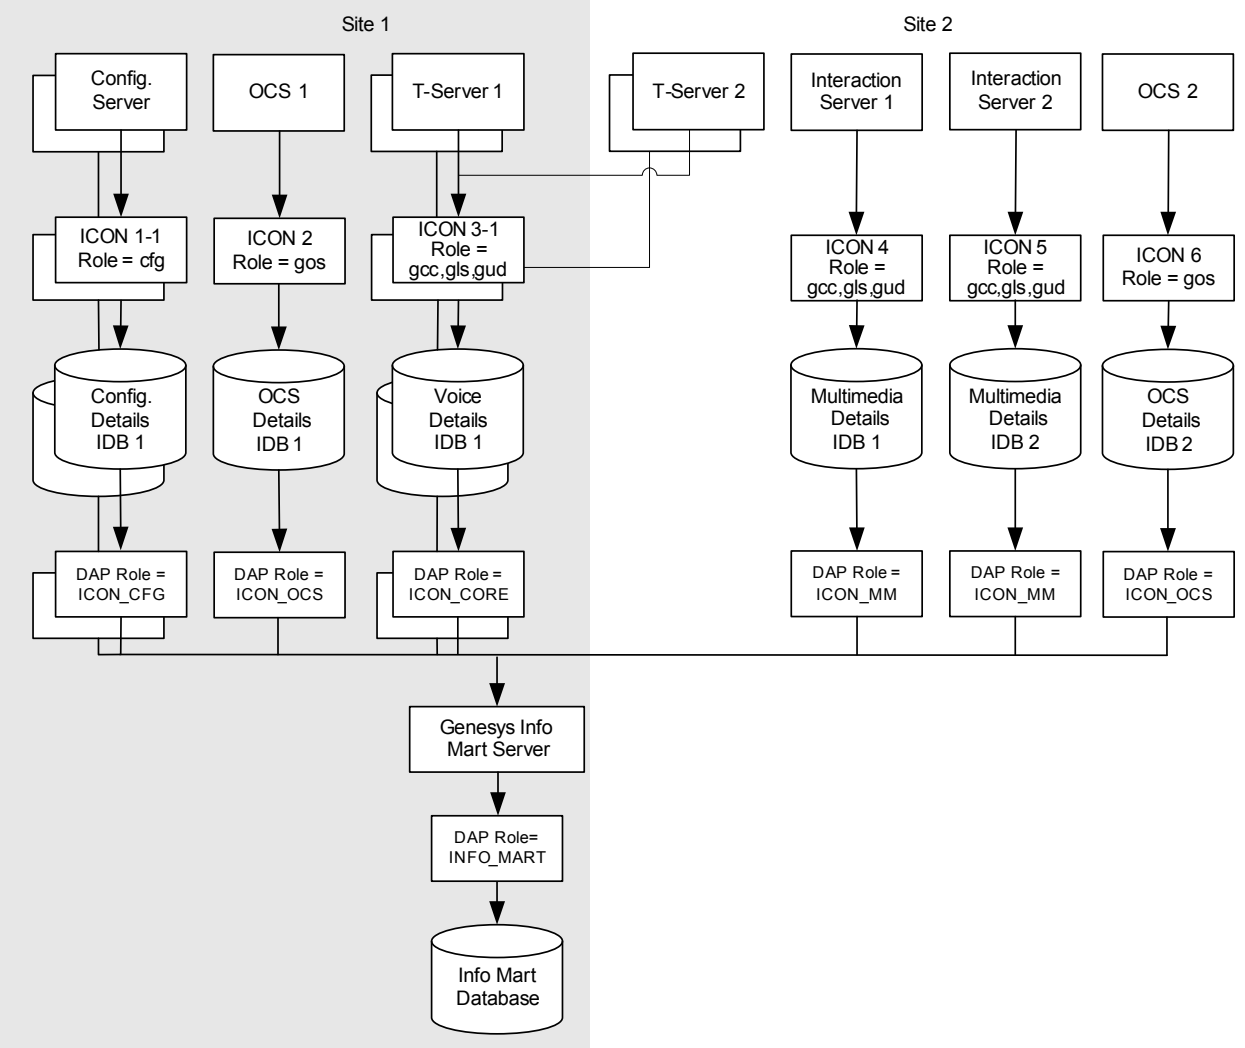

Figure 6: Mixed Topologies Example—Multi-Site, All Details

Recommendations on Hosting

Genesys Info Mart has a strong dependency on the availability of the contact center configuration information that is stored in the Configuration details IDB. To maximize the availability of the data, Genesys strongly recommends that you:

- Use a separate ICON application instance to store the contact center configuration history.
- Install this ICON application on the same host as the Genesys Configuration Server that provides the source event stream.

- Configure this ICON application to share the same Genesys DB Server that the Genesys Configuration Server uses. If this is not possible, install the DB Server that ICON will use on the same host as the Configuration Server's DB Server.
- Create the IDB in the same database server instance as the Configuration Server database.
- Use the ICON feature to resynchronize configuration data as soon as data inconsistency is suspected between the Configuration Database and IDB.

Genesys also recommends that you consider HA architecture for configuration data. For more information, see "HA for Configuration Details" on page 415.

To achieve a high level of data availability, Genesys recommends that you co-locate certain components on the same host computers.

Table 2 summarizes the hosting recommendations for a sample single-site deployment that provides HA reporting of Configuration details and Voice or Multimedia details, as well as non-HA reporting of Outbound Contact details. Extrapolate from these hosting recommendations to scale the topology for a multi-site deployment or for additional data sources and data domains.

Table 2: Hosting Recommendations for Configuration and Voice/Multimedia Details (HA) and Outbound Contact Details (Non-HA)

Host	Applications and Databases	Comment
A	• Primary Configuration Server • ICON-1A	
B	• Backup Configuration Server • ICON-1B	Host B provides backup support, in case Host A fails.

Table 2: Hosting Recommendations for Configuration and Voice/Multimedia Details (HA) and Outbound Contact Details (Non-HA) (Continued)

Host	Applications and Databases	Comment
C	• Configuration Database • DB Server that provides access to Configuration Database • IDB-1A (the IDB that is populated by ICON-1A)	Genesys recommends that you co-locate the Configuration details IDB and its data source. A failure of Host C results in Configuration Server being unable to write to Configuration Database. Running on a separate host, Configuration Server can continue to write configuration changes to its history log while Host C is being restored; the maximum amount of configuration changes that Configuration Server can write to its history log is configurable. Genesys recommends that you set the `max-record` configuration option in Configuration Server to the maximum valid value. Genesys recommends that the DB Server that provides access to IDB be located on the same host as IDB. To prevent the situation in which Configuration Server continues to write to Configuration Database while ICON cannot write to IDB, you might consider using the same DB Server to access both databases. On the other hand, to avoid a heavy load on the DB Server that provides access to Configuration Database, you might consider deploying a separate DB Server instance on either this or another host. **Note:** You can deploy an HA pair of DB Servers to access Configuration Database. If you do, Genesys recommends that you install the primary DB Server on Host C. You can install the backup DB Server either in a separate directory on Host C or on a different computer, such as Host B.
D	• DB Server that provides access to IDB-1B • IDB-1B (the IDB that is populated by ICON-1B)	Host D provides backup support in case the IDB on Host C is no longer being populated.
E	• Primary T-Server or Interaction Server • ICON-2A	As a general recommendation, select a network location for an ICON server to be the same host or to be close to the T-Server or Interaction Server host. If ICON is located away from its data source, the connection between the two servers is more likely to break. A loss of connection results in missing notifications about interaction or agent activity; this data cannot be restored.
F	• Backup T-Server or Interaction Server • ICON-2B	Host F provides backup support, in case Host E fails.

Table 2: Hosting Recommendations for Configuration and Voice/Multimedia Details (HA) and Outbound Contact Details (Non-HA) (Continued)

Host	Applications and Databases	Comment
G	• DB Server that provides access to IDB-2A • IDB-2A (the IDB that is populated by ICON-2A)	Genesys recommends that the DB Server that provides access to IDB be located on the same host as IDB. If ICON is located away from IDB, your configuration must account for data latency. A loss of connection between ICON and its IDB does not necessarily result in a loss of data, because ICON continues to write data to the persistent storage until the database connection is restored.
H	• DB Server that provides access to IDB-2B • IDB-2B (the IDB that is populated by ICON-2B)	Host H provides backup support, in case the IDB on Host G is no longer being populated.
I	• OCS • ICON-3	As a general recommendation, select a network location for an ICON server to be the same host or to be close to the OCS host. If ICON is located away from its data source, the connection between the two servers is more likely to break. A loss of connection results in missing notifications about Outbound Contact activity; this data cannot be restored.
J	• DB Server that provides access to IDB-3 • IDB-3 (the IDB populated by ICON-3)	Genesys recommends that the DB Server that provides access to IDB be located on the same host as IDB. If ICON is located away from IDB, your configuration must account for data latency. A loss of connection between ICON and its IDB does not necessarily result in a loss of data, because ICON continues to write data to the persistent storage until the database connection is restored.
K	• Genesys Info Mart Server application	Host K presents a single point of failure. Failure of the Genesys Info Mart Server does not necessarily result in data loss, because the data might still be in the IDB(s); however, failure of the Genesys Info Mart Server might delay data availability.

Table 2: Hosting Recommendations for Configuration and Voice/Multimedia Details (HA) and Outbound Contact Details (Non-HA) (Continued)

Host	Applications and Databases	Comment
L	• Info Mart database	Host L presents a single point of failure. Failure of the Info Mart database might result in data delays, but it does not usually result in data loss, provided that you have implemented adequate database management strategies to protect data. Genesys recommends that you perform frequent database backups or use failover strategies, such as clustering or mirroring to minimize delays and data loss. For example, if the Info Mart database fails, but the information in IDBs remains intact, Genesys Info Mart will process the IDB data later, when the Info Mart database is restored from backup.

Chapter

4 Database Considerations

This chapter describes database issues that you must consider before you deploy Genesys Info Mart. It contains the following sections:

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

Databases in Your Genesys Info Mart Deployment

Genesys Info Mart ships predefined jobs that run on the Genesys Info Mart Server to extract, transform, and load (ETL) data. These ETL jobs access several databases, as described in "Genesys Info Mart Jobs" on page 34. There

are also several supported database schemas (Genesys Info Mart and tenant user schemas), as described in "Genesys Info Mart Database" on page 37.

Notes:

- When you install Genesys Info Mart, you select a single RDBMS type: Microsoft SQL Server, Oracle, or, starting with release 8.1.3, PostgreSQL. All source databases (one or more IDBs) and the target database (Info Mart) must reside in databases of this same RDBMS type. The only possible exception is the Interaction Database (IDB) from which Genesys Info Mart extracts configuration history.
- Genesys Info Mart does not support database compression, except for compression that Reporting and Analytics Aggregates (RAA) may implement. Do not remove any compression that has been defined by RAA.
- Be sure to consult the Genesys Info Mart Release Advisory for important information about known RDBMS issues and potential solutions to those issues.

Database Capacity

As discussed in Chapter 1 on page 31, Genesys Info Mart 8.x reads from and writes to the Info Mart database. To determine the database capacity requirements for the Info Mart database in your environment, answer the following questions:

- How much space does the database require?
- How much space is needed for future growth?
- How powerful should the database server be?
- How do you plan to use the Genesys Info Mart database?
- How do you plan to manage the Genesys Info Mart indexes?
- What will your purging strategy be?

Use the following information to help answer these questions.

Storage Capacity

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

If your deployment includes Genesys Interactive Insights (GI2) or the RAA package, you will also need to provide storage for the aggregate tables and

indexes that you create in order to improve query performance, and to provide storage for an extended period of time that suits your requirements. See the RAA documentation set (described on page 471) for more information.

Genesys Info Mart provides a maintenance job that purges data in the Info Mart database. The maintenance job automatically purges data in accordance with configurable data-retention policies. For more information, see "Purging Info Mart Data" on page 405.

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

Genesys provides an interactive tool, the *Genesys Info Mart 8.1 Database Size Estimator,* to help you estimate the size of your Info Mart database.

Processing Capacity

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

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

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

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

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

Database Partitioning

Genesys Info Mart 8.x supports the use of partitioning in the Info Mart database in Oracle (range partitioning only), Microsoft SQL Server, and PostgreSQL deployments. In general, if partitioning is used, fact tables and associated indexes in GIDB and the dimensional model are partitioned.

Starting with release 8.1.2, certain Control tables are partitioned as well. Configuration object tables and configuration relationship fact tables are not partitioned, and all dimension tables are also not partitioned.

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

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

Configuration options enable you to specify the size of the partitions in GIDB (see `partitioning-interval-size-gidb` on page 298) as well as the dimensional model and Control tables (see `partitioning-interval-size-gim` on page 298). Another configuration option, `partitioning-ahead-range` (page 297), enables you to control how far ahead the Genesys Info Mart jobs will create partitions, in preparation for future ETL cycles. (`Job_InitializeGIM` creates the partitions in the first instance, then `Job_MaintainGIM` creates them on an ongoing basis.)

Note: By default, Genesys Info Mart runs the maintenance job daily. If your database is partitioned, ensure that you do not jeopardize routine maintenance of the partitions by inappropriately changing the configuration options that control scheduling of the ETL cycle and the maintenance job.

For more information about the scheduling-related configuration options, see "schedule Section" on page 314.

Source Data Retention and Purging

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

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

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

Generally, you should aim to achieve a balance that enables you to retain enough data in your Interaction Concentrator databases without affecting either

the operating performance of your source database or the extraction process of Genesys Info Mart.

For Voice details and Outbound Contact details, Genesys recommends that you consider using IDB partitioning, which is supported for Oracle deployments, starting from Interaction Concentrator release 8.1.1. Partitioning improves overall IDB performance and streamlines maintenance. However, given the way that partitioning support has been implemented, with only a limited number of partitions expected to be used, and given the long-living nature of multimedia interactions, Genesys recommends that you not use partitioned IDBs for Multimedia details. For more information about IDB partitioning, see the *Interaction Concentrator 8.1 Deployment Guide.*

Genesys provides specific recommendations regarding source data retention and purging frequency. For more information, see the chapter about maintaining Interaction Concentrator and data sources in the *Genesys Info Mart 8.1 Operations Guide*.

Database Object Owners and User IDs

Genesys Info Mart requires access to source and target databases to perform a variety of operations. Logically, there are three types of users for which you must provision user accounts:

- Info Mart—The user's account that is used by Genesys Info Mart jobs to access the IDB schemas and the Info Mart database schema for the purposes of:
  - Creating and modifying the Info Mart database schema
  - Modifying the IDB schemas to work with Genesys Info Mart
  - Extracting data from IDB
  - Processing the extracted data and storing it in the Info Mart database schema
  - Maintaining the Info Mart database schema
- Tenant Admin—The user's account that is used by a Tenant Administrator to access the Info Mart database schema and Tenant User schema(s) for the purposes of:
  - Creating and dropping the Tenant User schema(s) and views
  - Granting privileges to the read-only tenant users to access tenant-specific views in the Info Mart database schema through the Tenant User schema(s)
- Tenant User—The user's account that is used by a read-only tenant user to view tenant-specific data in the Info Mart database schema through views in the Tenant User schema.

Each `User ID` must have the necessary privileges to perform the required operations against the applicable database. For information about the required privileges, see "Database Privileges" on page 94.

- In general, the `User ID` that you specify to connect to each database does not need to be the same as the `Owner ID`. You can create the database objects in a schema that is different from the `User ID` that accesses them, provided that the `User ID` has the required privileges, with the following exception: In PostgreSQL deployments, starting with release 8.1.4, the name of the Tenant User schema must be the same as the name of the corresponding tenant `User_ID`.

Note: In Microsoft SQL Server, all database objects are contained in *schemas,* instead of being owned by a database *owner.* The SQL Server logins are mapped to database users, who can own objects in the various schemas. In addition, a *default schema* is configured for each database user, to contain unqualified database objects.

In this document, for Microsoft SQL Server, the term *Owner ID* refers to the database user who owns the object—for example, the database user who created a view.

Identifying the owner and user accounts that you will use in your Genesys deployment is an important step in your deployment and installation planning. After you have identified the accounts that you will use, record the IDs and passwords on the applicable worksheet that is provided in "Installation Worksheets" on page 421. You will need this information to create the database schemas and to specify connection parameters when you configure the database access points (DAPs).

Database Privileges

Table 3 summarizes the user account privileges that are required for the respective source and target database schemas.

Note: Genesys recommends that you explicitly grant the required privileges to the respective users. Do not assume that granting DBA privileges will be sufficient to enable the user account to perform all required operations.

Table 3: Required User Account Privileges

User Account	Database	Required Privileges[a]
Info Mart	IDB	• SELECT on all tables and views • CREATE and DROP tables, views, and indexes • INSERT, UPDATE, and DELETE on all IDB tables
	Info Mart schema	• CREATE, DROP, and ALTER tables, views, synonyms,[b] indexes, procedures, and, if applicable, partitions • SELECT on all tables and views • INSERT, UPDATE, and DELETE on all tables • RDBMS-specific privileges that are required to truncate tables • CREATE and EXECUTE on all stored procedures • (For Oracle only) EXECUTE on DBMS_LOCK • (For databases that implement partitioning) RDBMS-specific privileges that are required to CREATE, DROP, and ALTER tables, indexes, and partitions
Tenant Admin	Info Mart schema	• EXECUTE on all stored procedures • CREATE, SELECT, and INSERT on all tables • CREATE, DROP, and SELECT on all views and sequences • GRANT access to views for Tenant users
	Tenant User schema(s)	• CREATE and DROP views
Tenant User[c]	Info Mart schema	• SELECT on the tenant-specific view in the applicable Tenant User schema **Note:** Privileges for the Tenant User to access views in the Info Mart schema are granted by the Tenant Admin when the make_gim_view_for_tenant.sql script is run.
	Tenant User schema	• SELECT on all views

a. Privileges are called *permissions* in Microsoft SQL Server.

b. Synonyms are created only for Oracle databases.

c. Applicable only for read-only views, which are strictly required only for multi-tenant deployments.

Database User Authentication

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

Chapter

5 Multimedia Interactions

In this document, *multimedia interactions* refers to all interactions for which the data source for Genesys Info Mart is Interaction Server. Multimedia interactions include Genesys eServices/Multimedia e-mail and chat interactions, as well as 3rd Party Media interactions (formerly referred to as *Open Media*). For a more detailed definition of the term, see "Multimedia Interactions" on page 40.

The most important difference between voice and multimedia interactions is that voice interactions are short-lived and multimedia interactions can be long-lived. Therefore, multimedia interactions might be transformed while they are still active.

Genesys Info Mart processes data that is related to all multimedia interactions in a similar manner. However, from the point of view of internal processing, 3rd Party Media interactions require separate handling. This chapter describes special considerations for 3rd Party Media interactions, as well as considerations for configuring your Genesys Info Mart deployment.

This chapter contains the following sections:

Media Types

Genesys Info Mart can support any media type that is properly configured as a `Media Type Business Attribute` in the Configuration Layer.

Genesys Info Mart stores the media types that are available to describe interaction facts in the `MEDIA_TYPE` dimension table in the Info Mart database. By default, the `MEDIA_TYPE` table includes the Genesys eServices/Multimedia

e-mail and chat media types. It can be extended to include any number of deployment-specific 3rd Party Media media types.

Dynamic Addition of Unknown Media Types

When Genesys Info Mart is transforming agent activity data or any interaction data—including 3rd Party Media interaction data—and a media type arrives that is new to Genesys Info Mart, the transformation job automatically adds the new media type to the `MEDIA_TYPE` dimension table and includes it when transforming data. Genesys Info Mart generates a log event to alert you about the addition.

Note: The new media type must have been configured already in the Configuration Layer. If not, Interaction Server rejects the media type before it reaches Genesys Info Mart.

Online and Offline Interactions

Genesys Info Mart distinguishes between two types of interactions:

- Online—The interaction involves an online session with a customer in real time (for example, `chat`).
- Offline—The interaction does not involve an online session with a customer in real time (for example, `email`).

Genesys Info Mart uses the `IS_ONLINE` column in the `MEDIA_TYPE` table to indicate whether a particular media type is associated with online interactions (`IS_ONLINE=1`) or with offline interactions (`IS_ONLINE=0`).

Downstream reporting applications can use the `IS_ONLINE` flag to produce reports that distinguish between online and offline interactions.

For internal processing, Genesys Info Mart uses the `IS_ONLINE` flag to determine whether to process the interaction as online or offline. The distinction affects:

- Detection of when an interaction is abandoned—An offline interaction cannot be abandoned by the customer.
- Detection of when an interaction is serviced—An online interaction is serviced by an agent when the agent joins the interaction (and the session). An offline interaction has some level of service when an agent accepts and begins to work on the interaction, but it is not fully serviced until the agent sends a reply back to the customer.
- The setting of the Technical Descriptor for transfers to a queue.
- Determination of `ONLINE_DURATION`.

Configuration Considerations for 3rd Party Media

When Genesys Info Mart dynamically adds an unknown media type to the `MEDIA_TYPE` table (see "Dynamic Addition of Unknown Media Types"), the new media type is added as an offline media type (`IS_ONLINE=0`). If you want

Genesys Info Mart to identify interactions that are associated with this media type as online interactions, you must manually change the setting of the IS_ONLINE flag in the MEDIA_TYPE table. You can change the setting at any time. The change takes effect as soon as it is committed in the database, but it does not apply to data that has already been processed.

Genesys recommends that you set an alarm on the log event that Genesys Info Mart generates when the media type is added (message number 55-20150), to prompt you to verify this flag and, if required, change the setting to suit your deployment and reporting requirements.

Configuration Recommendation for Online Media Types

To ensure that online interactions that are associated with 3rd Party Media are processed consistently, Genesys recommends that you manually add the online media types to the MEDIA_TYPE table before Genesys Info Mart starts processing, so that you can set IS_ONLINE=1 from the start.

For more information about manually adding online media types or changing the IS_ONLINE flag, see Procedure: Setting up media types for online interactions, on page 344.

Interaction Types and Subtypes

Genesys Info Mart can support any multimedia interaction subtype that is properly configured in the Configuration Layer.

When a multimedia interaction subtype arrives that is new to Genesys Info Mart, the transformation job adds the new interaction subtype automatically to the INTERACTION_TYPE dimension table and includes it when transforming data. Genesys Info Mart generates a log event to alert you about the addition.

Note: The new interaction subtype must have been configured already in the Configuration Layer. If not, Interaction Server rejects the interaction type before it reaches Genesys Info Mart.

Genesys Info Mart uses the IGNORE column in the INTERACTION_TYPE table to determine whether to include interactions of a particular subtype in transformation processing:

- 0 (false)—Genesys Info Mart transforms all interactions with this interaction type.
- 1 (true)—Genesys Info Mart does not transform any interactions with this interaction type. No records are generated in fact tables for interactions of this type.

By default, all newly added multimedia interaction subtypes are set to be transformed. You can disable transformation of multimedia interactions of a

specified subtype by setting a value of 1 (true) for the IGNORE field in the INTERACTION_TYPE table.

Note: If you disable a subtype, both the parent interactions of that subtype and any child interactions of such parent interactions are disabled, even if the child interactions themselves are of a different subtype, one that is configured to be transformed.

Genesys recommends that you set an alarm on the log event that Genesys Info Mart generates when the interaction subtype is added (message number 55-20151), to prompt you to verify the IGNORE flag and, if required, change the setting to suit your deployment and reporting requirements.

You can change the setting at any time. The change takes effect as soon as it is committed in the database, but it does not apply to data that has already been processed.

Multi-Tenant Considerations

The IGNORE setting for an interaction type or subtype and the IS_ONLINE setting for a media type apply across the deployment, without regard to tenant. In other words, Genesys Info Mart does not support ignoring a particular interaction type for one tenant, but processing it for another. Similarly, Genesys Info Mart does not support transforming interactions that are associated with a particular media type as online for one tenant and offline for another.

Archive Queues

Your multimedia deployment might include business processes that archive completed interactions in queues before eventually terminating the interactions in Interaction Server. For example, the Genesys intelligent Workload Distribution (iWD) solution employs this approach.

Issues for Reporting

Archive queues create issues with performance and data quality in regular reporting. Without special handling, placement of an interaction in an archive queue is treated like a transfer to any other queue. This results in misleading queue- and interaction-duration metrics, as well as significant performance penalties in transformation and, if applicable, aggregation when extremely long-living interactions are eventually removed from the archive queue or terminated in Genesys Info Mart in accordance with the days-to-keep-active-facts configuration option (see page 281).

Configuring Genesys Info Mart for iWD and Archive Queues

Starting with release 8.1.3, Genesys Info Mart enables you to specify which `Script` objects of type `Interaction Queue` are archive queues that require special handling.

By default, Genesys Info Mart identifies the queues named `iWD_Completed` and `iWD_Canceled` as archive queues for completed and canceled interactions, respectively. Genesys Info Mart considers interactions that are placed into these queues to be completed and terminates the interactions, with appropriate technical descriptors in the `INTERACTION_RESOURCE_FACT` (IRF) record of the handling resource that placed the interaction in the queue.

If you want to identify additional or alternative queues as archive queues for special handling, adjust the settings of the `completed-queues` (see page 310) or `canceled-queues` (see page 309) options, respectively, as applicable.

Limitation

Genesys Info Mart permanently terminates an interaction that is placed in one of the archive queues identified for special handling. If an agent or strategy subsequently restarts the interaction from the queue or updates the interaction in the queue (for example, by adding or updating user data), Genesys Info Mart ignores the new activity.

Chapter

6 Attached Data

This chapter describes how Genesys Info Mart processes attached user data. It contains the following sections:

Genesys Info Mart and Attached Data

Genesys Info Mart uses attached data key-value pairs (KVPs) to populate several of its fact and dimension tables, including any number of custom-created tables. Interactive Voice Response (IVR) applications, Enterprise Routing Solution, Network Routing Solution, Outbound Contact solution, Genesys eServices/Multimedia solution, and Agent Desktop applications all attach KVPs to interactions. The KVPs that these applications attach depend on the following factors:

- Your deployment's interaction flows
- The information that is required by the resources that handle the interactions
- The information that you want to report

Genesys recommends that you configure upstream applications to attach KVPs as early in the interaction flow as possible. In this way, key interaction attributes are captured, even if the interaction is abandoned.

Note: To ensure that applications use certain attributes consistently (for example, `CustomerSegment`, `ServiceType`, and `ServiceObjective`), use Configuration Manager to configure values for them. For more information about how to configure these attributes, see *Framework 8.x Configuration Manager Help*.

Business analysts, report developers, and database administrators (DBAs) need to collaborate to plan user-data reporting. As part of deployment planning:

1. Review the high-level algorithm for processing user data in Genesys Info Mart. See "Processing User Data".

Business Analyst Responsibility

2. Evaluate what business attributes need to be stored for your contact center reporting purposes.

Report Developer Responsibilities

3. Research how to enable respective applications to attach required KVPs. See "User-Data Sources and KVPs" on page 107 and "Using UserEvent-Based KVP Data" on page 114 for information. Refer to the documentation for Genesys solutions, if necessary. Ensure that the respective KVPs have been defined for the required business attributes and that the respective applications have been configured to attach the required KVPs.
4. Research how to enable Interaction Concentrator (ICON) to store required KVPs in Interaction Database (IDB). Refer to "Customizing Your ICON Attached Data Specification File" on page 176, Appendix B on page 437 and the Interaction Concentrator documentation, as necessary.
5. Based on expected degrees of cardinality, decide which custom KVPs to store as facts and which to store as dimensions.
6. Based on business rules and reporting requirements for the values of KVPs at different stages of the interaction, decide which propagation rule to use for each KVP. For more information, see "Propagation Rules" on page 118.
7. Decide what default values to use for custom user-data facts. For information about when Genesys Info Mart uses the default values, see "Transformation" on page 106.

DBA Responsibilities

8. Decide what custom fact and dimension tables, if any, you require to store user data values. See "Storing User Data" on page 115.
9. For custom user-data facts, decide what KVP values, if any, you want to store as numeric or, starting with release 8.1.201, as date/time data (as supported by your RDBMS). Interaction Concentrator stores user data as character data, but Genesys Info Mart stores KVP values in custom

user-data fact tables as character, numeric, or date/time data types, depending on the way that you define the column in the Info Mart schema (see Step 11).

10. Map the KVPs to custom user-data tables and columns in the Info Mart database. Starting with release 8.1.2, you can use the User Data Assistant to help you perform the mapping and prepare the required SQL script. For information about the User Data Assistant, see "Enabling Storage of User Data" on page 148.

    To map the KVPs manually, fill the User Data Mapping worksheet, to use with your user data template script at the time of deployment. See "Mapping Call-Based Attached Key-Value Pairs" on page 117 for information and Appendix A on page 421 for a sample worksheet.

11. Prepare your copy of the `make_gim_UDE_template.sql` or `make_gim_UDE_template_partitioned.sql` template script to use to create custom user-data tables and columns when you deploy the database. See "Customizing the User Data Template" on page 184.

Processing User Data

There are two types of KVP data (referred to as *user data,* when discussed collectively elsewhere in this document):

- *Call-based* attached data
- *UserEvent-based* KVP data, which allows the agent to associate KVP data with a voice interaction after the voice interaction has ended (that is, after the call is released)

Genesys Info Mart uses the same, unified mechanism to process these two data types.

By content, user data can also be divided as follows:

- High-cardinality user data—Data for which there can be a very large number of possible values. Each interaction has no more than one value for each user-data type. A Customer ID number is an example of high-cardinality user data.
- Low-cardinality user data—Data that has a limited range of possible values. There might be multiple values of a specific type for a single interaction. A "new customer" flag, which has only two values—`Y` and `N`—in a respective database column is an example of a low-cardinality user data. Service type is another example of data that has a limited number of possible values.

The following high-level algorithms help you understand how Genesys Info Mart processes user data.

Extraction

1. Genesys Info Mart Server extracts user data along with other interaction data from one or more IDBs.

2. Global Interaction Database (GIDB), which is a set of tables within the Info Mart database schema, stores the extracted user data for future processing.

Transformation

3. Genesys Info Mart processes the user data and creates records in relevant user data tables (predefined or custom). Genesys Info Mart uses customer-configured mapping rules to identify in which user data tables to store certain KVP values. Typically, low-cardinality data is expected to be stored in dimension tables, while high-cardinality data is expected to be stored in fact extension tables.

   If no value is received for a customer-mapped KVP that has been included in the schema, Genesys Info Mart uses the default value that you specify in the user-data script (see Step 11 on page 105).

4. In the case of user data that you have configured Genesys Info Mart to store as date/time, Genesys Info Mart converts the KVP value to a date/time using the Genesys Info Mart default format for date/time (`yyyy-mm-ddThh24:mi:ss.ff`), unless you have specified an alternative conversion expression. For more information about customizing the date conversion, see "Customizing the User Data Template" on page 184.

5. If the attempt to convert a KVP value or to store it in a user-data table column fails, the transformation job itself does not fail. Instead, Genesys Info Mart logs a message about the failure and, for each invalid KVP value, inserts a record into the `STG_TRANSFORM_DISCARDS` table and uses the default value that you specified in the user-data script.

   In releases earlier than 8.1.3, this exception-handling behavior applied only to user-data facts. Starting with release 8.1.3, this exception-handling behavior applies to user-data dimensions as well.

6. During data transformation, Genesys Info Mart Server identifies whether the newly extracted user data should be associated with any `INTERACTION_RESOURCE_FACT` (IRF) records. Starting with release 8.1.2, and provided that the `link-msf-userdata` configuration option has been enabled on the applicable `DN` or `Script` objects (see pages 263 and 266), the Genesys Info Mart Server also identifies whether the newly extracted user data should be associated with any `MEDIATION_SEGMENT_FACT` (MSF) records.
   - If the interaction data is processed in the same cycle with the user data, Genesys Info Mart creates an association between a user data record and a newly created IRF or MSF record.
   - If the interaction data was processed in a previous cycle (that is, the user data arrived after call completion), Genesys Info Mart updates records in user-data tables that are associated with the IRF records.

Limitations for User Data in MSFs

When Genesys Info Mart release 8.1.2 or later is configured to store user data for interactions that are in mediation, user data is reported for MSFs in exactly the same way as for IRFs, except for the following limitations:

- If more than one change in user data is reported for a multimedia interaction that exited one queue and entered another queue in the same second, Genesys Info Mart cannot distinguish which change relates to which MSF. In all the MSFs for that interaction, Genesys Info Mart reports the value of the last change that occurred in that second. This value might not be correct for all the queues.
- The `PARTY` propagation rule is not suitable for certain situations, as described in the Notes on page 118.

User-Data Sources and KVPs

This section provides some guidelines about the KVPs that contact centers typically use for reporting purposes. KVPs are listed by the Genesys application that attaches them:

- IVR Application, page 108
- Universal Routing, page 110
- eServices/Multimedia-Specific Attached Data, page 112
- Outbound Contact Solution, page 113
- Agent Desktop Applications, page 113

Source Attributes in Events

For call-based attached data, KVPs can be reported in the `UserData`, `Reasons`, or `Extensions` attributes of TEvents and Interaction Server events. The `source` that is specified in the ICON attached-data specification file controls which event attribute ICON will store (for example, `source="userdata"`).

Prior to release 8.1.101.07, Genesys Info Mart always filtered the call-based attached data in IDB to extract only KVPs from the `UserData` attribute. Starting with release 8.1.101.07, the `filterUserData` startup parameter enables you to control whether Genesys Info Mart will continue to extract KVPs from only the `UserData` attribute of TEvents and Interaction Server events (`filterUserData=true`, the default behavior) or whether it will also consider KVPs from the `Reasons` and `Extensions` attributes of TEvents (`filterUserData=false`).

Turning off filtering of user data has performance implications, because it increases the amount of user data that Genesys Info Mart will have to process.

For information about setting the `filterUserData` startup parameter, see "Modifying JVM Startup Parameters" on page 360.

IVR Application

You must configure your applications to send the `IApplication` KVP—and you must configure ICON to store it—even if you do not want to store the `IApplication` KVP in Info Mart user-data tables for your reporting purposes. Genesys Info Mart uses the `IApplication` KVP value internally during transformation.

Other KVPs that your IVR applications attach depend on the following factors:

- The technologies that your IVR application supports
- Whether the applications are self-service-oriented
- Whether the applications work in conjunction with Enterprise Routing Solution

Based on these factors, you might choose to modify your IVR applications so that they attach the following KVPs:

- `IPurpose`
- `IResult`
- `IResultReason`
- `ITextToSpeech`
- `ISpeechRecognition`
- `CustomerID`
- `CaseID`
- `Revenue`
- `Satisfaction`
- `CustomerSegment`
- `ServiceType`
- `ServiceSubType`
- `Business Result`

You might also decide to attach user-defined KVPs.

Note: If the IVR DNs act as agents by logging into a queue, IVR applications can associate KVP data with a voice interaction by sending UserEvents after the voice interaction has ended (that is, after the call is released). The UserEvent has to be sent within the timeout that is specified in the Genesys Info Mart application configuration. `IPurpose` cannot be sent in UserEvents.

IPurpose IVR KVP

Genesys Info Mart uses the `IPurpose` KVP to determine whether an IVR application represents a self-service application or only a part of the mediation process:

- For a self-service IVR, Genesys Info Mart creates a separate row in the IRF table that represents the IVR activity as interaction handling (not as mediation). In other words, the IRF table is populated with facts for this self-service IVR in the same manner as for an agent.
- For a nonself-service IVR, no separate IRF row is created; the IVR activity is represented as mediation (not as interaction handling) as part of another row in the IRF table.

The presence of the `IPurpose` key with the value of `1 (Self Service)` forces Genesys Info Mart to treat an IVR as a handling resource. Otherwise, Genesys Info Mart treats the IVR as a mediation resource.

Notes:

- In an environment in which IVR applications rely on Universal Routing to select a target, you can modify your Universal Routing Server (URS) routing strategies to attach the `IPurpose` KVP on behalf of the self-service IVR application. See page 111 for more information.
- If you do not modify your self-service IVR applications or routing strategies to attach `IPurpose` KVP, you will see a high number of customer-abandoned interactions. To mitigate this, configure the ETL to treat all IVR applications as self-service.

  Do this by setting the `default-ivr-to-self-service` configuration option to `true` in the `[gim-transformation]` section; in this way, you can configure Genesys Info Mart to treat all IVR resources as self-service IVRs.
- If a self-service IVR uses a Two-Step or Mute transfer to transfer calls to an agent, configure the IVR application to set the value of the `IPurpose` key to `1` for consultation calls as well. Alternatively, set the T-Server option `consult-user-data` to `inherited` or `joint`, so that T-Server will propagate all user data, including the `IPurpose` KVP, from the original call to the consultation call.

In the following deployments, an IVR application can attach the `IPurpose` key with the value of `1 (Self Service)` to indicate to the reporting system that the corresponding IVR is a self-service resource:

- IVR In Front of the Switch—An IVR and IVR ports exist as configuration objects in the Configuration Database, and IVR ports are associated with DN objects that are configured under the IVR Server's virtual switch.

- IVR Behind the Switch—An IVR and IVR ports exist as configuration objects in the Configuration Database, and IVR ports are associated with DN objects that are configured under the premise switch.

When it arrives at your IVR port, the call is associated with a corresponding DN object in the Genesys environment; this association clearly indicates to the ETL that the call is at an IVR.

The IVR application can set the `IPurpose` key to the `Self Service` value and attach this data to the original call while the call is at the IVR port.

As a result, Genesys Info Mart creates a record in the `INTERACTION_RESOURCE_FACT` table to represent the self-service IVR application that is handling the customer interaction.

Universal Routing

The KVPs that Universal Routing attaches depend on the following factors:

- The type of routing strategies that you deploy
- Whether routing strategies work in conjunction with IVR applications

You can configure URS to attach the following strategy name and routing target KVPs automatically, by setting the URS `report_targets` configuration option to `true`:

- `RTenant`
- `RStrategyName`
- `RTargetTypeSelected`
- `RTargetObjectSelected`
- `RTargetAgentSelected`
- `RTargetPlaceSelected`

Note: By default, ICON stores values for these keys in the IDB `G_ROUTE_RESULT` table.

Your routing strategies can use the `MultiAttach` object and `FindServiceObjective` function in Interaction Routing Designer (IRD) to attach the following KVPs that represent requested skills, business attributes, and service objectives:

- `RRequestedSkillCombination`

Note: If you do not use the IRD `MultiAttach` object to define the requested skill combination, ensure that you represent the skill combination as a list of comma-separated skill names, each with an optional minimum proficiency; wordspacing between the list items is not significant. For example, the formats of the following skill combinations are valid:

- `skill1`
- `skill1=1`
- `skill1,skill2`
- `skill1, skill2`
- `skill1=1, skill2`
- `skill1,skill2=1`
- `skill1=1, skill2=2`
- `skill1=1,skill2=2`

A skill combination is not the same as a skill expression. Logical operators and comparitors—such as <, >, |, and &—are not valid.

- `CustomerSegment`
- `ServiceType`
- `ServiceObjective`

You might also decide to attach the following KVPs or user-defined KVPs:

- `CustomerID`
- `CaseID`
- `Revenue`
- `Satisfaction`
- `ServiceSubType`

Attached Data for Self-Service IVRs

When used in conjunction with self-service IVR applications, your routing strategies might also choose to attach the `IPurpose` KVP on behalf of the IVR application. (The `IPurpose` KVP that is attached by the IVR application takes priority.) The `Self Service` value for the `IPurpose` key indicates to the reporting system that the corresponding IVR is a self-service resource in the following deployments:

- IVR In Front of the Switch
- IVR Behind the Switch

IVR In Front of the Switch

In this deployment (as defined on page 109), a call also involves a routing point, which is configured as a DN of the `Routing Point` type under the IVR Server's virtual switch.

Either the IVR application or the routing strategy that is associated with the routing point (or both) can set the `IPurpose` key to the `Self Service` value. As

a result, Genesys Info Mart creates a record in the IRF table to represent the self-service IVR.

IVR Behind the Switch

In this deployment (as defined on page 110), a call might involve a routing point, which is configured as a `DN` object of the `Routing Point` type under the premise switch.

The `IPurpose` key with the `Self Service` value is set as follows, in any combination:

- The routing strategy that is associated with the routing point at the premise switch attaches the KVP before it routes the call to the IVR DN.
- The IVR application attaches the KVP while the call is at the IVR port.

As a result, Genesys Info Mart creates a record in the IRF table to represent the self-service IVR.

eServices/Multimedia-Specific Attached Data

Events from the eServices/Multimedia solution include a number of attributes that are specific to multimedia interactions. ICON stores these attributes in the `GM_F_USERDATA` and `GM_L_USERDATA` tables. By default, ICON stores the KVPs and event attributes that Genesys Info Mart requires, even if you do not explicitly specify them in the ICON attached-data specification file. Genesys Info Mart does not process custom KVPs that you configure ICON to store in the `GM_F_USERDATA` or `GM_L_USERDATA` tables.

Table 4 on page 112 describes important multimedia-specific KVPs that Genesys Info Mart processes and that ICON stores by default.

Table 4: Multimedia-Specific Interaction Attributes

Attribute	Description
Subject	The subject of the multimedia interaction.
FromAddress	The “from” address of the multimedia interaction.
InteractionSubType	The interaction subtype of the multimedia interaction. This subtype is a component of the value for the `INTERACTION_TYPE_KEY`. The `INTERACTION_TYPE` dimension includes both interaction type and subtype,
StopReason	eServices/Multimedia allows a reason name to be provided for each action. ICON records this reason name for the action that stops the interaction, identifying the reason the interaction was stopped. Genesys Info Mart uses this stop reason for internal purposes—for example, when setting the `TECHNICAL_DESCRIPTOR_KEY` in the IRF.

Outbound Contact Solution

Outbound Contact Server (OCS) automatically attaches the `GSW_CALL_ATTEMPT_GUID` call attempt ID for progressive and predictive dialing modes. For preview dialing mode, OCS provides the `GSW_CALL_ATTEMPT_GUID` KVP in the `UserEvent` with record information.

You must ensure that your desktop application attaches the `GSW_CALL_ATTEMPT_GUID` KVP. Genesys Info Mart uses it for internal processing. Downstream reporting applications can also use it to integrate contact attempt details with call details.

Agent Desktop Applications

Agent desktop applications might attach various KVPs, depending on your configuration of business attributes in Configuration Manager. For example, desktop applications can attach the following KVPs if they have not already been attached by some other application (such as IVR applications or Enterprise Routing Solution):

- `CaseID`
- `CustomerID`
- `Revenue`
- `Satisfaction`
- `Business Result`

You might also decide to attach some of the user-defined KVPs.

Note: Agent desktop applications can associate KVP data with a voice interaction by sending UserEvents after the voice interaction has ended (that is, after the call is released). The UserEvent has to be sent within the timeout that is specified in the Genesys Info Mart application configuration.

If you want to track the reasons for agents being in `NotReady` states, ensure that relevant KVPs are available to your agents through their desktop applications.

OCS automatically attaches the `GSW_CALL_ATTEMPT_GUID` call attempt ID for progressive and predictive dialing modes. For preview dialing mode, you must ensure that your desktop application attaches the `GSW_CALL_ATTEMPT_GUID` KVP to the actual interaction. The `GSW_CALL_ATTEMPT_GUID` KVP is provided by OCS in the `UserEvent` with record information. For voice interactions, the KVP must be attached before the voice call is released.

For eServices/Multimedia, Interaction Concentrator automatically stores information about the reason that processing of an interaction stopped. If you want to track the reasons for agents stopping multimedia interactions, ensure that the `Stop Reason` key with relevant values is available to your agents

through their desktop applications. Interaction Concentrator 8.x also stores information about the party that issued the request to stop processing an interaction, when the party is known.

Using UserEvent-Based KVP Data

Some agent desktop applications issue UserEvents to set KVP data after the agent's participation in the voice interaction has completed (that is, after the call is released). You can configure an ICON application that captures Voice details to store UserEvent-based KVP data in its IDB. When you configure the ICON application, you use ICON application configuration options—instead of the attached data specification XML file—to specify which KVPs ICON should store. Then you can configure Genesys Info Mart to extract this data from the IDB `G_CUSTOM_DATA_S` table.

Note the following about Genesys Info Mart processing of UserEvent-based KVP data:

- This functionality is supported for voice interactions only.
- This functionality is supported for logged-in agents and IVR applications that emulate logged-in agents.
- Data from only the `G_CUSTOM_DATA_S` table in IDB is extracted. UserEvent-based KVP data is *not* extracted from `G_CUSTOM_DATA_P`, nor are custom agent states extracted from the `G_CUSTOM_STATES` table in IDB.
- Applications that issue UserEvents must be sure to set the fields inside the UserEvent properly. Unlike with call-based attached data, T-Server does not validate the contents of the UserEvents, nor does it propagate their KVP data values among related calls, such as consultations, transfers, or conferences.

Genesys Info Mart stores the extracted UserEvent data in the same fact and dimension tables as the data that is sourced from call-based attached data. During deployment planning, you decide which Info Mart fact or dimension column should receive data from each UserEvent-based KVP that is of interest for reporting. During deployment configuration, you must configure ICON application options to specify which KVPs should be stored in `G_CUSTOM_DATA_S`. Also, you must configure Genesys Info Mart mapping between those KVPs and the Info Mart facts and dimensions (see "User Data Mapping Tables" on page 115).

Note: If you report on Outbound Contact details, you must configure ICON to store UserEvent-based KVP data for the `GSW_CALL_ATTEMPT_GUID` KVP.

For more information about how Genesys Info Mart populates its facts and dimensions from UserEvent-based KVP data and call-based attached data, see

the section about populating Genesys Info Mart data in the *Genesys Info Mart 8.1 User's Guide*.

Storing User Data

Genesys Info Mart 8.x provides a flexible storage mechanism for a potentially very large number of KVPs that are attached by Genesys solutions.

User data can be stored as facts or dimensions. Genesys Info Mart release 8.x provides you with the flexibility to store the same key as a fact and as a dimension. Starting with release 8.1, Genesys Info Mart also provides the flexibility to store KVP values in user-data fact tables as a character data type, a numeric data type, or a date/time data type.

High-Cardinality User Data

High-cardinality user data is stored as facts. Although there are no absolute limits on the quantity of high-cardinality user data that you can store, be mindful of database storage space and database performance.

The following fact extension tables are used for storage of predefined high-cardinality user data:

- `IRF_USER_DATA_GEN_1`.
- Custom fact extension tables. (Use the sample script that is provided for the `IRF_USER_DATA_CUST_1` table to add these tables.)

Low-Cardinality User Data

Low-cardinality user data is most efficiently stored as dimensions. The following dimension tables are used for storage of predefined low-cardinality user data:

- `INTERACTION_DESCRIPTOR`.
- Custom dimension tables. (Use the sample script that is provided for the `USER_DATA_CUST_DIM_1` table to add these tables.)

When you add custom tables to the Info Mart database, keep in mind the following limitation: The upper limit for low-cardinality user data is 800 custom dimension tables.

Low-Cardinality vs. High-Cardinality User Data

There is no simple rule about where the cutoff is between low-cardinality data and high-cardinality data, in terms of numbers of possible values for a KVP. For any given set of KVPs that will be stored in the same table, the number of combinations is a more important consideration than the number of possible values for each KVP. Genesys recommends that the number of rows in any dimension table not exceed 50,000.

User Data Mapping Tables

The Info Mart database schema contains two mapping tables that are required to process and store user data:

- `CTL_UD_TO_UDE_MAPPING`

- CTL_UDE_KEYS_TO_DIM_MAPPING

The mapping tables instruct Genesys Info Mart on how to populate the user-data KVPs and where in the Info Mart database to store them.

Figure 7 is a snapshot of a sample populated mapping table. For detailed descriptions of the table columns, refer to the *Genesys Info Mart 8.1 Reference Manual* for your particular RDMBS type.

KEY_NAME	TABLE_NAME	COLUMN_NAME	PROPAGATION_RULE	DEFAULT_VALUE	ACTIVE_FLAG
CUSTOMER_SEGMENT	INTERACTION_DESCRIPTOR	CUSTOMER_SEGMENT	CALL		1
SERVICE_TYPE	INTERACTION_DESCRIPTOR	SERVICE_TYPE	CALL		1
SERVICE_SUBTYPE	INTERACTION_DESCRIPTOR	SERVICE_SUBTYPE	CALL		1
BUSINESS_RESULT	INTERACTION_DESCRIPTOR	BUSINESS_RESULT	CALL		1
CASE_ID	IRF_USER_DATA_GEN_1	CASE_ID	CALL		1
CUSTOMER_ID	IRF_USER_DATA_GEN_1	CUSTOMER_ID	CALL		1
CUSTOM_DATA_1..16	IRF_USER_DATA_CUST_1	CUSTOM_DATA_1_1..16	PARTY		0
DIM_ATTRIBUTE_1_1..5	USER_DATA_CUST_DIM_1	DIM_ATTRIBUTE_1_1..5	PARTY		0

Figure 7: Example of Populated Mapping Table

To use user data in your reports, you must ensure that the upstream data sources (for example, T-Server) are configured to propagate user data as required; configure ICON to store this data in IDB; modify a Genesys-provided user-data script to work for your preferred KVP names; adjust the Info Mart database to be able to store this data; and specify mapping rules for Genesys Info Mart to extract and process this data.

For information about how to customize the Genesys-provided user-data script for the initial deployment of Genesys Info Mart, see “Customizing the User Data Template” on page 184. Later, as your reporting needs evolve, you can similarly add more custom tables to the Info Mart database to store new user data KVPs.

Special Requirements

Be aware that your downstream reporting application might have specific requirements for user data that Genesys Info Mart does not store by default. For example, if you want to include reporting on social media metrics—such as Sentiment, Actionability, or Influence—in a Genesys Interactive Insights (GI2) deployment, GI2 has specific requirements for the user-data tables and columns that you must use to store the required KVPs.

For the convenience of GI2 customers, the sample ICON attached-data specification file that Genesys Info Mart provides includes commented-out lines that specify the social media–related KVPs that GI2 requires. For more information about configuring Genesys Info Mart to store the user data that GI2 requires for social media metrics, see the section about social media user data in the chapter about customizing the GI2 universe and reports in the *Genesys Interactive Insights 8.1 User’s Guide.*

Viewing Current KVP Mappings

To view the current mapping of user-data KVPs to Info Mart tables and columns, execute the following SQL command against the Info Mart database:

```
SELECT * FROM CTL_UD_TO_UDE_MAPPING
```

The result will show both predefined and custom mappings in your deployment.

Mapping Call-Based Attached Key-Value Pairs

Genesys Info Mart extracts attached data KVPs from the ICON Voice details and ICON Multimedia details data sources—specifically, from the following IDB tables:

- G_CUSTOM_DATA_S
- G_USERDATA_HISTORY
- G_SECURE_USERDATA_HISTORY
- GM_F_USERDATA and GM_L_USERDATA

ICON automatically stores predefined router-specific KVPs in the G_ROUTE_RESULT table, provided that you configure Universal Routing Server (URS) to attach the KVPs to interactions. For more information, see "Universal Routing" on page 110.

ICON stores voice and multimedia attached KVPs in the G_USERDATA_HISTORY and G_SECURE_USERDATA_HISTORY tables, based on the options that you configure in the ICON application and in the ICON attached data specification (adata_spec) XML file. For multimedia interactions, ICON automatically stores predefined KVPs, including workbin usage details, in the G_USERDATA_HISTORY table.

ICON stores multimedia-specific attached data in the GM_F_USERDATA and GM_L_USERDATA tables. The storage and mapping of the attributes that Genesys Info Mart uses is predefined, and Genesys Info Mart does not process any custom KVPs in these tables.

By default, Genesys Info Mart stores the values of the KVPs that are listed in the "KVPs that are mapped by default" section of Table 7 on page 122. If you want to store additional, custom user data, you must define mapping and propagation rules and store them in Control tables in the Info Mart database. Genesys Info Mart then extracts the user data details on the basis of the KVP name and stores the value in the table and column that you designate in the mapping tables (CTL_UD_TO_UDE_MAPPING and CTL_UDE_KEYS_TO_DIM_MAPPING).

As part of the mapping, you specify the propagation rule that Genesys Info Mart will use to determine what value to store if more than one value is extracted for the same key in the same interaction. For more information, see "Propagation Rules".

You can also specify what value is stored as a default if a particular KVP is missing for an interaction.

Propagation Rules

For user-data keys that might have multiple values over the life of an interaction, the propagation rules enable you to specify which KVP value will be stored for a particular IRF or, starting with release 8.1.2, MSF, based on who changed (added, updated, or deleted) the KVP value or when it was changed.

Valid values for the propagation rules are the following:

- `CALL`—Genesys Info Mart stores the latest KVP value that is associated with the interaction when the interaction leaves the resource that is the subject of the IRF or MSF record, regardless of who changed the KVP value.

  The `CALL` propagation rule is suitable for most business requirements for most KVPs.

- `PARTY`—Genesys Info Mart stores the latest KVP value as changed by that party to the interaction, regardless of when it was changed.

  Use the `PARTY` propagation rule to capture KVP values that are set after the interaction leaves the agent (for example, during after-call work [ACW]) or for user data that should be associated only with the subject of the IRF or MSF record and not propagated to other resources.

  Notes:
  - `PARTY` is the only propagation rule that enables you to capture KVP values that are set after the interaction leaves the handling resource that is the subject of the IRF record.
  - Do not use the `PARTY` rule for user data that is associated with virtual queues for multimedia interactions. User-data transformation for the `PARTY` rule relies on party information from the target IRF that is not available at mediation time.
  - The `PARTY` rule does not work for user data in MSFs for voice or multimedia interactions that are cleared or abandoned in a virtual queue. In these cases, Genesys Info Mart cannot find user data for the MSF, and it will appear as if no user data was changed. Genesys Info Mart will store the default value for the KVP, if defined.

- `IRF`—For a KVP value that was changed during the timespan of the IRF (specifically, between the time that the interaction started mediation to the handling resource and the time that the interaction leaves the handling resource), Genesys Info Mart stores the latest KVP value that is associated with the interaction, regardless of who changed the KVP value.

  The initial state of user data for the applicable IRF record is empty.

  For user data in MSFs, the IRF to which the MSF belongs provides the context in which the KVP value is set. For example, if MSF1 represents

the mediation that occurs before the interaction is distributed to the resource that is the subject of IRF1, Genesys Info Mart will store the latest change, if any, that occurred during the MSF1 mediation interval, which is calculated from the mediation start time of IRF1 and the end time that is reported in MSF1.

- `IRF_FIRST_UPDATE`—For a KVP value that was changed during the extended timespan of the IRF for a handling resource, Genesys Info Mart stores the first update to the KVP value that is associated with the IRF, regardless of who changed the KVP value. The extended timespan starts from the time that the interaction started mediation to any handling resource (in other words, not necessarily to the resource that is the subject of the IRF) up to the time that the interaction leaves the handling resource in question. In a scenario with call redirection, the timespan of the IRF for the handling resource to which the interaction is eventually routed includes the durations of all preceding IRFs that have a technical result of `Redirected/RoutedOnNoAnswer` or `Redirected/Unspecified`.

  This propagation rule was introduced in release 8.1.1.

  For user data in MSFs, the IRF to which the MSF belongs provides the context in which the KVP value is set. For example, if MSF*n* represents the mediation that occurs immediately before the interaction is distributed to IRF*n,* Genesys Info Mart will store the first change that was made during the extended timespan of IRF*n,* up to and including MSF*n*.

To illustrate the effect of the propagation rules, Tables 5 and 6 provide the reporting results for the various propagation rules in typical sample scenarios.

Table 5 provides examples for user data in IRF records.

Table 5: Reporting Results for Propagation Rules—Sample Scenarios for IRF

Scenario*	User Data Result for Associated IRF Record, by Propagation Rule			
	CALL	PARTY	IRF	IRF_FIRST_UPDATE
1. A strategy attaches a KVP with `Value0` and routes the call to Agent1, who updates the KVP to `Value1` and then transfers the call to Agent2.	Agent1: `Value1` Agent2: `Value1`	Agent1: `Value1` Agent2: Empty	Agent1: `Value1` Agent2: Empty	Agent1: `Value0` Agent2: Empty
* T-Server settings: `merged-user-data=merged-over-main, consult-user-data=inherited`				

Table 5: Reporting Results for Propagation Rules—Sample Scenarios for IRF (Continued)

Scenario*	User Data Result for Associated IRF Record, by Propagation Rule			
	CALL	PARTY	IRF	IRF_FIRST_UPDATE
2. Same as Scenario 1, except that Agent2 subsequently updates the KVP to `Value2` after the call is released.	Agent1: `Value1` Agent2: `Value1`	Agent1: `Value1` Agent2: `Value2`	Agent1: `Value1` Agent2: Empty	Agent1: `Value0` Agent2: Empty
3. Strategy1 attaches a KVP with `Value0` and routes the call to Agent1, who updates the KVP to `Value1` and then initiates a two-step conference to Agent2, using Strategy2, which updates the KVP to `Value2`.	Agent1: `Value2` Agent1 (Initiated Consult): `Value2` Agent2: `Value2`	Agent1: `Value1` Agent1 (Initiated Consult): Empty Agent2: Empty	Agent1: `Value2` Agent1 (Initiated Consult): `Value2` Agent2: `Value2`	Agent1: `Value0` Agent1 (Initiated Consult): `Value2` Agent2: `Value2`
4. A route-on-no-answer (RONA) variant of Scenario 3: Strategy1 attaches a KVP with `Value0` and attempts to route the call to Agent0. When Agent0 does not answer, Strategy1 routes the call to Agent1, who updates the KVP to `Value1` and then initiates a two-step conference to Agent2, using Strategy2, which updates the KVP to `Value2`.	Agent0: `Value0` Agent1: `Value2` Agent1 (Initiated Consult): `Value2` Agent2: `Value2`	Agent0: Empty Agent1: `Value1` Agent1 (Initiated Consult): Empty Agent2: Empty	Agent0: `Value0` Agent1: `Value1` Agent1 (Initiated Consult): `Value2` Agent2: `Value2`	Agent0: `Value0` Agent1: `Value0` Agent1 (Initiated Consult): `Value2` Agent2: `Value2`
* T-Server settings: `merged-user-data=merged-over-main, consult-user-data=inherited`				

Table 6 on page 121 provides examples for user data in MSF records, provided that Genesys Info Mart release 8.1.2 or later has been configured to record the association.

Table 6: Reporting Results for Propagation Rules—Sample Scenarios for MSF

Scenario*	User Data Result for Associated MSF Record, by Propagation Rule			
	CALL	PARTY	IRF	IRF_FIRST_ UPDATE
Note: The MSFs for Queue1 and Queue2 happen before the resulting IRFs for Agent1 and Agent2, respectively.				
1. A strategy attaches a KVP with `Value0` and places a call in Queue1. The call stays in the queue for a few days, before it is distributed to Agent1, who updates the KVP to `Value1` and then transfers the call to Agent2 via Queue2.	Queue1: `Value0` Queue2: `Value1`	Queue1: Empty Queue2: Empty	Queue1: `Value0` Queue2: Empty	Queue1: `Value0` Queue2: Empty
2. Same as Scenario 1, except that a strategy updates the KVP to `Value2` when it places the interaction in `Queue2`.	Queue1: `Value0` Queue2: `Value2`	Queue1: Empty Queue2: Empty	Queue1: `Value0` Queue2: `Value2`	Queue1: `Value0` Queue2: `Value2`
* T-Server settings: `merged-user-data=merged-over-main, consult-user-data=inherited`				

Common Attached Data KVPs

Table 7 on page 122 describes the attached data KVPs that contact centers most commonly use for reporting purposes.

Some KVPs have numeric values, as specified in Table 7. The values of all of the other KVPs in Table 7 are strings, for which the maximum length in IDB is 255 characters. However, depending on the RDBMS, Genesys Info Mart further restricts the maximum length of certain KVP values, and you might similarly need to restrict the length of KVP values for customized user-data dimension tables. For more information, see "RDBMS Considerations" on page 124.

Some of the KVPs are Genesys-defined, while others are user-defined. Both types of KVPs can populate predefined as well as custom, deployment-specific facts and dimensions, according to the user-defined mapping rules.

Note: Your applications do not need to attach all of the KVPs that are listed in Table 7.

Use the "Worksheet for Mapping User Data" on page 428 or, starting with release 8.1.2, the User Data Assistant to map the KVP names in your contact center to the target Info Mart tables and column names.

Table 7: Commonly Used Attached Data KVPs

KVP Name	KVP Description
KVPs that are mapped by default	
Business Result *	The business result of the interaction.
CaseID *	The case identifier in an external case management application.
CustomerID *	The customer identifier in an external customer relationship management (CRM) application.
CustomerSegment *	Identifies a segment of the customer base to which the customer has been assigned. Customer segments are typically based on criteria such as revenue potential, service plan, or demographic attributes.
`GSW_CALL_ATTEMPT_GUID` *	In Outbound Contact deployments, stores the `GSW_CALL_ATTEMPT_GUID` call attempt ID that is assigned by OCS.
IPurpose *	The presence and value of this KVP affects how the ETL populates the IRF table. • `0` = Not self-service. The IVR application is considered to be a mediation resource. • `1` = Self-service. The IVR application is considered to be a handling resource. Genesys Info Mart creates a record in the IRF table. For more information, see page 109.
* The mapping of the KVP to fact or dimension tables is predefined, as shown in the mapping worksheet. ** The mapping of the KVP to fact or dimension tables is user-defined. Configure the mapping in the Control tables, as described in the section about storing user data.	

Table 7: Commonly Used Attached Data KVPs (Continued)

KVP Name	KVP Description
Revenue*	The amount of revenue that was generated for the customer interaction. **Note:** In deployments that use Reporting and Analytics Aggregates (RAA) aggregation (including GI2 deployments), ensure that the resources that set the value of this KVP have been configured or trained, as applicable, to use only integer-type values. Any non-numeric values will cause aggregation to fail. For more information about the problems that result from non-numeric KVP values, as well as information about recovering from the job failure, see the information about checking for incorrect data type in the troubleshooting chapter in the *Reporting and Analytics Aggregates 8.1 Deployment Guide.*
Satisfaction *	The numerical customer-satisfaction score for the customer interaction. **Note:** In deployments that use RAA aggregation (including GI2 deployments), ensure that the resources that set the value of this KVP have been configured or trained, as applicable, to use only integer-type values. Any non-numeric values will cause aggregation to fail. For more information about the problems that result from non-numeric KVP values, as well as information about recovering from the job failure, see the information about checking for incorrect data type in the troubleshooting chapter in the *Reporting and Analytics Aggregates 8.1 Deployment Guide.*
ServiceObjective *	The time objective (in seconds) to service the interaction, based on the customer segment, service type, and media type.
ServiceSubType *	The detailed type of service that the customer is requesting.
ServiceType *	The type of service that the customer is requesting.
Routing KVPs	
RRequestedSkillCombination *	The agent skills that are required to service the interaction.
RStrategyName *	The name of the routing strategy that is servicing the interaction.
RTargetObjectSelected *	The name of the target object that the router selects.
* The mapping of the KVP to fact or dimension tables is predefined, as shown in the mapping worksheet. ** The mapping of the KVP to fact or dimension tables is user-defined. Configure the mapping in the Control tables, as described in the section about storing user data.	

Table 7: Commonly Used Attached Data KVPs (Continued)

KVP Name	KVP Description
RTargetTypeSelected *	The type of routing target that the router selects—for example, `0` = Agent, `1` = Place, `2` = Agent Group, `3` = Place Group, `100` = Default Route. Valid values are defined by ICON. For more information, see the list of Route Target Type dictionary values in the *Interaction Concentrator 8.x Physical Data Model* for your RDBMS.
KVPs that can be mapped to custom user-data tables	
IApplication **	The IVR application that is servicing the interaction. **Note:** Genesys Info Mart uses this KVP during processing. You must configure ICON to store this KVP in IDB even if you do not require `IApplication` for your reporting and do not configure it to be stored in user-data tables in the Info Mart database.
IResult **	The technical result of the IVR application. • `1` = Completed • `2` = Abandoned • `3` = Transferred
IResultReason **	The reason for the IVR technical result. Values should be of *low cardinality.*
ISpeechRecognition **	Indicates whether IVR speech recognition was used. • `0` = No • `1` = Yes
ITextToSpeech **	Indicates whether IVR text-to-speech was used. • `0` = No • `1` = Yes
* The mapping of the KVP to fact or dimension tables is predefined, as shown in the mapping worksheet. ** The mapping of the KVP to fact or dimension tables is user-defined. Configure the mapping in the Control tables, as described in the section about storing user data.	

RDBMS Considerations

Depending on the RDBMS, you might have to consider RDBMS limitations when you specify customized mapping of KVPs to user-data dimension tables.

Microsoft SQL Server and PostgreSQL

An RDBMS limitation restricts the maximum length of index keys:

- On Microsoft SQL Server, to 900 bytes
- On PostgreSQL, to 2730 bytes

To avoid failures of the transformation job, the total length of values for KVPs that are configured to populate any one of the user-data dimension tables must not exceed the RDBMS limit.

For the KVPs that are mapped to any one of the user-data dimension tables by default, Genesys Info Mart restricts the length of individual user-data KVP values to 170 bytes, to ensure that combinations of KVPs do not exceed the limit.

This limitation does not apply to KVPs that are configured to populate user-data fact tables.

Oracle

There are no RDBMS limitations that you have to consider.

Chapter

7 Outbound Contact Data

This chapter describes how Genesys Info Mart processes Outbound Contact data. It contains the following sections:

Genesys Info Mart and OCS Record Field Data

When support for Outbound Contact details is configured in your deployment, Genesys Info Mart writes data about every outbound contact attempt. Outbound Contact Server (OCS) may attach a number of custom attributes (*record fields*) to each contact attempt. ICON then stores Record Field data in the Outbound Contact extension tables in IDB—that is, the tables that start with the `GO_` prefix (for example, `GO_FieldHist`). These IDB tables become the source of the Record Field data for Genesys Info Mart.

Note: For a complete list of the Outbound Contact extension IDB tables that Genesys Info Mart uses, see page 443.

Genesys Info Mart stores Record Field data that is defined in OCS calling lists in several places in the Info Mart database:

- Predefined dimensions, such as `Record Type`, `Record Status`, and `Contact Info Type`. These are mandatory record fields.
- Predefined facts, such as `Record ID`, `Chain ID`, `Chain N`, `Dialing From`, `Dialing Until`, and `Contact Info`. These are mandatory record fields.

- User-defined dimensions, such as the columns in the `RECORD_FIELD_GROUP_1` and `RECORD_FIELD_GROUP_2` tables. These are nonmandatory record fields.
- User-defined facts, such as `Record Field 1` through `Record Field40` in the `CONTACT_ATTEMPT_FACT` table. These are nonmandatory record fields.

Predefined and Custom Fields

Some `Field` objects are populated in the Configuration Database at the time of this database creation, to represent record fields that are typical to most campaign environments; these fields are referred to as *predefined* record fields. To reflect the record fields that are typical to campaigns in your specific contact center, you may have to create other `Field` objects in the Configuration Database; these fields are referred to as *user-defined* or *custom* fields. This division between predefined and custom fields is true for both mandatory and nonmandatory record fields. In other words, some fields that are mandatory from the perspective of Genesys Info Mart are not predefined in the Configuration Database.

Mandatory Record Field Data

Table 8 on page 129 shows how Genesys Info Mart uses the data from each mandatory `Field` object. Some fields map directly to Info Mart table columns, whereas others are used indirectly in calculations.

All the `Field` objects that are listed in Table 8 require the `icon_attribute` option in their configuration. You *must* configure each predefined field and each user-defined field that are expected in the Info Mart database. Proper configuration of the `Field` objects ensures that ICON stores the field value in a specified table in IDB. Genesys Info Mart requires the ICON application to store mandatory field data in its database, regardless of whether the field maps directly to an Info Mart table column.

Genesys Info Mart Server retrieves the field value from IDB and stores the value in a specified field of a specified Info Mart table or uses the value in calculations of other fields. Note that some predefined fields that are used in calculations also require a value that indicates a positive result.

For more information about how to configure `Field` objects so that ICON will store the required Record Field data, see Procedure: Configuring the storage of OCS record field data, on page 267.

Table 8: Mandatory Record Field Data

OCS Mandatory Field Name	Column Name in Info Mart CONTACT_ATTEMPT_FACT Table
agent_id	No direct mapping
app_id	No direct mapping
attempt*	ATTEMPT_ORDINAL
call_result*	No direct mapping
call_time	No direct mapping
campaign_id	No direct mapping
chain_id*	CHAIN_ID
chain_n	CHAIN_N
contact_info*	CONTACT_INFO
contact_info_type	CONTACT_INFO_TYPE_KEY
daily_from	DAILY_FROM_SECONDS DAILY_FROM_TIME CONTACT_DAILY_FROM_TIME
daily_till	DAILY_UNTIL_SECONDS DAILY_UNTIL_TIME CONTACT_DAILY_UNTIL_TIME
dial_sched_time	DIAL_SCHED_TIME CONTACT_DIAL_SCHED_TIME
group_id	No direct mapping
record_id	RECORD_ID
record_status	RECORD_STATUS_KEY
record_type	RECORD_TYPE_KEY
switch_id	No direct mapping
* OCS automatically attaches the value of this field to events, in a KVP with a key name that starts with GSW_. Except for the GSW_CALL_ATTEMPT_GUID KVP, Genesys Info Mart does not obtain record field data from attached data.	

Table 8: Mandatory Record Field Data (Continued)

OCS Mandatory Field Name	Column Name in Info Mart CONTACT_ATTEMPT_FACT Table
treatments	No direct mapping
tz_dbid	TIME_ZONE_KEY CONTACT_DAILY_FROM_TIME CONTACT_DAILY_UNTIL_TIME CONTACT_DIAL_SCHED_TIME
* OCS automatically attaches the value of this field to events, in a KVP with a key name that starts with GSW_. Except for the GSW_CALL_ATTEMPT_GUID KVP, Genesys Info Mart does not obtain record field data from attached data.	

Note: The following columns in the CONTACT_ATTEMPT_FACT table are no longer populated, although they remain in the schema:

- IXN_START_TIME
- IXN_START_TIME_KEY
- CONTACT_IXN_START_TIME
- CONTACT_WITHIN_DAILY_RANGE

Nonmandatory (Custom) Record Field Data

Genesys Info Mart can optionally store a limited number of nonmandatory record fields in the following tables.

- CONTACT_ATTEMPT_FACT
- RECORD_FIELD_GROUP_1
- RECORD_FIELD_GROUP_2

The CONTACT_ATTEMPT_FACT table can store up to 40 nonmandatory fields in the following formats:

- 20 integers: NUMBER(10)
- 10 floating-point numbers: NUMBER(14,4)
- 10 strings: VARCHAR(255)

RECORD_FIELD_GROUP_1 and RECORD_FIELD_GROUP_2 each store up to 10 nonmandatory fields (strings).

If you want to report on nonmandatory record fields, you must configure each nonmandatory field properly so that OCS attaches the value of the field, and you must configure ICON to store this field in a designated table. (For more information, see Procedure: Configuring the storage of OCS record field data,

on page 267.) In addition to configuring each field similarly to mandatory field configuration, configure options in each Field object to indicate the Genesys Info Mart table and column into which the data should be loaded.

You can use any field that you choose. The data type of the Field object must match the data type of the target Info Mart database table and column. Interaction Concentrator stores all custom field data as strings. The Genesys Info Mart ETL performs all necessary data conversions between strings and other target data types.

Each Field object maps to one and only one table and column in the Info Mart database. Nulls are loaded for any unmapped columns in the Info Mart CONTACT_ATTEMPT_FACT table. The Unspecified value is loaded for any unmapped columns in the Info Mart RECORD_FIELD_GROUP_1 and RECORD_FIELD_GROUP_2 tables.

RECORD_FIELD_GROUP_1 and RECORD_FIELD_GROUP_2 column values should be of low cardinality. Storing record fields with high-cardinality values will cause a decrease in the performance for both the ETL and your report queries.

For information about nonmandatory fields that have special meaning for Genesys Info Mart, see "Right Person Contacted Record Field" and "Conversion Record Field" on this page. You can also use the worksheet in Appendix A to plan your mapping of OCS Record Fields to the Info Mart table columns.

Right Person Contacted Record Field

Although Right Person Contacted is not a mandatory field, it has significance to Genesys Info Mart. It can be any Field object that you designate by adding the right_person option to the default section of the Annex tab on the Field object. The option value specifies the value of the field when the right person is contacted—for example, TRUE, YES, or 1. If the value of this field matches the configured option value (which is case-insensitive), Genesys Info Mart sets the RPC_FLAG in its CONTACT_ATTEMPT_FACT table to 1. For more information, see Procedure: Configuring the mapping of OCS record fields, on page 268.

If you want to report on right person contacted, you must configure ICON to store nonmandatory field data in its database. For information about how to configure ICON to store field data, see Procedure: Configuring the storage of OCS record field data, on page 267.

Conversion Record Field

Although Conversion is not a mandatory field, it has significance to Genesys Info Mart. It can be any Field object that you designate by adding the conversion option to the default section of the Field object's Annex tab. The option value specifies the value of the field when the purpose of the Outbound

Contact attempt has been achieved—for example, `TRUE`, `YES`, or `1`. If the value of this field matches the configured option value (which is case-insensitive), Genesys Info Mart sets the `CONVERSION_FLAG` in its `CONTACT_ATTEMPT_FACT` table to `1`.

If you want to report on conversion, you must configure ICON to store nonmandatory field data in its database. For information about how to store field data, see Procedure: Configuring the mapping of OCS record fields, on page 268.

Part

2 Deploying Genesys Info Mart

Part 2 of this document provides information about the configuration and installation of Genesys Info Mart and the supporting components and objects.

This information appears in the following chapters:

- Chapter 8, “Configuration Tasks,” on page 135
- Chapter 9, “Preparing Interaction Concentrator,” on page 159
- Chapter 10, “Preparing Genesys Info Mart Database,” on page 181
- Chapter 11, “Configuring DAPs,” on page 207
- Chapter 12, “Configuring the Genesys Info Mart Application,” on page 231
- Chapter 13, “Configuring Supporting Objects,” on page 253
- Chapter 14, “Genesys Info Mart Configuration Options Reference,” on page 271
- Chapter 15, “Installing Genesys Info Mart Components,” on page 321

Chapter

8 Configuration Tasks

This chapter summarizes the actions that are required to deploy Genesys Info Mart and to implement specific functionality. It contains the following sections:

Overview of Deployment Tasks

The deployment of Genesys Info Mart is a complex task because it involves a number of components, such as Interaction Concentrator, the Genesys Info Mart Server, the Genesys Info Mart Manager or Genesys Info Mart Administration Console, and databases. It also involves components in the Genesys Configuration Layer, Management Layer, and Media Layer. Before you start the deployment, review the recommended deployment task flow carefully, and make sure that you understand all of the activities that are involved.

The following table summarizes, at a high level, the task flow for activities that are required to plan and execute the Genesys Info Mart deployment.

Task Summary: Deploying Genesys Info Mart

Objective	Related Procedures and Actions
1. Plan the deployment.	Review the information in Part 1 of this guide to familiarize yourself with Genesys Info Mart architecture, requirements, and functionality, and to plan the implementation of Genesys Info Mart features in your deployment. If you want to include secure connection features, review the information in the part about Transport Layer Security (TLS) and in the chapter about client-side port definition in the *Genesys 8.x Security Deployment Guide,* to identify deployment requirements. See also the summary task table for "Enabling Secure Connection Features" on page 154. Appendix A on page 421 provides worksheets to assist you in identifying and recording the required configuration information for your environment.
2. Prepare Interaction Concentrator to provide source data for Genesys Info Mart.	1. Prepare the Interaction Concentrator server (ICON) application(s), which capture and store data from data sources such as Configuration Server and T-Server. For detailed information, see "Preparing ICON" on page 162. **Note:** You might have to configure additional, ICON-related settings on supporting objects to enable specific functionality. For more information, see Objective 6 on page 138. 2. Prepare the Interaction Database (IDB) instance(s) from which Genesys Info Mart will obtain data. For detailed information, see "Preparing IDBs" on page 178. 3. Prepare the database access point (DAP) application(s) that enable ICON to access the IDB(s) it populates. For more information, see the *Interaction Concentrator 8.x Deployment Guide*. 4. Ensure that all required data sources have been enabled, to identify them to Genesys Info Mart as available. Starting with release 8.1.2, also ensure that the ICON applications are enabled. To enable specific functionality, such as support for Outbound Contact details or High Availability (HA), see also the summary task tables in "Enabling Specific Functionality" on page 140.

Task Summary: Deploying Genesys Info Mart (Continued)

Objective	Related Procedures and Actions
3. Prepare the target Genesys Info Mart database and views.	1. Create and configure database schemas to process and store detailed reporting data. For detailed information, see: ◆ "Preparing the Info Mart Database" on page 191 ◆ "Using Database Links" on page 194 ◆ "Database Tuning" on page 200 2. After you have installed Genesys Info Mart (Step 8 on page 139), see also Procedure: Creating Genesys Info Mart tenant views—Oracle, on page 351.
4. Configure the database access points (DAPs) that Genesys Info Mart uses to access source and target databases.	Create DAPs, or modify the configuration of existing DAPs, to enable: • Genesys Info Mart Server to access the IDBs. • Genesys Info Mart Server to access the Info Mart database. • The Genesys Info Mart Administration Console to access the Info Mart database in order to monitor ETL jobs. For detailed information, see Chapter 11 on page 207.
5. Configure the Genesys Info Mart Server application.	1. Import the required `Application Template` for the Genesys Info Mart application. 2. Create and configure an `Application` object for Genesys Info Mart Server. The required configuration settings depend directly on the Genesys Info Mart features that you want to implement and on your choice of an end-user reporting tool, such as Genesys Interactive Insights (GI2). For more information, see Chapter 12 on page 231. To enable specific functionality, such as support for Outbound Contact details or high availability (HA), see also the summary task tables in "Enabling Specific Functionality" on page 140.

Task Summary: Deploying Genesys Info Mart (Continued)

Objective	Related Procedures and Actions
6. Configure other objects that are required by ICON and Genesys Info Mart.	Configure necessary options in individual configuration objects, for use by ICON or Genesys Info Mart Server: • `Switch` (for Voice and Multimedia details). For more information, see Procedure: Configuring the switch for ICON and Genesys Info Mart reporting, on page 255. • `Media Type Business Attribute` (for Multimedia details). For more information, see Procedure: Setting Media Type Business Attribute object options for Genesys Info Mart reporting, on page 260. • `DN` (for Voice and Multimedia details). For more information, see Procedure: Configuring a DN for ICON and Genesys Info Mart reporting, on page 261. • `Script` (for Multimedia details). For more information, see Procedure: Setting Script object options for Genesys Info Mart reporting, on page 265. • `Field` (for Outbound Contact details). For more information, see "Configuring Field Objects" on page 266. For detailed information about ICON and Genesys Info Mart requirements for configuration settings on these objects, see Chapter 13 on page 253.
7. Prepare the Genesys Info Mart Server host.	See "Preparing the Genesys Info Mart Server Host" on page 323.

Task Summary: Deploying Genesys Info Mart (Continued)

Objective	Related Procedures and Actions
8. Install the Genesys Info Mart components.	1. Install the Genesys Info Mart application on its host. For more information: • For Windows, see Procedure: Installing the Genesys Info Mart application (Windows), on page 327. • For UNIX, see Procedure: Installing the Genesys Info Mart application (UNIX), on page 328. 2. Verify the host setup. For more information, see "Verifying Host Requirements" on page 329. 3. Install a management GUI, either: • (Starting with release 8.1.4) Install the Genesys Info Mart Manager plug-in for Genesys Administrator Extension (GAX). For more information, see Procedure: Installing the Genesys Info Mart Manager plug-in, on page 331. • Install the Administration Console GUI on the same host on which Configuration Manager is installed. For more information, see Procedure: Installing the Genesys Info Mart Administration Console (Windows), on page 333. **Note:** The Administration Console host must use the Microsoft Windows operating system.
9. Complete the basic deployment.	Perform post-installation activities: • Start Genesys Info Mart. For more information, see Chapter 17 on page 357. • Review Genesys Info Mart logs to verify that the deployment is complete and configuration is correct. • In the management GUI, review the status of `Job_InitializeGIM` to verify successful initialization of the database and, starting with release 8.1.2, successful update of the IDBs. • (Required for multi-tenant deployments; optional for single-tenant deployments) Create tenant-specific, read-only views on the Info Mart database. For more information, see "Creating Genesys Info Mart Read-Only Tenant Views" on page 348. • Access the management GUI, to continue managing Genesys Info Mart jobs. For more information, see "Accessing Genesys Info Mart Manager" on page 340 or "Accessing the Genesys Info Mart Administration Console" on page 341. For more information about how to manage jobs, see the *Genesys Info Mart 8.1 Operations Guide*.
(Optional) Enable aggregation. **Note:** This is required for GI2.	For more information, see Task Summary: Enabling Aggregation in a Genesys Info Mart Deployment, on page 157.

Enabling Specific Functionality

This subsection provides targeted task flows to enable Genesys Info Mart to support various specific functions.

Enabling Storage of Configuration Details

The following table summarizes the task flow to enable Genesys Info Mart to capture and store Configuration details.

Task Summary: Enabling the Storage of Configuration Details

Objective	Related Procedures and Actions
1. Prepare the ICON application to capture and store Configuration details.	When you configure the ICON application, ensure that: • At a minimum, you set the following option values: ◦ The `role` option includes the value `cfg`. **Note:** If you follow the Genesys recommendation to provide a dedicated Configuration details ICON (and IDB), ensure that the `role` option specifies only `cfg`. ◦ The mandatory options that are listed in Step 2 on page 163 are set as specified. • You configure an ADDP connection to Configuration Server (or to an HA pair of Configuration Servers). For more information about how to configure the ICON application, see Procedure: Configuring the ICON application, on page 163. For more information about all of the ICON configuration options that affect ICON processing and data storage, see the chapter about configuration options in the *Interaction Concentrator 8.x Deployment Guide*.

Task Summary: Enabling the Storage of Configuration Details (Continued)

Objective	Related Procedures and Actions
2. Prepare the IDB instance(s) from which Genesys Info Mart will obtain Configuration details.	(Mandatory for Genesys Info Mart 8.1.0 and 8.1.1 deployments; optional for 8.1.2 and later) After you have run the ICON-provided scripts to create the IDB(s), execute one of the following SQL scripts, which Genesys Info Mart provides: • `update_idb_for_gim.sql` • `update_idb_for_gim_mm.sql` For more information, see Procedure: Preparing IDBs, on page 179.
3. Enable ICON to access the Configuration details IDB.	When you create and configure the DAP application that enables the Configuration details ICON to access IDB, ensure that you add the DAP to the `Connections` tab of the ICON `Application` object.
4. Configure and install the Genesys Info Mart application.	When you configure the Genesys Info Mart application, ensure that: • The values of Genesys Info Mart configuration options are suitable for your deployment. For more information, see “Options Summary Charts” on page 240. • You configure a connection to the ICON application (or HA set) that stores Configuration details. For more information about how to configure the Genesys Info Mart application, see Chapter 12 on page 231.
5. Enable Genesys Info Mart to access the Configuration details IDB.	• When you configure the DAP application that enables Genesys Info Mart to access the Configuration details IDB, ensure that the `role` option includes the value `ICON_CFG`. For more information, see “Extraction DAPs” on page 211. • On the `Connections` tab of the Genesys Info Mart `Application` object, add a connection to this DAP.

Enabling Storage of Voice Details

The following table summarizes the task flow to enable Genesys Info Mart to capture and store Voice details.

Task Summary: Enabling the Storage of Voice Details

Objective	Related Procedures and Actions
1. Prepare the ICON application to capture and store Voice details.	When you configure the ICON application, ensure that: • At a minimum, you set the following option values: ◦ The `role` option includes the values `gcc`, `gls`, `gud`. ◦ The mandatory options that are listed in Step 2 on page 163 are set as specified. ◦ Other voice-related options that are described in Table 9 on page 165 are set as specified. For more information, see Procedure: Configuring the ICON application, on page 163. For more information about all of the ICON configuration options that affect ICON processing and data storage, see the chapter about configuration options in the *Interaction Concentrator 8.x Deployment Guide*. • You configure ADDP connections to all of the T-Servers (or HA pairs of T-Servers) in the deployment that ICON might be required to monitor.
2. Prepare the IDB instance(s) from which Genesys Info Mart will obtain Voice details.	(Mandatory for Genesys Info Mart 8.1.0 and 8.1.1 deployments; optional for 8.1.2 and later) After you have run the ICON-provided scripts to create the IDB(s), execute the following SQL script, which Genesys Info Mart provides: • `update_idb_for_gim.sql` For more information, see Procedure: Preparing IDBs, on page 179.
3. Enable ICON to access the Voice details IDB(s).	When you create and configure the DAP application that enables the Voice details ICON to access IDB, ensure that you add the DAP to the `Connections` tab of the ICON `Application` object.
4. Configure and install the Genesys Info Mart application.	When you configure the Genesys Info Mart application, ensure that: • The values of Genesys Info Mart configuration options are suitable for your deployment. For more information, see "Options Summary Charts" on page 240. • You configure a connection to all of the ICON applications (or HA sets) that store Voice details. For more information about how to configure the Genesys Info Mart application, see Chapter 12 on page 231.

Task Summary: Enabling the Storage of Voice Details (Continued)

Objective	Related Procedures and Actions
5. Prepare other objects as required to enable ICON and Genesys Info Mart reporting for Voice details.	• Configure the required ICON-related and Genesys Info Mart–related options on the `Annex` tab of: ◦ The `Switch` object that handles voice interactions, as described in Procedure: Configuring the switch for ICON and Genesys Info Mart reporting, on page 255. ◦ `DN` objects for IVRs or virtual queues for voice interactions, as described in Steps 2 through 5 in Procedure: Configuring a DN for ICON and Genesys Info Mart reporting, on page 261. ◦ (Optional, for DN-specific queue thresholds) `DN` objects for ACD queues or virtual queues for voice interactions, as described in Steps 6 and • in Procedure: Configuring a DN for ICON and Genesys Info Mart reporting, on page 261. • Verify that all of the T-Servers that ICON is required to monitor for Voice details are enabled.
6. Enable Genesys Info Mart to access the Voice details IDB.	• When you configure the DAP application that enables Genesys Info Mart to access the Voice details IDB, ensure that the `role` option includes the value `ICON_CORE`. For more information, see "Extraction DAPs" on page 211. • On the `Connections` tab of the Genesys Info Mart `Application` object, add a connection to this DAP.
7. Optimize the extraction job.	After you have created the Info Mart database schema, modify the `GSYS_DNPREMOTELOCATION` table in the Info Mart database, as required, to optimize performance of the merge procedure during the extraction job. For more information, see Procedure: Configuring Info Mart database for merge, on page 193.

Enabling Storage of Multimedia Details

The following table summarizes the task flow to enable Genesys Info Mart to capture and store Multimedia details.

Task Summary: Enabling the Storage of Multimedia Details

Objective	Related Procedures and Actions
1. Prepare the ICON application to capture and store Multimedia details.	When you configure the ICON application, ensure that: • At a minimum, you set the following option values: ◆ The `role` option includes the values `gcc`, `gls`, `gud`. ◆ The mandatory options that are listed in Step 2 on page 163 are set as specified. ◆ Other Multimedia-related options that are described in Table 9 on page 165 are set as specified. For more information, see Procedure: Configuring the ICON application, on page 163. For more information about all of the ICON configuration options that affect ICON processing and data storage, see the chapter about configuration options in the *Interaction Concentrator 8.x Deployment Guide*. • You configure ADDP connections to all of the Interaction Servers (or HA pairs of Interaction Servers) in the deployment that ICON might be required to monitor.
2. Prepare the IDB instance from which Genesys Info Mart will obtain Multimedia details.	(Mandatory for Genesys Info Mart 8.1.0 and 8.1.1 deployments; optional for 8.1.2 and later) After you have run the ICON-provided scripts to create the IDB, execute the following SQL script, which Genesys Info Mart provides: • `update_idb_for_gim_mm.sql` For more information, see Procedure: Preparing IDBs, on page 179.
3. Enable ICON to access the Multimedia details IDB(s).	When you create and configure the DAP application that enables the Multimedia details ICON to access IDB, ensure that you add the DAP to the `Connections` tab of the ICON `Application` object.
4. Configure and install the Genesys Info Mart application.	When you configure the Genesys Info Mart application, ensure that: • The values of Genesys Info Mart configuration options are suitable for your deployment. For more information, see "Options Summary Charts" on page 240. • You configure a connection to all of the ICON applications (or HA sets) that store Multimedia details. For more information about how to configure the Genesys Info Mart application, see Chapter 12 on page 231.

Task Summary: Enabling the Storage of Multimedia Details (Continued)

Objective	Related Procedures and Actions
5. Configure other objects, as required, to enable ICON and Genesys Info Mart reporting for Multimedia details.	• (For reporting on virtual-queue activity) Verify that the required ICON-related options are set correctly on the switch and DNs. In the `gts` section on the `Annex` tab of: ◦ The `Switch` object that handles multimedia interactions, `support-dn-type-5=1` ◦ `DN` objects for virtual queues for multimedia interactions, `monitor=1` For more information about these options, see Table 17 on page 256 and Table 18 on page 263, respectively. • In the `gim-etl` section on the `Annex` tab of the `Switch` object that handles multimedia interactions, configure the Genesys Info Mart–related options, as described in Table 17 on page 256. • (Optional) To override Genesys Info Mart `Application` settings for configured thresholds by media type at the tenant, switch, DN, or script level, set the supported Genesys Info Mart–related options on the applicable objects. For a summary of the media-specific thresholds that you can configure at various levels, see Table 14 on page 241. For more information about how to configure the applicable objects, see: ◦ Procedure: Configuring the switch for ICON and Genesys Info Mart reporting, on page 255 ◦ Procedure: Setting Media Type Business Attribute object options for Genesys Info Mart reporting, on page 260 ◦ Procedure: Configuring a DN for ICON and Genesys Info Mart reporting, on page 261 ◦ Procedure: Setting Script object options for Genesys Info Mart reporting, on page 265 • Verify that all of the Interaction Servers that ICON is required to monitor for Multimedia details are enabled.
6. Enable Genesys Info Mart to access the Multimedia details IDB.	• When you configure the DAP application that enables Genesys Info Mart to access the Multimedia details IDB, ensure that the `role` option includes the value `ICON_MM`. For more information, see “Extraction DAPs” on page 211. • On the `Connections` tab of the Genesys Info Mart `Application` object, add a connection to this DAP.

Task Summary: Enabling the Storage of Multimedia Details (Continued)

Objective	Related Procedures and Actions
7. Customize the processing and storage of 3rd Party Media interaction data, if applicable.	• (Recommended) Before you start ETL processing, add to the `MEDIA_TYPE` dimension table any new or missing online media types that might be associated with 3rd Party Media interactions in your deployment. Ensure that the `IS_ONLINE` flag in the `MEDIA_TYPE` record is set to `1`. (You do not have to add offline media types in advance.) For more information about how to add online media types to the Info Mart database schema, see Procedure: Setting up media types for online interactions, on page 344. For more information about why Genesys recommends predefining expected online media types for 3rd Party Media interactions, see "Online and Offline Interactions" on page 98. • During runtime, if you want to change whether Genesys Info Mart processes a particular interaction subtype, alter the record for that interaction type in the `INTERACTION_TYPE` dimension table, to set the value of the `IGNORE` field appropriately. For more information about disabling transformation of interaction subtypes, see "Interaction Types and Subtypes" on page 99.

Enabling Storage of Outbound Contact Details

The following table summarizes the task flow to enable Genesys Info Mart to capture and store Outbound Contact–related data.

Task Summary: Enabling the Storage of Outbound Contact Details

Objective	Related Procedures and Actions
1. Prepare the ICON application to store Outbound Contact details.	When you configure the ICON application, ensure that: • At a minimum, you set the following option values: ◦ The `role` option includes the value `gos`. ◦ The mandatory options that are listed in Step 2 on page 163 are set as specified. ◦ `gos-write-duplicate-metrics=1`. For more information, see Table 9 on page 165. For more information about all of the ICON configuration options that affect ICON processing and data storage, see the chapter about configuration options in the *Interaction Concentrator 8.x Deployment Guide*. • You configure ADDP connections to all of the Outbound Contact Servers (or HA pairs of OCS instances) in the deployment that ICON might be required to monitor.
2. Prepare the IDB instance(s) from which Genesys Info Mart will obtain Outbound Contact details.	(Mandatory for Genesys Info Mart 8.1.0 and 8.1.1 deployments; optional for 8.1.2 and later) After you have run the ICON-provided scripts to create the IDB(s), execute one of the following SQL scripts, which Genesys Info Mart provides: • `update_idb_for_gim.sql` • `update_idb_for_gim_mm.sql` For more information, see Procedure: Preparing IDBs, on page 179.
3. Enable ICON to access the Outbound Contact details IDB(s).	When you create and configure the DAP application that enables the Outbound Contact details ICON to access IDB, ensure that you add the DAP to the `Connections` tab of the ICON `Application` object.
4. Configure Outbound Contact–related objects so that OCS will send the required data and ICON will store it.	• On the `Annex` tab of each `Field` configuration object that describes a single field within a record, in the `default` section, configure the `icon_attribute` option, to control ICON storage of the data. For more information, see Procedure: Configuring the storage of OCS record field data, on page 267. • Verify that all of the Outbound Contact Servers that ICON is required to monitor for Outbound Contact details are enabled.

Task Summary: Enabling the Storage of Outbound Contact Details (Continued)

Objective	Related Procedures and Actions
5. Map Outbound Contact–related objects to columns in the Genesys Info Mart database.	For every nonmandatory field: • In the `gim-etl-mapping` section, configure the `table-name` and `column-name` options. • In the `default` section, configure the `right_person` and `conversion` options, if applicable. For more information, see Procedure: Configuring the mapping of OCS record fields, on page 268.
6. Configure and install the Genesys Info Mart application.	When you configure the Genesys Info Mart application, ensure that: • The values of Genesys Info Mart configuration options are suitable for your deployment. For more information, see "Options Summary Charts" on page 240. • You configure a connection to all of the ICON applications (or HA sets) that store Outbound Contact details. For more information about how to configure the Genesys Info Mart application, see Chapter 12 on page 231.
7. Enable Genesys Info Mart to access the Outbound Contact details IDB.	• When you configure the DAP application that enables Genesys Info Mart to access the Outbound Contact details IDB, ensure that the `role` option includes the value `ICON_OCS`. For more information, see "Extraction DAPs" on page 211. • On the `Connections` tab of the Genesys Info Mart `Application` object, add a connection to this DAP.

Enabling Storage of User Data

The following table summarizes the task flow to enable Genesys Info Mart to capture and store user data (interaction-based attached data or `UserEvent`-based data) that is attached to voice or multimedia interactions.

Note: Starting with release 8.1.2, the Genesys Info Mart installation package (IP) includes the User Data Assistant, a tool to automate preparation of the user-data configuration files. The User Data Assistant is a Microsoft Excel file, `User_Data_Assistant.xlsm`, located in the `sql_scripts` folder in your Genesys Info Mart 8.1 IP. The User Data Assistant requires Microsoft Excel 2007 or later, with macros enabled. Full instructions on using the tool are within the tool itself.

The tool is intended to be used only for fresh deployments and not to modify existing user-data extension tables or an existing ICON attached-data specification file.

The tool does not support PostgreSQL.

Task Summary: Enabling the Storage of User Data

Objective	Related Procedures and Actions
1. Configure the ICON application to store user data.	When you prepare the ICON application, ensure that you set the following option values: • The `role` option includes the value `gud` • `adata-extensions-history = none` (for Voice details) • `adata-reasons-history = none` (for Voice details) • `adata-userdata-history = none` (for Voice or Multimedia details)
2. Specify the user data that ICON will store in IDB.	1. Identify the key-value pairs (KVPs) from various applications that Genesys Info Mart requires for data processing. For more information about the KVPs that contact centers typically use for reporting purposes, see "User-Data Sources and KVPs" on page 107. 2. For call-based attached data, modify the ICON attached data specification file to capture the KVPs that you require and to control in which IDB table(s) ICON will store the data. By default, ICON expects that the name of the attached data specification is `ccon_adata_spec.xml`. Starting with release 8.1.2, you can use the User Data Assistant to prepare the customized attached data specification file. Genesys Info Mart provides an example of a customized attached data specification. The sample specification file, `ccon_adata_spec_GIM_example.xml`, is included in the `sql_scripts` folder in your Genesys Info Mart 8.1 IP. The sample is reproduced in Appendix B on page 437. For more information about how to customize the attached data specification file, see Procedure: Customizing your ICON attached data specification file, on page 177. See also the section about storing attached data in the *Interaction Concentrator 8.x Deployment Guide*. /continued...

Task Summary: Enabling the Storage of User Data (Continued)

Objective	Related Procedures and Actions
2. Specify the user data that ICON will store in IDB (continued).	3. For `UserEvent`-based user data, set ICON configuration options, as required, to store user data in the `G_CUSTOM_DATA_S` table in IDB. For more information, see "Using UserEvent-Based KVP Data" on page 114. For more information about the required ICON configuration settings, see the descriptions of the options in the `custom-states` section in the *Interaction Concentrator 8.x Deployment Guide*.
3. Plan the Info Mart tables in which you want to store user data.	1. Identify the Info Mart fact or dimension tables in which you want user data to be stored. For more information, see "Storing User Data" on page 115. **Note:** Some target tables and columns are predefined. You can also store user data in custom-defined tables and columns, which you create later when you modify and execute SQL scripts. 2. Map the user-data KVPs to the Info Mart tables and columns that you have identified. Genesys provides a worksheet (see "Worksheet for Mapping User Data" on page 428) to assist you in completing the mapping. Starting with release 8.1.2, you can use the User Data Assistant to map the KVPs. For more information about how to map attached data and specify the propagation rules, see "Mapping Call-Based Attached Key-Value Pairs" on page 117.

Task Summary: Enabling the Storage of User Data (Continued)

Objective	Related Procedures and Actions
4. Modify the Info Mart database schema, as required, to store the custom user data.	1. Prepare the user-data SQL script to: • Specify creation of the target user-data fact and dimension tables and columns, in accordance with your mapping worksheet. • Map the user-data key names to the target tables and columns (in the `CTL_UD_TO_UDE_MAPPING` table), and the key names in the user-data dimension table(s) to the `IRF_USER_DATA_KEYS` table (in the `CTL_UDE_KEYS_TO_DIM_MAPPING` table). Starting with release 8.1.2, the User Data Assistant generates the required script, based on the mapping you provide in the tool. For information about manually preparing the script, see Procedure: Customizing the user-data template script, on page 185. Sample template scripts (`make_gim_UDE_template.sql` and `make_gim_UDE_template_partitioned.sql`) are available in the RDBMS-specific scripts folder in your Genesys Info Mart 8.1 installation package. 2. Execute the modified script when you create the rest of the Info Mart database schema, or when you complete the deployment after installing Genesys Info Mart. **Note:** If you are modifying an existing Info Mart database schema, Genesys recommends that you back up the Info Mart database before you execute the script.
5. (Optional, for release 8.1.2 and later) Enable storage of user data for interactions that are in mediation.	Configure the `link-msf-userdata` configuration option in the `[gim-etl]` section of `ACD Queue` and `Virtual Queue DNs` and `Interaction Queue` or `Interaction Workbin Script` objects. For more information, see Procedure: Configuring a DN for ICON and Genesys Info Mart reporting, on page 261 and Procedure: Setting Script object options for Genesys Info Mart reporting, on page 265.
6. (Optional, for release 8.1.101.07 and later) Enable Genesys Info Mart to extract KVPs that are sourced from the `Reasons` or `Extensions` attributes, in addition to the `UserData` attribute, in data-source events.	Turn off filtering of user data in IDB by setting the `filterUserData` startup parameter in the `gim_etl_server` file to `false`. For more information, see “Modifying JVM Startup Parameters” on page 360.

Task Summary: Enabling the Storage of User Data (Continued)

Objective	Related Procedures and Actions
7. (Optional) Streamline Genesys Info Mart processing of user data.	If your historical reporting involves the use of large quantities of user data, consider increasing the value of the `ud-io-parallelism` option (see page 312), to enhance performance of the transformation job.

Enabling High Availability

The following table summarizes the task flow to enable Genesys Info Mart to provide high availability (HA) of reporting data in a new or existing deployment.

Task Summary: Enabling High Availability in a New or Existing Deployment

Objective	Related Procedures and Actions
1. Provide redundancy for monitoring the data sources and storing source data.	Create and configure identical redundant ICON `Applications` for the type of details you require. For more information, see Procedure: Configuring the ICON application, on page 163. Except for identical configuration, there are no special requirements for HA ICON applications. However, especially ensure the following: • On the `Connections` tab of every ICON `Application` object, an ADDP connection has been configured to every data source that the ICON instance might be required to monitor. The connection list can include HA pairs of Configuration Servers, T-Servers, Interaction Servers, or Outbound Contact Servers. **Note:** Remember that you configure an overt connection only to the primary data source. For example, say the deployment consists of: ◦ An HA pair of Configuration Servers (primary Configuration Server 1 and backup Configuration Server 1a) ◦ Two HA pairs of T-Servers (primary T-Server 1, backup T-Server 1a, primary T-Server 2, and backup T-Server 2a) ◦ A single set of redundant ICONs for Configuration and Voice details (ICON-1 and ICON-2, storing data in IDB-1 and IDB2, respectively) Then both ICON-1 and ICON-2 must have connections to Configuration Server 1, T-Server 1, and T-Server 2. /continued...

Task Summary: Enabling High Availability in a New or Existing Deployment (Continued)

Objective	Related Procedures and Actions
1. Provide redundancy for monitoring the data sources and storing source data (continued).	• On the `General` tab of every data-source application whose availability and activity Genesys Info Mart must monitor, the `State Enabled` check box is selected. • Starting with release 8.1.2, on the `General` tab of every ICON in the Genesys Info Mart `Application` connections that Genesys Info Mart must monitor, the `State Enabled` check box is selected. • The value of the `dss-no-data-tout` option (see page 167) is suitable for your HA deployment; the default value is 5 minutes. **Note:** If you need to restart existing ICON application(s) for configuration information to take effect, Genesys recommends that you do not restart ICON at this time; see Step 5 on page 154.
2. Provide redundancy for the Genesys Info Mart database sources (IDBs).	Create and initialize identical redundant IDBs, together with their associated DAPs, to store the type of ICON details that you require. • For more information about how to prepare an IDB, see "Preparing IDBs" on page 178. • For more information about how to configure a DAP for Interaction Concentrator, see the *Interaction Concentrator 8.x Deployment Guide*.
3. Provide redundancy for Genesys Info Mart to access the redundant IDBs.	Configure identical redundant DAPs to enable the Genesys Info Mart Server to access the redundant IDBs. Create a separate DAP to enable access to each redundant IDB. Starting with release 8.1.2, ensure that all these extraction DAPs are enabled (`State Enabled` check box on the `General` tab of the `Application` object is selected). For more information about how to configure the required DAPs, see "Extraction DAPs" on page 211.

Task Summary: Enabling High Availability in a New or Existing Deployment (Continued)

Objective	Related Procedures and Actions
4. Configure the Genesys Info Mart Server application.	Configure the Genesys Info Mart `Application` object in the usual way. For more information, see Procedure: Creating the Genesys Info Mart application, on page 235. There are no special requirements for the Genesys Info Mart Server application in an HA environment. However, especially ensure the following: • On the `Connections` tab of the Genesys Info Mart `Application` object, a connection has been configured to every redundant ICON `Application` in the deployment. • On the `Connections` tab of the Genesys Info Mart `Application` object, a connection has been configured to every DAP that provides access to a redundant IDB in the deployment. • The values of the following options are suitable for your HA deployment: ◦ `extract-data-stuck-threshold` (see page 289) ◦ `max-time-deviation` (see page 296)
5. Start or, if necessary, restart the ICON applications.	For more information about how to start ICON, see the chapter about starting and stopping ICON in the *Interaction Concentrator 8.x Deployment Guide*. Also see the ICON documentation for information about when you might need to restart an existing ICON.
6. In a new deployment, start the Genesys Info Mart Server.	For more information about how to start Genesys Info Mart, see Chapter 17 on page 357. **Note:** The Genesys Info Mart Server supports dynamic changes. If you are adding HA to an existing deployment, it is not necessary to interrupt the Genesys Info Mart job schedule or to stop Genesys Info Mart for new connection information or other configuration option changes to take effect.
7. Verify the deployment.	Review logs to confirm the results of the configuration check, to verify connections to all redundant Interaction Concentrator instances and data sources, and to verify correct configuration.

Enabling Secure Connection Features

The following table summarizes the task flow to enable Genesys Info Mart to implement the features that Genesys provides to secure connections in the deployment. All of the security features are optional.

Task Summary: Enabling Secure Connections

Objective	Related Procedures and Actions
• Enable Transport Layer Security (TLS) protocol on the connections from Genesys Info Mart Server to Configuration Server and Message Server.	1. (For UNIX-based deployments only) Install the Genesys Security Pack on the Genesys Info Mart Server host, and set the applicable environment variable to specify the path to the Security Pack libraries. For more information, see the chapter about Security Pack installation in the *Genesys 8.x Security Deployment Guide.* 2. If certificates do not already exist, create and install certificates on the Genesys Info Mart Server host, as well as on the Configuration Server and Message Server hosts. Starting with release 8.1.3, Genesys Info Mart supports mutual TLS, which requires exchange of certificates from both the TLS Server and the TLS Client. For more information, see the chapters about certificate generation and installation and about Genesys TLS configuration in the *Genesys 8.x Security Deployment Guide.* 3. If necessary, modify the configurations of the Configuration Server and Message Server applications to: a. Add a new port for secure connections. On the Configuration Server `Application` object, select the `Auto Detect` mode. On the Message Server `Application` object, select the `Secured` mode. b. Use a host certificate. For full details, see the chapter about Genesys TLS configuration in the *Genesys 8.x Security Deployment Guide.* 4. On the Genesys Info Mart `Application` object, add connections to Configuration Server and Message Server (see Step 7 on page 237 in Procedure: Creating the Genesys Info Mart application). When you add the connection(s), ensure that you specify the port that you created in Step 3.
• (Starting with release 8.1.3) Enable compliance with Federal Information Processing Standards (FIPS).	Genesys Info Mart support for TLS complies with FIPS, but there are additional steps to enable FIPS mode. For information about setting up your Java environment to be compliant with FIPS, see the section about enabling FIPS in a Genesys Java environment in the *Genesys 8.1 Security Deployment Guide.*

Task Summary: Enabling Secure Connections (Continued)

Objective	Related Procedures and Actions
• Enable client-side port definition for the connection from Genesys Info Mart Server to Configuration Server.	1. When you install Genesys Info Mart, specify the connection parameters that Genesys Info Mart will use for the initial connection to Configuration Server. For full details, see the procedures about configuring a connection to Configuration Server in the chapter about client-side port definition in the *Genesys 8.x Security Deployment Guide*. 2. In the Genesys Info Mart `Application` object, add or modify the connection to Configuration Server, to specify the connection parameters (port number and, optionally, IP address) that Genesys Info Mart will use to reconnect to Configuration Server after a switchover or disconnection. You configure the parameters in the `Transport Protocol Parameters` text box on the `Advanced` tab of the connection properties. For full details, see the procedure about adding Configuration Server to client connections in the chapter about client-side port definition in the *Genesys 8.x Security Deployment Guide*.
• Enable client-side port definition for the connection from Genesys Info Mart Server to Message Server.	In the Genesys Info Mart `Application` object, add or modify the connection to Message Server, to specify the connection parameters (port number and, optionally, IP address) that Genesys Info Mart will use. You configure the parameters in the `Transport Protocol Parameters` text box on the `Advanced` tab of the connection properties. For full details, see the procedure about configuring client connections in the chapter about client-side port definition in the *Genesys 8.x Security Deployment Guide*.
• Enable use of the Secure Socket Layer (SSL) protocol on the JDBC connections between Genesys Info Mart Server and its source and target databases.	When you configure the extraction DAP(s) and the Info Mart DAP, use the `jdbc-url` option to specify the URL information as required by your RDBMS to implement JDBC over SSL. For more information about configuring the `jdbc-url` option, see Procedure: Configuring JDBC DAPs, on page 225 or Procedure: Configuring non-JDBC DAPs, on page 228. For more information about the parameters that you must specify, see your RDBMS vendor documentation.

Enabling Aggregation

The following table summarizes the task flow to enable Genesys Info Mart to support the aggregation process and to populate aggregate tables in the Info Mart database.

Note: If you plan to use Genesys Interactive Insights (GI2), you must enable aggregation.

Task Summary: Enabling Aggregation in a Genesys Info Mart Deployment

Objective	Related Procedures and Actions
1. Install the aggregation engine software.	Install GI2 or the separate Reporting and Analytics Aggregates (RAA) package. • For full information about how to install GI2, see the *Genesys Interactive Insights 8.x Deployment Guide.* • For full information about how to install the RAA package, see the *Reporting and Analytics Aggregates 8.x Deployment Guide.* **Note:** RAA is included as part of the GI2 installation package. If you intend to install GI2, you do not have to install RAA first.

Task Summary: Enabling Aggregation in a Genesys Info Mart Deployment (Continued)

Objective	Related Procedures and Actions
2. Configure the Genesys Info Mart `Application` object and other applications and objects in your Genesys environment, as required, to support aggregation.	Settings on the Genesys Info Mart `Application` object—as well as settings for ICON filtering options, switch and DN options, Business Objects components (for GI2), and other applications—affect aggregation. Be aware that GI2 has specific requirements for certain settings. • For full information about how to configure Genesys Info Mart for aggregation, see the *Reporting and Analytics Aggregates 8.x Deployment Guide.* **Note:** You must restart the Genesys Info Mart Server if you modify the following settings during runtime: ◆ `jdbc-url` option in the Info Mart DAP ◆ `aggregation-engine-class-name` in the Genesys Info Mart `Application` ◆ User name for the Info Mart database ◆ Password for the Info Mart database Genesys Info Mart generates a log event (message 55-20153) when these settings are changed. Genesys recommends that you set an alarm on the log message, to prompt you to restart the Genesys Info Mart Server. • For information about how to configure Business Objects and create the aggregation-related database objects for GI2, see the *Genesys Interactive Insights 8.x Deployment Guide.*
3. (Optional) Configure custom calendars.	See Procedure: Configuring custom calendars, on page 347. **Note:** If your reports use custom calendar dimensions, create the custom calendars before you start aggregation.
4. Start the aggregation engine.	If you have configured the Genesys Info Mart scheduler to control the aggregation process (`run-aggregates=true`), the aggregation job will start automatically at the scheduled time (as specified by `aggregate-schedule`). To run aggregation continuously, set `aggregate-schedule=0 0` and `aggregate-duration=24:00`. For full information about how to start the aggregation engine, see the *Reporting and Analytics Aggregates 8.x Deployment Guide.* For more information about how to schedule and manage the aggregation job in Genesys Info Mart, see the *Genesys Info Mart 8.1 Operations Guide*.

Chapter

9 Preparing Interaction Concentrator

This chapter describes how to prepare the Interaction Concentrator applications (ICONs) and Interaction Databases (IDBs) that provide data to Genesys Info Mart.

Refer to this chapter before you install your Genesys Info Mart 8.1 application.

This chapter contains the following sections:

For information about the various data-source topologies that are supported by Genesys Info Mart 8.1, see Chapter 3 on page 69.

For information about database access accounts and privileges, see "Database Object Owners and User IDs" on page 93.

Overview: Preparing Interaction Concentrator for Genesys Info Mart

The following table summarizes the task flow for preparing Interaction Concentrator to capture and store reporting data for Genesys Info Mart.

Task Summary: Preparing Interaction Concentrator for Genesys Info Mart

Objective	Related Procedures and Actions
Capture information to support detailed reporting about contact center configuration and activities for interactions of any type, as well as related agent activity.	Set up the ICON applications to capture and store data from instances or HA pairs of data sources—Configuration Server, for Configuration details; T-Server, for Voice details; Interaction Server, for Multimedia details; or Outbound Contact Server (OCS), for Outbound Contact details. For information about how to configure the ICON application as required for Genesys Info Mart, see Procedure: Configuring the ICON application, on page 163. For full details about creating and configuring an ICON `Application` object, see the *Interaction Concentrator 8.x Deployment Guide.*
(Optional) Provide support for high availability (HA) of reporting data.	Set up redundant ICON applications to store ICON details in a highly available (HA) set of redundant IDBs. For more information, see "Enabling High Availability" on page 152.
(Optional) Minimize the required database storage space, while supporting detailed reporting of interactions of any type and related agent activity.	Set up the ICON application to exclude extra details from being stored in IDB. For more information, see Procedure: Controlling IDB storage, on page 173.
(Optional) Capture attached data to support detailed reporting of interactions of any type.	Indicate what attached data ICON should store in IDB. You must include all key-value pairs (KVPs) that Genesys Info Mart requires for data processing, as described in "Genesys Info Mart and Attached Data" on page 103. Optionally, include any additional KVPs that you require for end reports. For more information, see "Enabling Storage of User Data" on page 148 and Procedure: Customizing your ICON attached data specification file, on page 177.

Task Summary: Preparing Interaction Concentrator for Genesys Info Mart (Continued)

Objective	Related Procedures and Actions
Enable extraction of ICON details from the IDBs for reporting purposes.	1. Create the IDB instances that are required for your deployment. For more information about the SQL scripts that Interaction Concentrator provides, as well as information about how to execute the scripts to create the IDB schema, see the chapter about creating IDB in the *Interaction Concentrator 8.x Deployment Guide*. 2. Ensure that the database access account that the ETL jobs will use to access IDB data is available and has the required user account privileges (see Table 3 on page 95). 3. If you are preparing an IDB for use with Genesys Info Mart 8.1.0 or 8.1.1, modify IDB so that the jobs that extract, transform, and load (ETL) data extract relevant reporting data. For more information, see Procedure: Preparing IDBs, on page 179. **Note:** Starting with release 8.1.2, Genesys Info Mart jobs automatically execute the scripts to modify IDB. However, if you are adding IDBs to an existing 8.1.2 or later deployment, you might still choose to perform this step manually, for reasons that are described in "Updating IDBs in Genesys Info Mart 8.1.2 and Later" on page 178.
Verify correct configuration for database connectivity between ICON and IDB.	• Verify configuration settings on the database access points (DAPs) that enable ICON(s) to access IDB(s). For more information, see the *Interaction Concentrator 8.x Deployment Guide*. • Verify that all required DAPs have been added to the `Connections` tab of the applicable ICON `Application` objects.
Prepare other objects, as required, to support detailed reporting about interactions of various media types.	Configure ICON-related settings on other objects: • `Switch` (for Voice or Multimedia details)—For more information, see Procedure: Configuring the switch for ICON and Genesys Info Mart reporting, on page 255. • `DN` (for Voice or Multimedia details)—For more information, see Procedure: Configuring a DN for ICON and Genesys Info Mart reporting, on page 261. • `Field` (for Outbound Contact details)—For more information, see Procedure: Configuring the storage of OCS record field data, on page 267.

Task Summary: Preparing Interaction Concentrator for Genesys Info Mart (Continued)

Objective	Related Procedures and Actions
Verify correct configuration for connectivity between ICON and its data sources.	Verify that all data sources from which information is required to be captured have been configured to be available. Specifically: • The data sources are included in the connections of the applicable ICON applications. • All data-source applications that you want Genesys Info Mart to consider to be active are enabled.

Preparing ICON

Your Genesys Info Mart deployment requires at least one ICON application and one IDB. However, depending on your chosen topology, you may have additional ICON applications or additional IDBs for separate storage of Configuration details, Voice details, Multimedia details, and Outbound Contact details.

Within Configuration Manager, you must configure each ICON `Application` object and related objects in the deployment, in accordance with the Genesys Info Mart deployment requirements that are described in this section.

The type of data that Genesys Info Mart will extract from a particular ICON and IDB depends on your topology and reporting requirements. The required configuration settings, therefore, also depend on your topology and reporting requirements. In some cases, you must configure settings on other configuration objects (`DN, Field, Switch, Script`), as well as on the ICON `Application` object. For more information about ICON-related settings on other configuration objects, see Chapter 13 on page 253.

Notes:

- The valid values that are listed for ICON application options in this chapter do not necessarily represent the entire set of values that are available in ICON; instead, these are the values that make sense in a reporting environment that is based on Genesys Info Mart.
- The values that Genesys Info Mart requires for certain ICON options (“mandatory options”) are not the default values.
- Changes to the values of certain ICON options—including mandatory options—require a restart of ICON. When you restart an existing ICON application, active calls are lost.

- All the IDBs from which Genesys Info Mart 8.x extracts data must have been populated only by ICONs that conform to Genesys Info Mart 8.x requirements. If you are adding Genesys Info Mart 8.x to an existing ICON deployment and you upgrade ICON or change ICON settings to meet Genesys Info Mart 8.x requirements, you must create new IDBs as part of the upgrade and start populating them only after all of the mandatory settings have been updated. Otherwise, Genesys Info Mart 8.x will not be able to process data from IDB. Any data that was collected in IDB before the upgrade or change in the ICON settings will not be processed by Genesys Info Mart.

For more information about the ICON options, see the *Interaction Concentrator 8.x Deployment Guide*.

Recommendations on ICON Deployment and Upgrade

- If you are deploying Interaction Concentrator at the same time as Genesys Info Mart, follow the installation and configuration instructions in the *Interaction Concentrator 8.x Deployment Guide,* while observing the Genesys Info Mart deployment requirements that are documented in this chapter.
- You can deploy Genesys Info Mart in an environment in which Interaction Concentrator 8.x has been deployed already, provided that ICON was configured in accordance with Genesys Info Mart requirements when Interaction Concentrator was originally deployed. If there is a large amount of data in IDB by the time Genesys Info Mart is added to the deployment, be aware that this approach might result in significant data latency while the Genesys Info Mart ETL processes the backlog.

Procedure: Configuring the ICON application

Purpose: To enable ICON to capture Configuration details, Voice details, Multimedia details, or Outbound Contact details for Genesys Info Mart.

Start of procedure

1. Create and configure an Interaction Concentrator `Application` object, as described in the configuration and installation chapter in the *Interaction Concentrator 8.x Deployment Guide*.

   Ensure that the `State Enabled` check box on the `General` tab is selected.
2. Review the options from the `callconcentrator` section, and modify settings as required for Genesys Info Mart. At a minimum, you must set the following option values for all types of ICON details:
   - `use-dss-monitor=true`
   - `calls-in-the-past=true`

- `om-force-adata=true`
- `gls-active-reason-codes=true`
- `partition-type=2`
- `role` (the value depends on the type of ICON details that you require that ICON application to store)

Table 9 on page 165 describes the required ICON options, as well as other ICON options and recommended settings that are important for Genesys Info Mart.

3. Configure a connection to every data source that will supply data for Genesys Info Mart. Ensure that all connections between ICON and its data sources use the Advanced Disconnect Detection Protocol (ADDP).

End of procedure

Next Steps

- (Optional) Procedure: Controlling IDB storage, on page 173
- (Optional, for Voice or Multimedia details) Procedure: Customizing your ICON attached data specification file, on page 177
- Procedure: Preparing IDBs, on page 179

ICON Configuration Options

Table 9 on page 165 describes required or recommended settings for options that you set on the ICON `Application` object.

The options are grouped by type of ICON details to which the option applies:

- All details (page 165)
- Voice or Multimedia details (page 168)
- Voice details only (page 171)
- Multimedia details only (page 172)
- Outbound Contact details only (page 173)

For more information about all of the ICON configuration options, see the chapter about configuration options in the *Interaction Concentrator 8.x Deployment Guide.*

Note: In the context of this section, the term *interaction* is used generically to refer to voice calls and multimedia interactions.

Table 9: ICON Application Options—callconcentrator Section, by Type of Details

Area of Functionality	Option Name	Recommended Value	Description
All details			
ICON role	role*	• `cfg`—For Configuration details • `gcc`, `gud`, `gls`—For Voice or Multimedia details • `gos`—For Outbound Contact details	Specifies what type of data this ICON instance processes and stores in an IDB: • `cfg`—The initial configuration state and a history of configuration changes, which are retrieved from Configuration Server. • `gcc`—Interaction-related and party-related information. • `gud`—T-Server (for Voice) or Interaction Server (for Multimedia) data about the attached data that is associated with interactions. • `gls`—T-Server (for Voice) or Interaction Server (for Multimedia) data about agent states and agent login sessions. • `gos`—Data that pertains to Outbound Contact calls and campaigns.

* Mandatory option—Genesys Info Mart will not function if the option value is not set as specified.

** Recommended option—Data quality might be compromised if you do not use the recommended value. For more information about data-quality considerations, see the *Genesys Info Mart 8.1 User's Guide*.

Table 9: ICON Application Options—callconcentrator Section, by Type of Details (Continued)

Area of Functionality	Option Name	Recommended Value	Description
Configuration information	cfg-annex (ICON 8.1.4 and later)	1 (true)	Specifies that ICON will gather and store data configured as options in `RPT_` configuration sections on the `Annex` tabs of Person, Agent Group, DN, DN Group, and Switch `Application` objects. Genesys Interactive Insights (GI2) uses the data associated with `Person`, `Agent Group`, and `DN Group` configuration objects to control visibility for certain data and reports.
	cfg-auto-resync**	1 (true)	Specifies that ICON will automatically initiate resynchronization of configuration data between Configuration Database and IDB when a serious discrepancy is detected. **Note:** This option requires Interaction Concentrator release 8.1.000.23 or later.

* Mandatory option—Genesys Info Mart will not function if the option value is not set as specified.

** Recommended option—Data quality might be compromised if you do not use the recommended value. For more information about data-quality considerations, see the *Genesys Info Mart 8.1 User's Guide*.

Table 9: ICON Application Options—callconcentrator Section, by Type of Details (Continued)

Area of Functionality	Option Name	Recommended Value	Description
Session monitoring	dss-no-data-tout**	`60` (seconds)	Specifies the time interval, in seconds, after which, if no data has been received from the data source, ICON updates the `NODATA_IUTC` field in the applicable `G_DSS_*_PROVIDER` table. The `nodata` record is a kind of heartbeat mechanism that enables Genesys Info Mart to distinguish between situations in which there is no data from a data source because there is no activity, and situations in which there is no data because the data source is unavailable.
	use-dss-monitor*	`true`	Specifies that ICON will populate the `G_DSS_*_PROVIDER` tables in IDB.
Interaction processing	calls-in-the-past*	`true` **Note:** Although this option applies only when ICON processes multimedia interactions, Genesys Info Mart requires that you set the value of this option to `true` for all types of ICON details.	Enables reporting for multimedia interactions that have started in the past. **Note:** Reporting on multimedia interactions that began while ICON was down or before ICON started raises unavoidable data-quality issues. For more information about the potential issues and about strategies for mitigating them when `calls-in-the-past=true`, see the section in the *Genesys Info Mart 8.1 Operations Guide* about restarting ICON, in the chapter about managing Interaction Concentrator and data sources.

* Mandatory option—Genesys Info Mart will not function if the option value is not set as specified.

** Recommended option—Data quality might be compromised if you do not use the recommended value. For more information about data-quality considerations, see the *Genesys Info Mart 8.1 User's Guide*.

Table 9: ICON Application Options—callconcentrator Section, by Type of Details (Continued)

Area of Functionality	Option Name	Recommended Value	Description
Interaction processing (continued)	om-force-adata*	`true` **Note:** Although this option applies only when ICON processes multimedia interactions, Genesys Info Mart requires that you set the value of this option to `true` for all types of ICON details.	Specifies that ICON will store a UserData snapshot in the `GM_F_USERDATA` table for interactions that started in the past.
IDB	partition-type*	`2` **Note:** For Genesys Info Mart purposes, this option effectively applies only to processing of multimedia interactions. However, Genesys Info Mart requires that you set the value of this option to `2` for all types of ICON details.	Specifies that, for multimedia interactions, the value of `gsys_partition` in the `G_IR` and `G_CALL` tables is taken from the `attr_itx_submitted_at` attribute of the Interaction Server event. The value populates the `CREATED_TS` field in the `GIDB_G_IR_MM` and `GIDB_G_CALL_MM` tables in the Info Mart database.
Voice or Multimedia details			
Agent state and login session	gls-active-reason-codes*	`true` **Note:** Although this option applies only when ICON processes voice or multimedia interactions, Genesys Info Mart requires that you set the value of this option to `true` for all types of ICON details.	Specifies that ICON captures active agent state reason codes, and temporarily stores the values in the `G_AGENT_STATE_RC_A` table in IDB. Genesys Info Mart requires knowledge of active reason codes to prevent inconsistencies in the `SM_RES_STATE_REASON_FACT` table when a reason code state ends after transformation of the interval in which the reason code started.

* Mandatory option—Genesys Info Mart will not function if the option value is not set as specified.

** Recommended option—Data quality might be compromised if you do not use the recommended value. For more information about data-quality considerations, see the *Genesys Info Mart 8.1 User's Guide*.

Table 9: ICON Application Options—callconcentrator Section, by Type of Details (Continued)

Area of Functionality	Option Name	Recommended Value	Description
Virtual queue	route-res-vqid-hist-enabled**	`true`	Specifies that ICON populates the `G_ROUTE_RES_VQ_HIST` table.
	vq-write-mode**	• `0`—For Voice details • `1`—For Multimedia details	Specifies when ICON creates the IDB record: • `0`—ICON creates a complete IDB record at the time that a particular association between an interaction and a virtual queue terminates, as indicated by either `EventDiverted` or `EventAbandoned`. • `1`—ICON initially creates an IDB record when a particular association between an interaction and a virtual queue starts (as indicated by `EventQueued`); after the association is terminated (as indicated by either `EventDiverted` or `EventAbandoned`), ICON updates the existing record. **Note:** If you are using virtual queues in multimedia interaction processing, it is very important to configure a `DN` object for this virtual queue under the `Switch` object for your multimedia switch in Configuration Manager. Otherwise, the ICON applications that report on Multimedia details are unable to indicate interactions that have entered, but not left, a virtual queue.

* Mandatory option—Genesys Info Mart will not function if the option value is not set as specified.

** Recommended option—Data quality might be compromised if you do not use the recommended value. For more information about data-quality considerations, see the *Genesys Info Mart 8.1 User's Guide*.

Table 9: ICON Application Options—callconcentrator Section, by Type of Details (Continued)

Area of Functionality	Option Name	Recommended Value	Description
Virtual queue (continued)	extended-route-result**	1—ICON stores extended routing results. (Required if detailed dispositions on routing from virtual queues is required for reporting.)	ICON stores extended routing results—the status of interactions distributed by Universal Routing Server (URS)—in the G_VIRTUAL_QUEUE and G_ROUTE_RESULT tables in IDB. To implement this feature, you must also set the report_targets and report_reasons URS configuration options to true.
Attached data	adata-userdata-history**	none	No value is recorded in IDB for a key that originates from the UserData TEvent attribute, but that is not included in the XML specification file that is specified by the adata-spec-name option value.
	cseq-adjustment** (ICON 8.1.000.37 and later)	2	Modifies sequence tracking for user data in the G_USERDATA_HISTORY table. With cseq-adjustment=2, ICON populates the value of the CSEQ field in the G_USERDATA_HISTORY table to match the behavior that Genesys Info Mart 8.x expects.

* Mandatory option—Genesys Info Mart will not function if the option value is not set as specified.

** Recommended option—Data quality might be compromised if you do not use the recommended value. For more information about data-quality considerations, see the *Genesys Info Mart 8.1 User's Guide*.

Table 9: ICON Application Options—callconcentrator Section, by Type of Details (Continued)

Area of Functionality	Option Name	Recommended Value	Description
Voice details only			
Agent state and login session	gls-acw-first	No recommended value, but the setting might affect reporting results: • `false` (default)—ICON associates after-call work (ACW) metrics with the voice interaction that immediately precedes the *completion* of the ACW (the last voice interaction). • `true`—ICON associates ACW metrics with the voice interaction that immediately precedes the *start* of the ACW (the first voice interaction). Subsequent voice interactions are considered to be related to ACW processing and should not interrupt measurement of ACW-related metrics.	Specifies which interaction ICON associates with ACW. When the agent logs out, changes his or her state to `Ready`, or goes `NotReady` for any reason other than to perform ACW, ICON reports the end of the current ACW state. **Note:** This option applies to all switches that ICON is configured to monitor, but the value does not override the value of the `gls-acw-first` configuration option (described on page 257) if it is configured within the `Switch` configuration object.
Attached data	adata-extensions-history**	`none`	No value is recorded in IDB for a key that originates from the `Extensions` TEvent attribute, but that is not included in the XML specification file that is specified by the `adata-spec-name` option value.

* Mandatory option—Genesys Info Mart will not function if the option value is not set as specified.

** Recommended option—Data quality might be compromised if you do not use the recommended value. For more information about data-quality considerations, see the *Genesys Info Mart 8.1 User's Guide*.

Table 9: ICON Application Options—callconcentrator Section, by Type of Details (Continued)

Area of Functionality	Option Name	Recommended Value	Description
Attached data (continued)	adata-reasons-history**	`none`	No value is recorded in IDB for a key that originates from the `Reasons` TEvent attribute, but that is not included in the XML specification file that is specified by the `adata-spec-name` option value.
Scenario recognition	dest-busy-processing**	`true`	ICON stores data in the `G_PARTY_HISTORY` table that identifies when a destination is busy.
	store-releasing-party	No recommended value	Specifies whether ICON stores data in the `G_CALL_STAT` table to identify which party released the call: • `false` (default)—ICON does not store the data. • `true`—ICON stores the data. **Note:** This ICON feature requires additional configuration on T-Server, and not all T-Servers support this feature. For more information, see the section in the *Interaction Concentrator 8.x User's Guide* about identifying who released the call.
Multimedia details only			
3rd Party Media	mcr-om-processing	1 (default)—This is the recommended setting because it enables recording of 3rd Party Media agent states and interaction data in IDB.	Specifies whether ICON stores information about 3rd Party Media interactions in IDB.

* Mandatory option—Genesys Info Mart will not function if the option value is not set as specified.

** Recommended option—Data quality might be compromised if you do not use the recommended value. For more information about data-quality considerations, see the *Genesys Info Mart 8.1 User's Guide*.

Table 9: ICON Application Options—callconcentrator Section, by Type of Details (Continued)

Area of Functionality	Option Name	Recommended Value	Description
Outbound Contact details only			
Outbound Contact metrics	gos-write-duplicate-metrics**	1	ICON writes to IDB all metrics that are related to active outbound objects, exactly as OCS provides them, without filtering out duplicate metrics. ICON identifies active outbound objects by `CampaignGUID`, `ChainGUID`, and `CallAttemptGUID`.

* Mandatory option—Genesys Info Mart will not function if the option value is not set as specified.

** Recommended option—Data quality might be compromised if you do not use the recommended value. For more information about data-quality considerations, see the *Genesys Info Mart 8.1 User's Guide*.

Controlling IDB Storage

By default, ICON stores full details about voice and multimedia interactions, as well as associated agent-related data. To customize IDB as a source of data for Genesys Info Mart, you can enable certain filtering through ICON configuration options. If you do not store in IDB the details that Genesys Info Mart does not extract, you can minimize the required storage space for IDB and improve ETL data extraction performance.

For a list of the IDB tables from which Genesys Info Mart extracts data, see Appendix C on page 441.

The following procedure describes how you can set up the ICON application to exclude storage of details that Genesys Info Mart does not use.

Procedure: Controlling IDB storage

Purpose: To exclude the data that is not required for Genesys Info Mart from being stored in IDB.

This procedure is optional.

Prerequisites

- Procedure: Configuring the ICON application, on page 163

Start of procedure

1. Review the options from the `filter-data` section, described in Table 10, for configuring ICON storage.
2. Open the ICON `Application` object.
3. Create a new section that is named `filter-data`, if it does not already exist on the `Options` tab.
4. Open the `filter-data` section.
5. Configure the ICON options that control data storage, as recommended in Table 10. To turn a filter on, so that certain details are excluded from storage in IDB, set the value of the option to `1`.

   Note: Table 10 lists only those `filter-data` options that you can set to safely exclude data from IDB. To avoid compromising Genesys Info Mart data quality, do not set any other filtering options to `1` (`true`). By default, ICON does not exclude any data from storage (all of the `filter-data` option values are set to `0`).

   For more information, see the chapter about configuration options in the *Interaction Concentrator 8.x Deployment Guide*.
6. Repeat this procedure for every IDB in your environment that stores either Voice details or Multimedia details.

Table 10: ICON Storage—Application Options—filter-data Section

Option Name	Effect of Setting the Filter (option value = 1)
acd-party-metrics	ICON does not store party metrics for distribution devices—such as ACD queues, Routing Points, and virtual routing points—in the `G_PARTY_STAT` table in IDB. Genesys Info Mart does not extract data from the `G_PARTY_STAT` table. **Note:** The `acd-party-metrics` option applies to SIP and voice interactions only.
call-history	ICON does not store information about the call history in the `G_CALL_HISTORY` table in IDB.

Table 10: ICON Storage—Application Options—filter-data Section (Continued)

Option Name	Effect of Setting the Filter (option value = 1)
gls-ivr	ICON verifies whether the DN at which an agent logs in is an Interactive Voice Response (IVR) device. If it is, ICON does not store information about this agent's activity in the following IDB tables: • G_LOGIN_SESSION • GX_SESSION_ENDPOINT • G_AGENT_STATE_HISTORY • G_AGENT_STATE_RC • G_DND_HISTORY • GS_AGENT_STAT • GS_AGENT_STAT_WM Furthermore, for parties that are associated with an IVR device, ICON does not record the agent's ID in the G_PARTY table. **Note:** See the description of the ivr option (page 263) for more information about how to configure a DN as an IVR resource. For more information about how ICON identifies an IVR, see the *Interaction Concentrator 8.x Deployment Guide*.
gls-metrics	ICON does not store information about agent states in the following IDB tables: • GS_AGENT_STAT • GS_AGENT_STAT_WM
ir-history	ICON does not store information about the interaction record history in the G_IR_HISTORY table in IDB.
observer-party	ICON does not store data about a party that has the role of Observer in an interaction. ICON collects data about every other party that is involved with the interaction and stores this information in the following IDB tables: • G_PARTY • G_PARTY_HISTORY • G_PARTY_STAT **Note:** Genesys Info Mart does not process observer parties. Starting with release 8.1.2, Genesys Info Mart ignores observer-party data if it is present in IDB. However, in earlier releases, Genesys Info Mart does process observer party data if the data is in IDB, and this can create data problems in certain scenarios. Therefore, in 8.1.0 and 8.1.1 deployments, Genesys strongly recommends that you filter out observer parties by setting observer-party=1.

Table 10: ICON Storage—Application Options—filter-data Section (Continued)

Option Name	Effect of Setting the Filter (option value = 1)
udata-history-terminated	ICON does not insert new records in the following IDB tables, at interaction termination time: • G_USERDATA_HISTORY • G_SECURE_USERDATA_HISTORY However, ICON does continue to write information about the creation, addition, and removal of KVPs to these tables.

End of procedure

Next Steps

- (Optional, for Voice or Multimedia details) Customizing Your ICON Attached Data Specification File
- Procedure: Preparing IDBs, on page 179

Customizing Your ICON Attached Data Specification File

When applications attach KVPs to interactions, ICON records them in the appropriate IDB table. When you deploy the ICON application, you configure certain ICON options and you also create an XML-based specification file, to indicate which KVPs ICON should store and in which IDB tables and columns they should be stored.

Genesys Info Mart ships an attached data specification file (`ccon_adata_spec_GIM_example.xml`), which specifies the KVPs that are related to Genesys Info Mart functionality. You modify the file to specify additional custom KVPs that you want ICON to store. Starting with release 8.1.2, in Oracle and Microsoft SQL Server deployments, you can use the User Data Assistant, `User_Data_Assistant.xlsm`, to help prepare the customized attached data specification.

When you install Genesys Info Mart, the `ccon_adata_spec_GIM_example.xml` and `User_Data_Assistant.xlsm` files are copied to the `sql_scripts` folder in the installation directory. They overwrite any files that have the same name.

The `ccon_adata_spec_GIM_example.xml` and the `User_Data_Assistant.xlsm` files are also available in the `sql_scripts` folder on the Genesys Info Mart CD.

Use the following general procedure to customize the attached data specification for your deployment.

Procedure: Customizing your ICON attached data specification file

Purpose: To customize your ICON attached data specification file to indicate the KVPs that you want ICON to store and in which IDB tables and columns they should be stored.

For a detailed description of the attached data KVPs that Genesys Info Mart recognizes, see Table 7 on page 122.

Prerequisites

- The `ccon_adata_spec_GIM_example.xml` file or, for release 8.1.2 and later, the `User_Data_Assistant.xlsm` file is available. To obtain the files, do one of the following:
  - Install Genesys Info Mart, and locate the required file(s) in the `sql_scripts` folder in the installation directory.
  - Locate the required file(s) in the `sql_scripts` folder on the Genesys Info Mart CD.
- If you plan to use the User Data Assistant, you have Microsoft Excel 2007 or later, with macros enabled.

Start of procedure

1. Modify the ICON attached-data specification file to include the predefined and custom KVPs that are required for your deployment. Do one of the following:
   - (For Oracle and Microsoft SQL Server deployments) Use the User Data Assistant to automatically generate the customized attached-data specification file. For more information, see the instructions in the tool.
   - Prepare the attached-data specification file manually:
     i. Edit the `ccon_adata_spec_GIM_example.xml` file to include KVP names for the additional, custom attached data elements that you want the ETL to extract.
     ii. Comment out the attached data elements that you do not want the ETL to extract.
2. Copy this file to your ICON installation directory on the ICON host that stores attached data (that is, where the ICON `role` contains `gud`).
3. If you did not name the file `ccon_adata_spec.xml`, update the `adata-spec-name` option in the ICON `Application` object, to point to this file.
4. Restart the ICON application so that the configuration changes take effect.

5. Repeat these steps for each ICON application from which Genesys Info Mart will extract Voice or Multimedia details (that is, where the ICON `role` contains `gud`).

End of procedure

Next Steps

- To continue preparing Interaction Concentrator, see Procedure: Preparing IDBs, on page 179.
- To continue preparing Genesys Info Mart to store user data, see Procedure: Customizing the user-data template script, on page 185. For more information, see also "Enabling Storage of User Data" on page 148.

Preparing IDBs

For each IDB in the deployment, you must run the ICON-provided SQL scripts to create the IDB after you have configured and installed ICON. Refer to the *Interaction Concentrator 8.x Deployment Guide* for the list of initialization scripts, their location, and the order in which to execute them.

Before an IDB is used by Genesys Info Mart, the IDB schema must be modified to enable the ETL jobs to work with it. During subsequent Genesys Info Mart migrations, the IDB schema might need to be updated further to work with the new Genesys Info Mart release.

Updating IDBs in Genesys Info Mart 8.1.1 and Earlier

On Initial Deployment and Subsequent Addition of IDBs

- If you are preparing IDB(s) for use with Genesys Info Mart releases earlier than 8.1.2, you must complete Procedure: Preparing IDBs, on page 179, which describes the scripts you must execute against the IDB.

On Migration

- If the IDB schema needs to be updated further when you subsequently migrate to a Genesys Info Mart release earlier than 8.1.2, you must similarly execute the update scripts. For information about determining whether you need to update the IDB schema during migration and resulting considerations about timing, see the chapter about Genesys Info Mart 8.x migration procedures in the *Genesys Migration Guide.*

Updating IDBs in Genesys Info Mart 8.1.2 and Later

On Initial Deployment

- Starting with release 8.1.2, `Job_InitializeGIM` automatically executes the required scripts to update the IDB schema(s) when Genesys Info Mart is first deployed. The job executes the scripts on all the IDBs for which there are database access points (DAPs) in the Genesys Info Mart application's connections.

On Migration

- If the IDB schema needs to be updated further when you subsequently migrate to a Genesys Info Mart 8.1.2 or later release, `Job_MigrateGIM` automatically executes the scripts. For information about identifying whether the scripts will be executed automatically during migration and resulting considerations about timing, see the chapter about Genesys Info Mart 8.x migration procedures in the *Genesys Migration Guide*.

On Subsequent Addition of IDBs

- If you add an Interaction Concentrator instance to an existing Genesys Info Mart 8.1.2 or later deployment, Genesys recommends that you manually execute the applicable update script, as described in Procedure: Preparing IDBs, before you add the new extraction DAP to the Genesys Info Mart application's connections. If you do not do so, the ETL cycle will be interrupted as Genesys Info Mart enters the migration state, as described in "IDB Schema Compatibility" on page 403, and you will have to run `Job_MigrateGIM` to execute the required script automatically.

From the point of view of executing the update, it makes no difference whether the IDB is partitioned or not. The `update_IDB_for_GIM` scripts detect if an IDB is partitioned and automatically adjust the SQL statements as required.

Procedure: Preparing IDBs

Purpose: To prepare the IDB so that the ETL jobs are able to use it.

This procedure is mandatory for Genesys Info Mart 8.1.0 and 8.1.1 deployments and optional for Genesys Info Mart 8.1.2 and later deployments.

Prerequisites

- All IDB instances that are required for your deployment have been created by using the ICON-provided SQL scripts. (Refer to the *Interaction Concentrator 8.x Deployment Guide* for the list of initialization scripts, their location, and the order in which to execute them.)
- The database access account that the ETL jobs will use to access IDB data is available and has the required user account privileges (see Table 3 on page 95).

  The user account does not have to be the same as the owner account. For more information about the rules and recommendations that pertain to database access accounts for Genesys Info Mart, see "Database Object Owners and User IDs" on page 93.

Start of procedure

1. For each IDB from which Genesys Info Mart will extract ICON details, log in to IDB using the database access account that you used to create the IDB. Refer to the "Installation Worksheets" beginning on page 421 to determine the ID to use.
2. Run the Genesys Info Mart–provided SQL script to add to IDB the views and indexes that Genesys Info Mart requires. The scripts are provided in the `sql-scripts` folder in your Genesys Info Mart 8.1 installation package.
   - For a Voice details IDB, use `update_idb_for_gim.sql`.
   - For a Multimedia details IDB, use `update_idb_for_gim_mm.sql`.
   - For a Configuration details or an Outbound Contact details IDB, use either `update_idb_for_gim.sql` or `update_idb_for_gim_mm.sql`.

Note: If you do not run the applicable script, Genesys Info Mart will be unable to extract data from the IDB.

End of procedure

Next Steps

- Continue on to "Preparing Genesys Info Mart Database" on page 181.

Chapter

10 Preparing Genesys Info Mart Database

This chapter describes how to prepare the target Genesys Info Mart database and predefined views. It describes how to modify and run the SQL scripts needed to create the Info Mart database and views. This chapter also describes how to tune your relational database management system (RDBMS) server for optimal performance.

This chapter contains the following sections:

The Genesys-supplied SQL scripts are provided in the `sql-scripts` folder in your Genesys Info Mart 8.1 installation package. They are also available as a separate SQL Scripts installation package. Use your database-specific tool (for example, SQL *Plus) to run the supplied SQL scripts.

Note: The Genesys Info Mart–provided SQL scripts do not qualify database objects by their schema or owner. When you run the SQL scripts, make sure that you use the ID of the schema or owner when you log in to the database. (You noted the schema or owner ID and password of each database in the appropriate section of the "Installation Worksheets" on page 421.)

Overview

As Genesys Info Mart extracts source data and transforms and loads this data into the target tables, it operates with a single Genesys Info Mart database. The Info Mart database consists of the target fact and dimension tables (Dimensional model), Global Interaction Database (GIDB) tables, Merge tables (used for voice interactions only), Control tables, Staging tables, and Temporary tables.

The following sections describe how to prepare the Genesys Info Mart database, including:

- The Info Mart database itself
- The Info Mart database read-only views

Perform the following steps for the target database:

1. Set up database access accounts and privileges.
2. Create the database schema.

The performance of the Genesys Info Mart jobs is greatly affected by the performance of the RDBMS server. Before you create the target database, tune your RDBMS server for optimal performance (see “Database Tuning” on page 200).

Note: The Genesys-provided SQL scripts create objects without specifying tablespaces or storage parameters. Work with your database administrator or data-warehousing specialist to develop a database implementation that is optimal for your environment, and make the necessary changes to the SQL scripts. See Chapter 4 on page 89 for more information.

Task Flow for Preparing Genesys Info Mart Database

Task Summary: Preparing the Genesys Info Mart Database summarizes the task flow to prepare the Genesys Info Mart database.

Task Summary: Preparing the Genesys Info Mart Database

Objective	Related Procedures and Actions
1. Set up RDBMS resources to support the reporting database.	Create the database instance. For more information, refer to the documentation for your RDBMS. **Notes:** • For Microsoft SQL Server, you must use a case-insensitive collation for the SQL Server instance and the Info Mart database. • For PostgreSQL, Genesys recommends that you use lower case for all schema names. If you cannot do so, specify the value in case-sensitive mode by surrounding the value with a set of double quotes. See also the important Notes on page 90.
2. Create and configure the Info Mart database schema to process and store detailed reporting data.	Prepare the databases by running Genesys-provided scripts: • For reporting environments that require user-data storage, customize the user-data template script. See Procedure: Customizing the user-data template script, on page 185. • For all deployments, create the Info Mart database schema. For more information, see Procedure: Preparing the Info Mart database, on page 191. • For deployments that include Voice details, customize the `GSYS_DNPREMOTELOCATION` table in GIDB, as required, to optimize performance of the merge operation. For more information, see Procedure: Configuring Info Mart database for merge, on page 193.
3. (Optional, but recommended) Prepare database links, to streamline performance.	Configure links between the IDB and Info Mart RDBMS servers. For more information, see "Using Database Links" on page 194.
4. Tune the Info Mart database.	To improve ETL execution time, adjust the settings for the target database, as applicable to your RDBMS environment: • For Microsoft SQL Server, Procedure: Tuning the Info Mart database on Microsoft SQL Server, on page 200 • For Oracle, Procedure: Tuning the Info Mart database on Oracle, on page 202 • For PostgreSQL, Procedure: Tuning the Info Mart database on PostgreSQL, on page 204

Task Summary: Preparing the Genesys Info Mart Database (Continued)

Objective	Related Procedures and Actions
5. After you have installed Genesys Info Mart, complete database preparation.	• If applicable, create read-only, tenant-specific views. For more information, see Procedure: Creating Genesys Info Mart tenant views—Oracle, on page 351. • (Optional) Create custom calendar dimensions. For more information, see Procedure: Configuring custom calendars, on page 347. **Note:** If you plan to employ custom calendars in a deployment that includes Genesys Interactive Insights (GI2) or the Reporting and Analytics Aggregates (RAA) package, configure the custom calendars before you start aggregation.

Customizing the User Data Template

Genesys Info Mart 8.1 provides two SQL scripts as a template for modifications you can make to the Info Mart database schema to customize user-data reporting:

- `make_gim_UDE_template.sql`—For use with nonpartitioned databases
- `make_gim_UDE_template_partitioned.sql`—For use with partitioned databases

You modify the applicable user-data template script, as required, to create custom user-data extension tables and columns and to specify storage of custom KVPs. Starting with release 8.1.2, in Oracle and Microsoft SQL Server deployments, you can use the User Data Assistant, `User_Data_Assistant.xlsm`, to help prepare the customized script.

Figure 8 shows the relationships that the user-data template script creates, to illustrate how custom user data is stored in the Info Mart database and populated in interaction records. Inclusion of the `MEDIATION_SEGMENT_FACT` (MSF) table in Figure 8 illustrates optional storage of user data for interactions that are in mediation, a feature that was introduced in release 8.1.2.

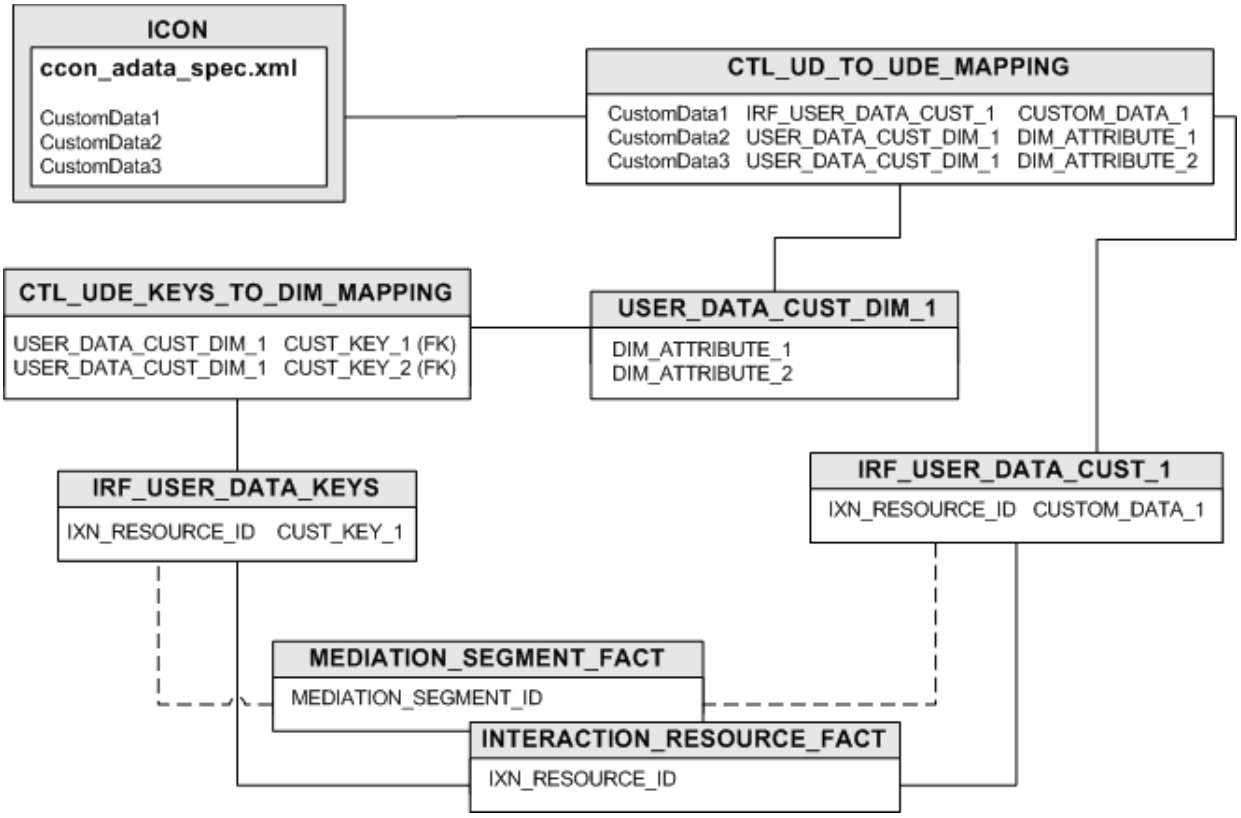

Figure 8: Custom User Data Storage

You can define the names of the custom tables and columns as you choose to see them in the Info Mart database. In custom fact tables, you can also specify the data types—character, numeric, or date/time—for the columns that store KVP values.

To make the best use of the flexible user-data storage that Genesys Info Mart release 8.x provides, Genesys recommends that you use table and column names that reflect the meaning of the user-data KVPs in your deployment. Meaningful names of columns in recognizable user-data extension tables makes it easier to write unambiguous reporting queries.

For more information about planning user-data storage in your deployment, see “Storing User Data” on page 115.

Procedure: Customizing the user-data template script

Purpose: To customize the Genesys-provided user-data script in order to specify user-defined KVP names and define custom user-data extension tables.

You can modify the user-data script and use it to customize the user-data tables in the Info Mart database schema at any time.

Even if you use the User Data Assistant to automatically draft the customized script in 8.1.2 or later deployments, Genesys recommends that you review the information in the following manual procedure, so that you can verify the validity of the generated script before you execute it.

Prerequisites

- The worksheet for mapping user-data keys that are used for reporting in your environment is complete. For the mapping worksheet, see page 428. Alternatively, if you are using the User Data Assistant, the `Business Analyst` and `Report Developer` tabs have been completed.
- The ICON application has been configured to store the required user-data KVPs.

Start of procedure

1. Locate a copy of the applicable template script (`make_gim_UDE_template.sql` or `make_gim_UDE_template_partitioned.sql`) in the RDBMS-specific scripts folder on the Genesys Info Mart 8.1 product CD.
2. Save a copy of the script to a local machine.

Create a User Data Fact Table

3. Modify your copy of the script to provide columns in a user-data fact table to store high-cardinality KVPs that you will use in your reports. By default, the script creates a table named `IRF_USER_DATA_CUST_1`.

   You can change the name of the `IRF_USER_DATA_CUST_1` table to any name that you want to see in the Info Mart database. However, if you change the name, ensure that you change all instances in the script, including the parts of the script that are described in Steps 11 and 4.

   - If you are modifying the script to prepare for the initial deployment, simply replace default names for the columns that store KVP values, such as `CUSTOM_DATA_1`, with names that are more meaningful in your deployment. Genesys recommends that you use the actual names of the high-cardinality KVPs.
   - If you want to store particular KVP values as numeric data, change the data type for the columns that store those KVP values to any numeric data type that is supported by your RDBMS. The template script includes examples of columns with numeric data types.
   - (For release 8.1.201 and later) If you want to store particular KVP values as date/time data, change the data type for the columns that store those KVP values to one of the following data types:
     - For Oracle, `DATE` or `TIMESTAMP`
     - For Microsoft SQL Server, `DATETIME`

— For PostgreSQL, `TIMESTAMP`

The template script includes examples of columns with date/time data types.

If the date/time that you want to store is in the Genesys Info Mart default format for date/time (`yyyy-mm-ddThh24:mi:ss.ff`), you do not need to perform any further mapping. If the date/time is in another date format, you must specify the conversion expression when you map the KVP to the fact table column (see Step 12 on page 189).

- If you are modifying the script to update the database schema after Genesys Info Mart has already been deployed, you must:
  — Delete the `DROP TABLE` SQL statement, which appears in the template script before the `CREATE TABLE` statement.
  — Change the `CREATE TABLE` SQL statement to an `ALTER TABLE` one.

Create an Index on the User Data Fact Table

4. Modify the script, as required, to create an index for the user-data fact table that you created in Step 3 on page 186.

Create a User Data Dimension Table

5. Continue modifying your copy of the script to provide columns in a user-data dimension table to store low-cardinality KVPs that you will use in your reports. By default, the script creates a table named `USER_DATA_CUST_DIM_1`.

   You can change the name of the `USER_DATA_CUST_DIM_1` table to any name that you want to see in the Info Mart database. However, if you change the name, ensure that you change all instances in the script, including the parts of the script described in Step 8 on page 188 through Step 11.

   - If you are modifying the script to prepare for the initial deployment, simply replace default column names, such as `DIM_ATTRIBUTE_1`, with names that are more meaningful in your deployment. Genesys recommends that you use the actual names of the low-cardinality KVPs.
   - If you are modifying the script to update the database schema after Genesys Info Mart has already been deployed, you must:
     — Delete the `DROP TABLE` SQL statement, which appears in the template script before the `CREATE TABLE` statement.
     — Change the `CREATE TABLE` SQL statement to an `ALTER TABLE` one.

Warning! Do not modify the data types for the fields. Genesys Info Mart does not support numerical data types for user-data dimensions.

Create an Index on the User Data Dimension Table

6. Modify the script, as required, to create an index for the user-data dimension table that you created in Step 5.

Create Additional User Data Fact and Dimension Tables

7. If necessary, repeat Steps 3 through 6 to add SQL commands to create additional custom user-data fact and dimension tables.

Add Foreign Key References in the IRF_USER_DATA_KEYS Table

8. Modify the script, as required, to create foreign key reference(s) for the user-data dimension table(s) in the IRF_USER_DATA_KEYS table.

   The script includes the following placeholder:

   - CUSTOM_KEY_1—The name of the foreign key that Genesys Info Mart will use to reference the user-data dimension table that you created in Step 5 on page 187. Genesys recommends that you use a key name that provides an obvious association with the table name. You map this key to the referenced table later (Step 10).

   Warning! Do not change the data type, the mandatory status, or the default value of the fields that you add. (The default value -2 indicates NO_VALUE.)

   Note: Adding columns to a big IRF_USER_DATA_KEYS table can consume significant DBMS resources and time. If you are modifying the script to prepare for the initial deployment, consider adding redundant columns in advance. Later, you can map new user-data dimensions to existing IRF_USER_DATA_KEYS columns, as required.

Add Initial Records to the User Data Dimension Table

9. For the user-data dimension table(s) that you created in Step 5 on page 187, modify the script, as required, to populate the table(s) with mandatory values for predefined keys (for example, UNKNOWN). By default, the script inserts the required values into a table named USER_DATA_CUST_DIM_1.

Map the User Data Dimension Table to the Foreign Key

10. Modify the script to add to the CTL_UDE_KEYS_TO_DIM_MAPPING table the mapping between the user-data dimension table(s) and the foreign key(s) that you added to the IRF_USER_DATA_KEYS table (Step 8). The script includes the following placeholders:

    - USER_DATA_CUST_DIM_1—The user-data dimension table name (which you defined in Step 5 on page 187)
    - ID—The primary key for the user-data dimension table
    - CUSTOM_KEY_1—The foreign key for the user-data dimension table (which you specified in the IRF_USER_DATA_KEYS table in Step 8)

Map User Data Keys to User Data Fact and Dimension Table Columns

11. For each column that you defined for user-data fact and dimension tables (see Steps 3 through 7 on page 186), modify the script to add to the CTL_UD_TO_UDE_MAPPING table the mapping between user-data keys and the user-data table columns, to specify default values, and, for custom date conversion in user-data fact tables, to specify the custom conversion

expression. Use the worksheet that you prepared for user-data mapping (see page 428) to identify the required script changes. To customize the date conversion expression, see Step 12.

The script includes the following placeholders for user-data fact tables:

- `CustomData`*N*—The key name (as stored by ICON)
- `IRF_USER_DATA_CUST_1`—The user-data fact table name (which you defined in Step 3 on page 186)
- `CUSTOM_DATA_`*N*—The column name (which you defined in Step 3 on page 186)

The script includes the following placeholders for user-data dimension tables:

- `CustomAttribute`*N*—The key name (as stored by ICON)
- `USER_DATA_CUST_DIM_1`—The user-data dimension table name (which you defined in Step 5 on page 187)
- `DIM_ATTRIBUTE_`*N*—The column name (which you defined in Step 5)

The script also requires you to specify the propagation rule, default value, and activity status for each KVP. For more information about values for these fields, see the column descriptions for the `CTL_UD_TO_UDE_MAPPING` table in the *Genesys Info Mart 8.1 Reference Manual* for your RDBMS.

Ensure that the default values that you specify are consistent with the data type for the column.

12. If you want Genesys Info Mart to store a date/time value expressed in a format other than the Genesys Info Mart default format for date/time (for example, `DD Mon YY` instead of `yyyy-mm-ddThh24:mi:ss.ff`), specify the conversion expression in the `CONVERT_EXPRESSION` field in the `CTL_UD_TO_UDE_MAPPING` table entry for the KVP. Genesys Info Mart includes the conversion expression in SQL statements to convert the data. The template script provides an example of the syntax to use.

Note: The example in the template script expresses the date/time in the Genesys Info Mart default format for date/time (`yyyy-mm-ddThh24:mi:ss.ff`) and is therefore an unnecessary entry in the mapping table.

- **For Oracle:**

  ```
  TO_DATE(${},'yyyy-mm-dd'T'hh24:mi:ss')
  ```

  or

  ```
  TO_TIMESTAMP(${},'yyyy-mm-dd'T'hh24:mi:ss.ff')
  ```

  where `${}` is a placeholder for the KVP value to be converted.

  To customize the conversion, use one of the standard RDBMS-provided functions (`TO_DATE` or `TO_TIMESTAMP`), replacing the date-format expression with a valid format, as defined in the RDBMS documentation.

- **For Microsoft SQL Server:**
  `${schema}.GIM_TO_TIMESTAMP_ISO8601(${})`
  where:
  — `${schema}` is a placeholder for the Info Mart database schema name; Genesys Info Mart gets the value of the `${schema}` parameter from the `[gim-etl].default-schema` option in the Info Mart database access point (DAP).
  — `${}` is a placeholder for the KVP value to be converted.
  — `GIM_TO_TIMESTAMP_ISO8601` is an out-of-box function, which first executes another out-of-box function, `GIM_IS_ISO8601_DATE`, to check format before calling the Microsoft SQL Server system function `convert(datetime, ${}, 126)`, to convert the `datetime` expression in format `yyyy-mm-ddThh24:mi:ss.ff`.
  It is not possible to use `convert()` without first checking the format of the expression, because of a Microsoft SQL Server limitation that makes transactions unusable after most conversion errors. To customize the conversion, you must define your own conversion function in the database, and then call that function in the `CONVERT_EXPRESSION` field, using the syntax shown above. Use the `GIM_TO_TIMESTAMP_ISO8601` and `GIM_IS_ISO8601_DATE` functions, which are defined in the `make_gim.sql` and `make_gim_partitioned.sql` scripts, as examples for your custom functions, in conjunction with RDBMS documentation about syntax requirements and date formats. Do not modify the out-of-box functions.
- **For PostgreSQL:**
  `TO_TIMESTAMP(${},'yyyy-mm-dd'T'hh24:mi:ss.ms')`
  where `${}` is a placeholder for the KVP value to be converted.

Save the Script

13. Save the modified copy of the script.

End of procedure

Next Steps

- Execute the modified script when you create the Info Mart database schema (see Step 4 on page 192) or as required to update an existing database schema.
  To verify correct mappings, execute the following SQL command against the Info Mart database:
  `SELECT * FROM CTL_UD_TO_UDE_MAPPING`

Compare the results of the query against the mapping that you prepared before customizing the script, as described in the Prerequisites on page 203.

Note: Before you execute the script to update an existing database schema, Genesys recommends that you back up the Info Mart database.

- If you update an existing database schema that uses read-only tenant views to access Info Mart data for reports, you must re-create the read-only tenant views. For more information, see Procedure: Creating Genesys Info Mart tenant views—Oracle, on page 351.

Preparing the Info Mart Database

The RDBMS-specific SQL scripts that are provided with Genesys Info Mart create the Info Mart database schema. This includes merge tables, which are a required part of the Info Mart database schema in any deployment. Genesys Info Mart provides separate scripts for partitioned and nonpartitioned database schemas.

The Genesys Info Mart database scripts do not create the additional database objects that are required to support aggregation. For more information about database preparation for deployments that use Genesys Interactive Insights (GI2) or the separately installed Reporting and Analytics Aggregates (RAA) package, see the *Reporting and Analytics Aggregates 8.x Deployment Guide* and the *Genesys Interactive Insights 8.x Deployment Guide.*

Procedure: Preparing the Info Mart database

Purpose: To create the data schema for the Info Mart database.

Prerequisites

- Create a database instance for your Info Mart database.
- (Optional) If you plan to store user-defined attached data, customize the applicable user-data SQL script template, as instructed in Procedure: Customizing the user-data template script, on page 185.
  - For a nonpartitioned database, use `make_gim_UDE_template.sql`.
  - For a partitioned database, use `make_gim_UDE_template_partitioned.sql`.

- (Optional) If you plan to create multiple calendar dimensions to support your reporting, customize the `make_gim.sql` or `make_gim_partitioned.sql` script to create additional calendar tables. Alternatively, create the custom calendars after you have installed Genesys Info Mart. For more information, see Procedure: Configuring custom calendars, on page 347.

Start of procedure

1. Ensure that the database access account that you use to create the Info Mart database schema is available and has the required privileges (see Table 3 on page 95).

   Refer to the "Installation Worksheets" beginning on page 421 to determine the ID to use.
2. Log in to the Info Mart database, using the Info Mart user account.
3. Run the applicable SQL script to create the Info Mart database schema:
   - For a nonpartitioned database, use `make_gim.sql`. This script creates the Genesys Info Mart dimension and fact tables and related indexes.
   - For a partitioned database, use `make_gim_partitioned.sql`. This script creates the Genesys Info Mart dimension and fact tables and related indexes. For the tables and indexes that are partitioned, this script creates a single, outdated partition that is expected to be purged during the first run of the maintenance job.
4. (Optional) Run the modified `make_gim_UDE_template.sql` or `make_gim_UDE_template_partitioned.sql` SQL script that you have updated with required key-value pair (KVP) names. This script creates extension tables in the Info Mart database schema to store custom user-data, configures user-data mappings, and adds the specified dimension key fields to the `IRF_USER_DATA_KEYS` table.
5. Ensure that the database access account that the ETL jobs will use to access the Info Mart database is available and has the required user account privileges (see Table 3 on page 95).

   Refer to the "Installation Worksheets" beginning on page 421 to determine the ID to use.

   The user account does not have to be the same as the owner account. For more information about the rules and recommendations that pertain to database access accounts for Genesys Info Mart, see "Database Object Owners and User IDs" on page 93.
6. Update the `GSYS_DNPREMOTELOCATION` tables, as required, to optimize performance of the merge operation. For more information, see Procedure: Configuring Info Mart database for merge, on page 193.

End of procedure

Next Steps

- (Required for Voice details only) Procedure: Configuring Info Mart database for merge, on page 193.
- (Recommended, but optional) "Using Database Links" on page 194.
- Tune up your Info Mart database, as appropriate for your RDBMS environment. For more information, see "Database Tuning" on page 200.

Procedure: Configuring Info Mart database for merge

Purpose: To optimize performance of the merge operation.

Start of procedure

1. If any switches are not monitored by ICON, store those `Switch` object names in the `GSYS_DNPREMOTELOCATION` table of the Info Mart schema. Otherwise, merging of some interswitch voice interactions will be delayed until the configured IS-Link timeout occurs, and this delays transformation of those voice interactions.

   For example, suppose that you have four switches and two ICON instances:

   - `ICON1` monitors switch `SITE1_sw1`.
   - `ICON2` monitors switch `SITE2_sw2`.
   - `SITE3_sw3` and `SITE4_sw4` are not monitored by either ICON instance.

   To avoid delays in merging, add the following two records to the `GSYS_DNPREMOTELOCATION` table in the Info Mart schema:

   - `GSYS_DNPREMOTELOCATION.REMOTELOCATION=SITE3_sw3`
   - `GSYS_DNPREMOTELOCATION.REMOTELOCATION=SITE4_sw4`

   For each row in the `GSYS_DNPREMOTELOCATION` table, populate the `ID` field with a unique value.

2. Review the settings of the `max-call-duration`, `merge-chunk-size`, and `merge-failed-is-link-timeout` configuration options (see pages 290, 296, and 297), and modify them as required for your deployment.

   For more information about the role of `max-call-duration` in determining the IS-Link timeout, see "Merge Operation Stuck Threshold" on page 387.

End of procedure

Next Steps

- (Recommended, but optional) "Using Database Links" on page 194.
- Tune up your Info Mart database as appropriate for your RDBMS environment. For more information, see "Customizing the User Data Template" on page 184.

Using Database Links

To improve performance of the extraction job, Genesys recommends that you utilize RDBMS functionality to configure links between the RDBMS servers that host the IDB schemas and the RDBMS server that hosts the Info Mart database schema. If database links are not configured, the extraction job first moves data from IDB into Genesys Info Mart memory, and then from Genesys Info Mart memory into the Info Mart database. When database links are configured, the extraction job can copy data directly from the IDB into the Info Mart database.

Note: Genesys Info Mart does not support database links on PostgreSQL.

This section provides the following information:

Creating Database Links—Oracle

Note: In Oracle Real Application Cluster (RAC) deployments with database links, node synchronization activity is significant and can lead to performance issues. Genesys recommends that you do not use database links if you are using Oracle RAC.

The following command configures a link between Oracle databases:

```
CREATE PUBLIC DATABASE LINK <Link Name>
connect to <IDB User Name>
identified by <IDB User Password>
USING '<IDB Connect_String>';
```

where:

- `<Link Name>` is any valid database link name.
- `<IDB User Name>` is the `User Name` that is specified on the `DB Info` tab of the DAP that enables connection to the IDB.

- `<IDB User Password>` is the `Password` that is specified on the `DB Info` tab of the DAP that enables connection to the IDB.
- `<IDB Connect_String>` is the service name of the remote IDB database. For Oracle, this is the Transparent Network Substrate (TNS) name.

> **Note:** If you specify only the database name, Oracle implicitly appends the database domain to the connect string to create a complete service name. Therefore, particularly if the IDB and Info Mart databases are not in the same database domain, ensure that you specify the complete service name.

Execute the command on the server that hosts the Info Mart database schema. Repeat the command for each IDB from which Genesys Info Mart might need to extract data, using database connection information that matches the DAP that Genesys Info Mart uses to extract data from that IDB.

Figure 9 and Figure 10 on page 196 illustrate the relationship between the parameters in the database-link command and the properties of the Genesys Info Mart extraction DAP, for non-JDBC and JDBC DAPs, respectively.

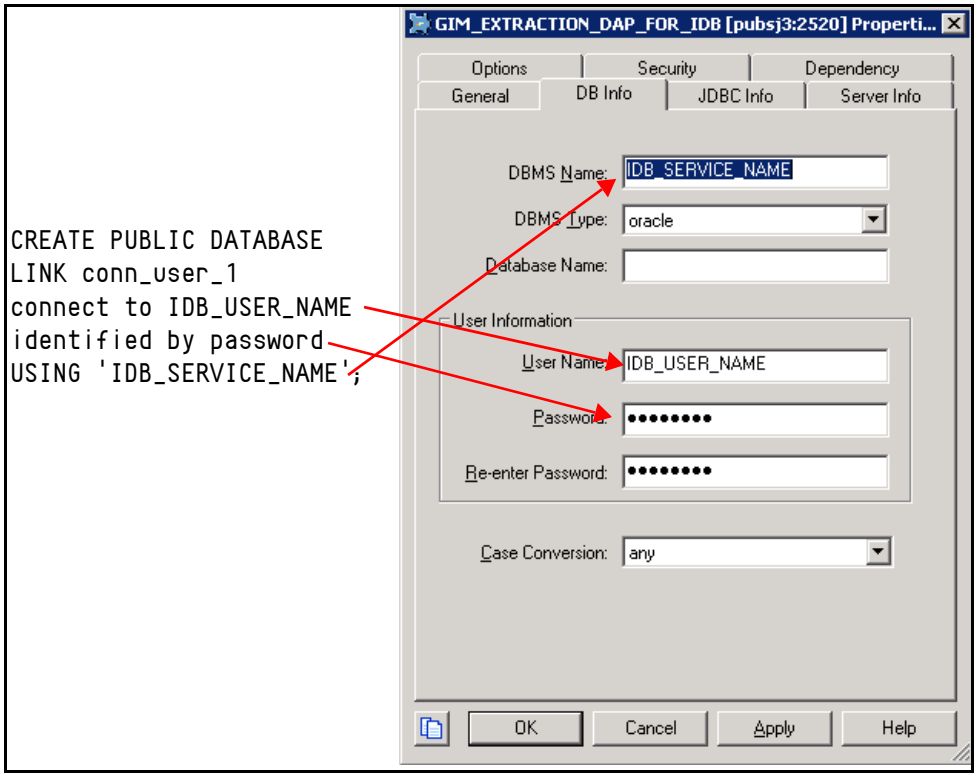

Figure 9: Oracle Database Links Configuration Example—Non-JDBC DAP

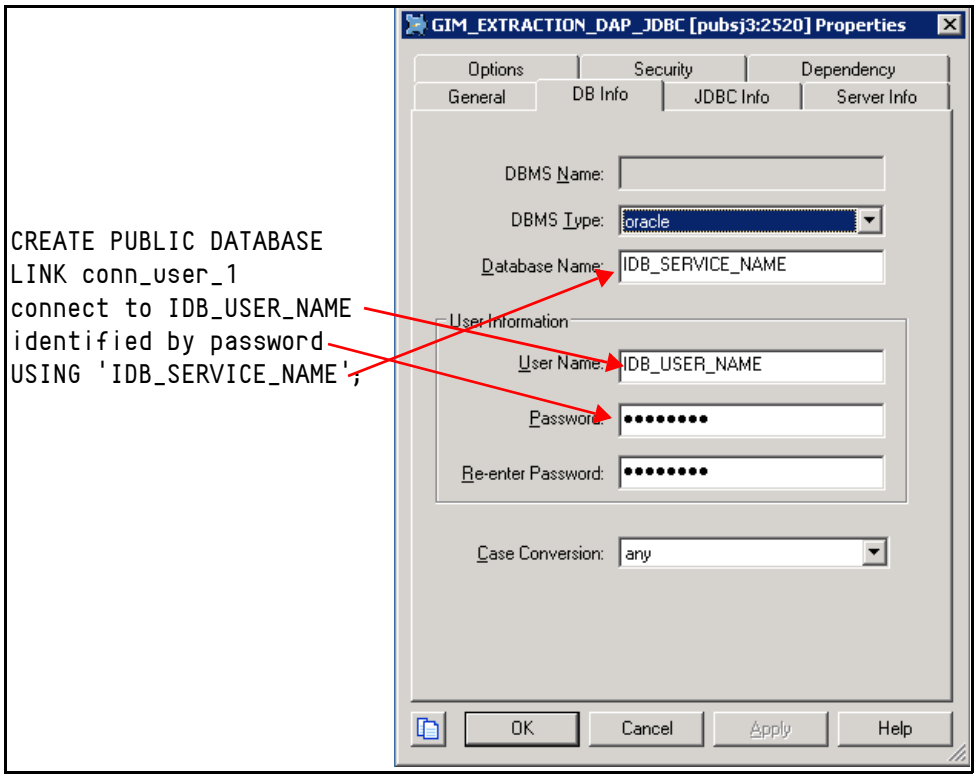

Figure 10: Oracle Database Links Configuration Example—JDBC DAP

Verifying Database Links

Verifying Connection Information

After the database links have been created, Genesys recommends that you verify that Oracle has stored the connection strings correctly. To do so, execute the following statement:

```
SELECT * FROM ALL_DB_LINKS
```

Then verify the connection information that is stored in the `ALL_DB_LINKS.HOST` column.

Verifying Use of Database Links

In release 8.1.0 and earlier, reviewing the logs at debug level enables you to verify that Genesys Info Mart is using the configured database links. During runtime, database links will be used in the SQL statements that Genesys Info Mart executes during extraction. For example:

```
INSERT INTO GIM-SCHEMA...
SELECT...FROM IDB-SCHEMA...@DB-LINK-NAME
```

where `DB-LINK-NAME` is the name of the database link you created. The SQL queries are reflected in the logs at debug level.

To confirm that Genesys Info Mart is actually using the links, search the debug-level log for the name(s) of the database link(s) that you created.

Starting with release 8.1.1, you can review the Genesys Info Mart logs at standard level, to verify the use of database links. For more information, see "Verifying Use of Database Links" on page 200.

Note: If you use database links in an Oracle deployment, ensure that you configure the limit for open links to provide sufficient connections for concurrent processing. Genesys recommends that, at a minimum, you set the value of the `open_links` Oracle initialization parameter to the value of the Genesys Info Mart `extract-data-thread-pool-size` configuration option (see page 289).

Creating Linked Servers—Microsoft SQL Server

The following procedure describes how to use SQL Server Management Studio to configure links between Microsoft SQL Server databases.

Procedure: Creating Linked Servers for Microsoft SQL Server

Start of procedure

1. Log in to the SQL Server Management Studio, using administrator credentials.
2. Connect `Object Explorer` to the Info Mart database.
3. Open `Object Explorer` and select `Server Objects > Linked Servers`.
4. Follow Microsoft instructions to create Linked Servers that link to the IDBs from which Genesys Info Mart will extract data.

   For illustrations of the elements in the DAP configuration that are used in the Linked Servers definition, see "Related DAP Configuration Example".
5. Log in to the Info Mart database, using the Info Mart user ID.
6. In SQL Server Management Studio, verify that you can access IDB tables from the Linked Servers.

End of procedure

Related DAP Configuration Example

Figure 11 and Figure 12 on page 199 illustrate the relationship between the components that are referenced in the Linked Servers definition and the properties of the Genesys Info Mart extraction DAP, for non-JDBC and JDBC DAPs, respectively.

For full information about configuring the DAPs that Genesys Info Mart requires, see Chapter 11 on page 207.

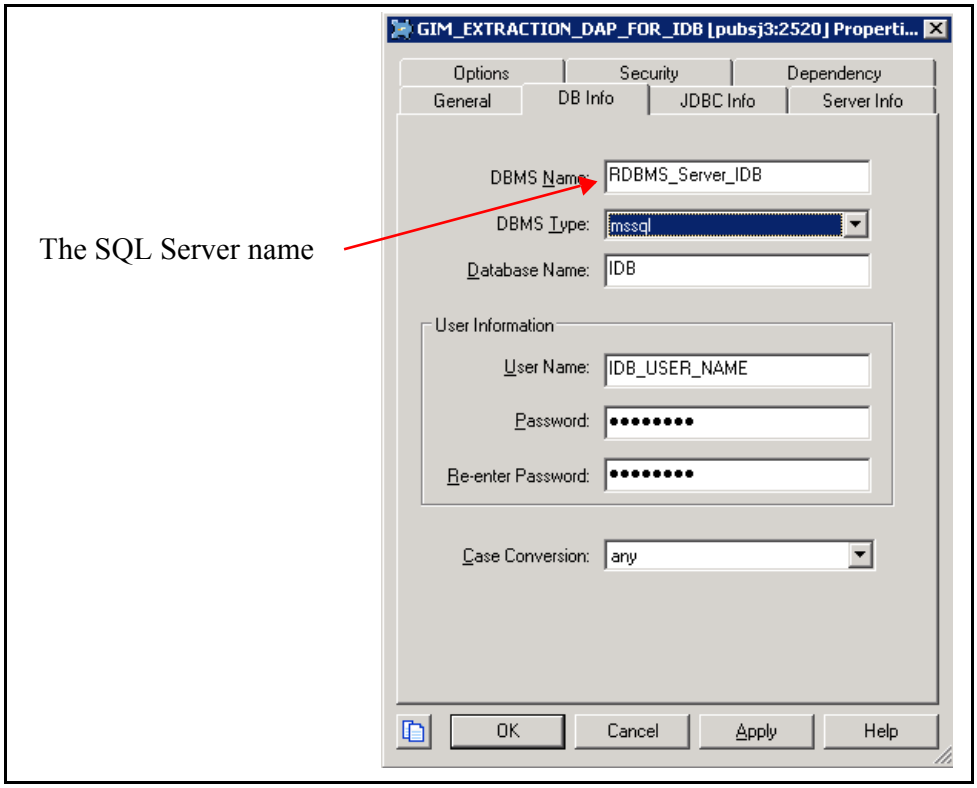

Figure 11: Linked Servers Configuration Example—Non-JDBC DAP

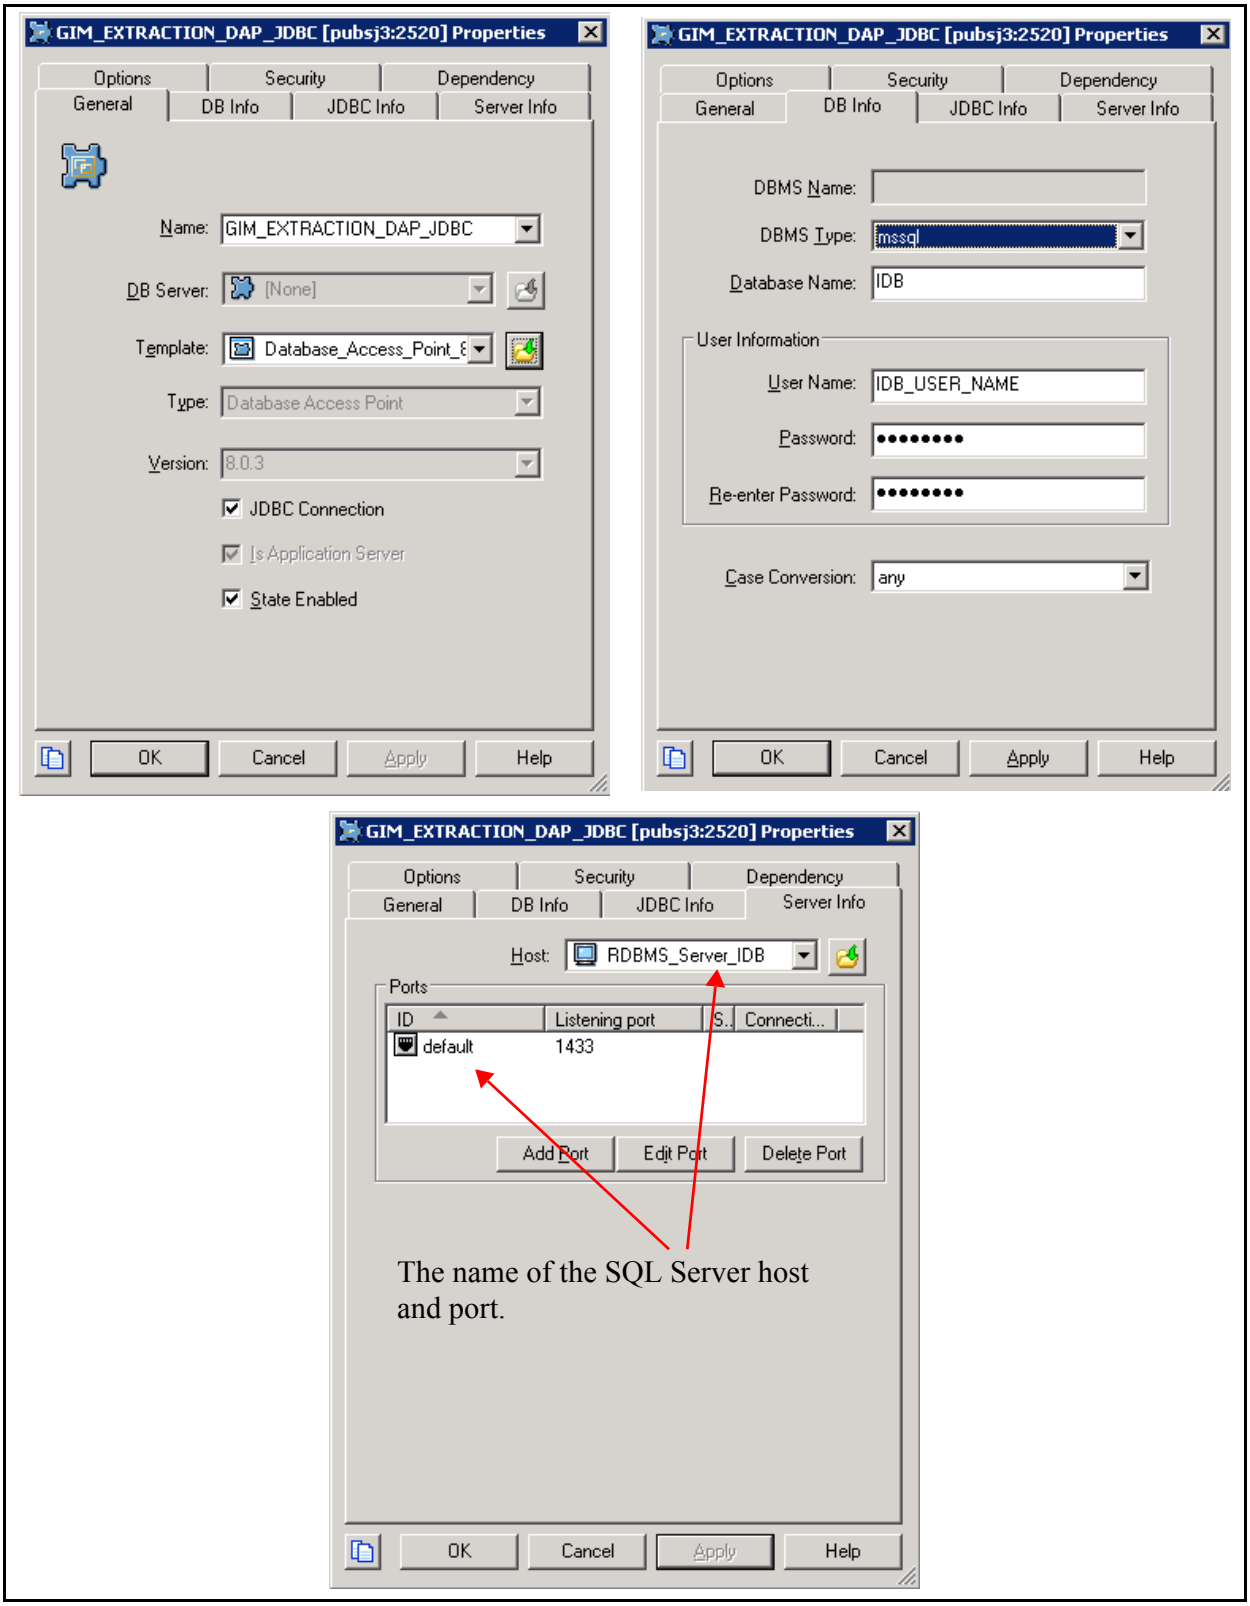

Figure 12: Linked Servers Configuration Example—JDBC DAP

Verifying Use of Database Links

Starting with release 8.1.1, at the start of each extraction cycle, Genesys Info Mart logs the DAPs and associated database links that will be used during that extraction cycle. Review the logs to verify the use of configured database links during runtime.

Figure 13 shows an excerpt from a sample log.

```
... INFO ... Logging DAPs:
... INFO ... DAP=DAP_WITH_DBLINK, role=ICON_CFG, dbname=..., schema=... -> DbLink=LINKNAME
... INFO ... DAP=DAP_WITHOUT_DBLINK, role=ICON_CFG, dbname=..., schema=...
```

Figure 13: Example of Database Link in Logs

Database Tuning

This section provides recommended parameter settings and tuning guidelines that you can use to improve the Info Mart ETL execution time (see "Info Mart Database").

It also includes suggestions for your database administrator for managing the target database after it is deployed (see "Additional Considerations" on page 206).

Info Mart Database

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

- Procedure: Tuning the Info Mart database on Microsoft SQL Server, on page 200
- Procedure: Tuning the Info Mart database on Oracle, on page 202
- Procedure: Tuning the Info Mart database on PostgreSQL, on page 204

Procedure: Tuning the Info Mart database on Microsoft SQL Server

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

Prerequisites

- Create the Info Mart database schema, as instructed in Procedure: Preparing the Info Mart database, on page 191.
- Use Microsoft SQL Server Enterprise Manager to review the settings of the Microsoft SQL Server properties.

Start of procedure

1. Allocate sufficient memory.

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

   a. Start with the total real memory on the database server.

   b. Subtract the memory that is required by the operating system and any other applications.

   c. Set the maximum memory that is allocated to Microsoft SQL Server to the result.

2. Select relevant server settings.

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

3. Configure connections.

   The ETL jobs use many concurrent database connections. Set `Maximum concurrent user connections` to `0` (unlimited).

Note: Consult with your database administrator to further fine-tune these and other parameters, if you find the ETL job performance unacceptable for your environment. See also "Additional Considerations" on page 206.

4. Review the settings of the following database properties:
   - Data Files: Select `Automatically grow file` and `Unrestricted file growth`.
   - Transaction Log: Select `Automatically grow file` and `Unrestricted file growth`.
   - Options—Settings: Select `Auto update statistics` and `Auto create statistics`.

- Options—Recovery Model: Take into account the following considerations.

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

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

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

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

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

   In Microsoft SQL Server deployments, the transformation job might record duplicate rows in Info Mart fact tables because of the inability of Microsoft SQL Server to finish rollback correctly when deadlock happens during execution of rollback. To minimize locking contention, Genesys recommends that you use the `READ COMMITTED` isolation level, with the `READ_COMMITTED_SNAPSHOT` database option set to `ON`.

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

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

End of procedure

Next Steps

- Configure the required database access points (DAPs). For more information, see Chapter 11 on page 207.

Procedure: Tuning the Info Mart database on Oracle

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

Prerequisites

- Create the Info Mart database schema, as instructed in Procedure: Preparing the Info Mart database, on page 191.

Start of procedure

Initialization Parameters

1. Set the Oracle initialization parameters:
   - `filesystemio_options = ASYNCH`
   - `processes` = at least `1000`
   - `sessions` = at least `1000`
   - `open_cursors` = at least `1000`
   - (For use with database links) `open_links` = at least the value of `extract-data-thread-pool-size` (see page 289)

     If your deployment uses database links, configure the `open_links` parameter to allow a sufficient number of connections through the database links. For more information about how to configure your deployment to use database links, see "Using Database Links" on page 194.

Memory

2. Allocate sufficient memory.

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

   As a quick approximation:

   a. Start with the total real memory on the database server.
   b. Subtract the amount of memory that is required by the operating system and any other applications.
   c. Split the result between the buffer cache and the PGA aggregate target.

Listener

3. Increase the size of the listener queue for TCP/IP connection requests.

   The need to modify the listener configuration depends on deployment-specific factors such as hardware and the configured or desired degree of parallelism. See the Oracle documentation for information about the conditions under which you should increase the queue size in the listener configuration file, `listener.ora`, as well as the syntax for doing so. See also "Performance tuning" on page 251 for a summary of Genesys Info Mart configuration options that control the degree of parallelism for ETL processing.

Case-Sensitivity

4. Genesys recommends that you use case-sensitive settings for `NLS_SORT` and `NLS_COMPARE`.

Fine-Tuning

5. Consult your database administrator to further fine-tune these and other parameters, including JVM startup parameters such as `queryParallelism` (see page 360), if you find the performance of the ETL jobs in your environment unacceptable. See also "Additional Considerations".

End of procedure

Next Steps

- Configure the required database access points (DAPs). For more information, see Chapter 11 on page 207.

Procedure: Tuning the Info Mart database on PostgreSQL

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

Prerequisites

- Create the Info Mart database schema, as instructed in Procedure: Preparing the Info Mart database, on page 191.

Start of procedure

Initialization Parameters

1. Set the database server configuration parameters in the `postgresql.conf` file:
   - `checkpoint_segments = 64`
   - `synchronous_commit = off`
   - `wal_buffers = 256MB`
   - `max_connections = 1000`
   - `constraint_exclusion = partition`
   - `default_statistics_target = 100`
   - `enable_mergejoin = off`
   - `temp_buffers = 128MB`
   - `max_prepared_transactions = 1000`
   - `work_mem` = A value in the range `256MB–1GB`*
   - `maintenance_work_mem` = A value in the range `64MB–256MB`*

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

Memory

2. Allocate sufficient memory.

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

   a. Start with the total real memory on the database server.

   b. Subtract the amount of memory that is required by the operating system and any other applications. The result is the available memory for PostgreSQL.

   c. Set `effective_cache_size` to 75 percent of available memory and `shared_buffers` to 25 percent of available memory.

Autovacuum

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

   - `autovacuum = on`
   - `autovacuum_analyze_threshold = 10000`
   - `autovacuum_freeze_max_age = 1000000000`
   - `autovacuum_max_workers = 10`
   - `autovacuum_naptime = 20s`
   - `autovacuum_vacuum_cost_delay = 10ms`
   - `autovacuum_vacuum_cost_limit = 1000`
   - `autovacuum_vacuum_scale_factor = 0.3`
   - `autovacuum_vacuum_threshold = 100000`
   - `vacuum_freeze_min_age = 10000000`
   - `vacuum_freeze_table_age = 800000000`

   A Genesys Info Mart job, `Job_UpdateStats`, invokes PostgreSQL functionality to supplement the autovacuum process. Ensure that you also configure the Genesys Info Mart application to run `Job_UpdateStats` regularly. For more information, see `run-update-stats` (page 318) and `update-stats-schedule` (page 319).

Fine-Tuning

4. Consult your database administrator to further fine-tune these and other parameters if you find the performance of the ETL jobs in your environment unacceptable. See also "Performance tuning" on page 251 for a summary of Genesys Info Mart configuration options that control the degree of parallelism for ETL processing, as well as "Additional Considerations" on page 206.

End of procedure

Next Steps

- Configure the required database access points (DAPs). For more information, see Chapter 11 on page 207.

Additional Considerations

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

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

Chapter

11 Configuring DAPs

This chapter describes how to configure the database access points (DAPs) that Genesys Info Mart uses to access source and target databases. It contains the following sections:

Overview

You configure DAPs to specify the database connection parameters and other options.

You must configure DAPs to access the applicable source and target databases in your environment, in order to extract, transform, load, aggregate (if applicable), and maintain all types of data.

Required DAPs

The number and type of DAPs that you need depends on your particular deployment topology. For a description of Genesys Info Mart–supported topologies, see Chapter 3 on page 69.

At a minimum, you must configure the following DAPs:

- DAPs to access the Interaction Databases (IDBs) from which Genesys Info Mart extracts data. In this document, these DAPs are referred to as *extraction DAPs*. You must configure one extraction DAP for each IDB.
- One DAP to access the Info Mart database. In this document, this DAP is referred to as the *Info Mart DAP*.

- If you plan to use the Genesys Info Mart Administration Console for day-to-day management of Genesys Info Mart jobs, one DAP to enable the Genesys Info Mart Administration Console to access the Info Mart database. In this document, this DAP is referred to as the *Administration Console DAP*.

Note: As you configure the DAPs, refer to the applicable worksheets (see Appendix A on page 421) that you completed earlier.

Reusing DAPs

Genesys Info Mart Server and the jobs that extract, transform, and load (ETL) data use Java Database Connectivity (JDBC) to access all databases. The Genesys Info Mart Administration Console connects to the Info Mart database through a non-JDBC DAP that is associated with a DB Server, and it does not use any of the options on the `JDBC Info` tab of the DAP `Application` object. Similarly, the DAPs that the Interaction Concentrator server (ICON) uses to connect to IDB are non-JDBC DAPs that are associated with a DB Server.

To simplify your deployment, you can reuse the non-JDBC DAPs in your deployment and make these DAPs suitable for Genesys Info Mart to access the same databases. In particular, you can:

- Reuse the existing DAPs that enable ICON(s) to access IDB(s), to enable Genesys Info Mart to access the same IDBs.
- Use the same DAP to enable both Genesys Info Mart and the Genesys Info Mart Administration Console to access the Info Mart database.

Minimal additional configuration is required for you to reuse non-JDBC DAPs, if DB Server is running on the same host as the DBMS server; if the port of the DBMS listener is the default; and (for Oracle) if the TNS name and the SID name are identical. Otherwise, `jdbc-*` options, which you configure in the `[gim-etl]` section on the `Options` tab, enable you to specify the required JDBC connection parameters.

For more information about the JDBC options that you can configure on the non-JDBC DAPs to enable Genesys Info Mart to reuse these DAPs, see Table 11 on page 214.

For more information about reusing the ICON DAP for Genesys Info Mart, see Procedure: Configuring non-JDBC DAPs to access source databases, on page 213.

For more information about how to configure a single DAP for both Genesys Info Mart and the Genesys Info Mart Administration Console to use, see Procedure: Configuring non-JDBC DAPs to access source databases, on page 213.

Warning! Before you decide to reuse the same DAP for more than one purpose, consult your database administration policies. Using the same DAP for multiple purposes has implications for security and access control. For example, the assigned database user needs sufficient privileges for all databases to which a given DAP provides access.

Additional DAPs

If you have a high availability (HA) deployment topology, you must create additional DAPs. See "Enabling High Availability" on page 152 for more information.

Task Flow for Enabling Database Access

Task Summary: Enabling Access to Source and Target Databases summarizes the task flow to enable Genesys Info Mart to access source and target databases.

Note: The task summary does not include information about how to configure the JDBC connection to use Secure Socket Layer (SSL). For information about how to set up a secure JDBC connection, see your RDBMS vendor documentation.

Task Summary: Enabling Access to Source and Target Databases

Objective	Related Procedures and Actions
Enable access to the source databases from which Genesys Info Mart needs to extract data.	For each IDB from which Genesys Info Mart might need to extract data, configure a DAP that specifies the connection information and other options for Genesys Info Mart to access the IDB. For each IDB, do one of the following: • Modify the non-JDBC DAP that you configured for ICON to access the IDB. For more information, see Procedure: Configuring non-JDBC DAPs to access source databases, on page 213. • Create and configure a new JDBC DAP. For more information, see Procedure: Configuring JDBC DAPs to access source databases, on page 212. /continued...

Task Summary: Enabling Access to Source and Target Databases (Continued)

Objective	Related Procedures and Actions
Enable access to the source databases (continued).	**Note:** When you configure the DAP, the role that you must specify (in the `[gim-etl]` section) depends on the kind of data that is stored in the IDB and that Genesys Info Mart will extract: • For Configuration details, set the `role` option to `ICON_CFG`. • For Voice details, set the `role` option to `ICON_CORE`. • For Multimedia details, set the `role` option to `ICON_MM`. • For Outbound Contact details, set the `role` option to `ICON_OCS`. Ensure that the role that you specify for Genesys Info Mart (in the `[gim-etl]` section in the DAP) is consistent with the role that is specified for ICON (in the `[callconcentrator]` section in both the ICON application and DAP). For more information about the corresponding ICON roles, see “Preparing ICON” on page 162.
Enable access to the target database in which Genesys Info Mart processes and stores data.	Configure a DAP that specifies the connection information and other options for Genesys Info Mart to access the Info Mart database. Do one of the following: • Create and configure a new, non-JDBC DAP that will also be used by the Genesys Info Mart Administration Console. For more information, see Procedure: Configuring non-JDBC DAPs, on page 228. • Create and configure a new, JDBC DAP. For more information, see Procedure: Configuring JDBC DAPs, on page 225. For more information, see Procedure: Configuring a dedicated Info Mart DAP, on page 218. **Note:** When you configure the DAP, ensure that the value of the `role` option in the `[gim-etl]` section includes `INFO_MART`.

Task Summary: Enabling Access to Source and Target Databases (Continued)

Objective	Related Procedures and Actions
Enable control and monitoring of jobs through the Genesys Info Mart Administration Console. **Note:** This step is not required for Genesys Info Mart Manager.	Configure a DAP that specifies the connection information and other options for the Genesys Info Mart Administration Console to access the Info Mart database, in order to monitor ETL jobs: 1. Install DB Server to handle database requests from the Genesys Info Mart Administration Console. Refer to *Framework 8.x DB Server User's Guide* for details. 2. Configure a DAP for Genesys Info Mart Administration Console to access the Info Mart database. Do one of the following: • Modify the non-JDBC DAP that you created to enable Genesys Info Mart to access the Info Mart database for data processing and storage. • Create and configure a new, non-JDBC DAP that will be used only by the Genesys Info Mart Administration Console. For more information, see Procedure: Configuring non-JDBC DAPs, on page 228. **Note:** When you configure the DAP, ensure that the value of the `role` option in the `[gim-etl]` section includes `ADMIN_CONSOLE`. See Procedure: Configuring a dedicated Administration Console DAP, on page 224. Set the `role` option to `ADMIN_CONSOLE`.

Configuring Required DAPs

Use the procedures in this section to create and configure the required DAPs:

- Extraction DAPs, page 211
- Info Mart DAP, page 218
- Administration Console DAP, page 223

Use "DAP Configuration Supporting Procedures" on page 225 when you are referred to them during the configuration.

Extraction DAPs

A separate DAP is required for each IDB from which Genesys Info Mart will extract data. You can create and configure a new, dedicated JDBC DAP, or you can reuse the non-JDBC DAP that you have already configured for ICON to access IDB.

Use the following procedures, as applicable, to enable Genesys Info Mart to access the IDBs:

- Procedure: Configuring JDBC DAPs to access source databases

- Procedure: Configuring non-JDBC DAPs to access source databases, on page 213

Procedure: Configuring JDBC DAPs to access source databases

Purpose: To configure a new, dedicated DAP that enables Genesys Info Mart to access IDB through JDBC.

Prerequisites

- Procedure: Preparing IDBs, on page 179

Start of procedure

1. Create a JDBC DAP, as described in Procedure: Configuring JDBC DAPs, on page 225.
2. Configure the required options in the `[gim-etl]` section on the `Options` tab:
   - `role`—Specifies which data domain Genesys Info Mart will extract through the DAP:
     - For Configuration details, `role=ICON_CFG`
     - For Voice details, `role=ICON_CORE`
     - For Multimedia details, `role=ICON_MM`
     - For Outbound Contact details, `role=ICON_OCS`
   - `default-schema`—(Optional) Specifies the database schema or owner name, if it is different from the database user ID.
   - `geo-location`—(Optional) Specifies the location of the IDB.

   For information about how to create the `[gim-etl]` section and options, see Step 8 on page 227.

   For more information about the options, including requirements and restrictions for various roles, see Table 11 on page 214.
3. Repeat this procedure for each IDB from which Genesys Info Mart will extract data and for which you are not reusing an existing Interaction Concentrator DAP.

   If your deployment topology involves HA, you must also repeat this procedure to provide a DAP for each redundant IDB.

End of procedure

Next Steps

- Create the Info Mart DAP. See “Info Mart DAP” on page 218.

Procedure: Configuring non-JDBC DAPs to access source databases

Purpose: To modify the configuration of an existing DAP that ICON uses to access an IDB, to enable Genesys Info Mart to access the same IDB.

Repeat this procedure for each IDB from which Genesys Info Mart will extract data and for which you want to reuse an existing ICON DAP.

Prerequisites

- Procedure: Preparing IDBs, on page 179.
- You have the required permissions to access and modify the ICON DAP(s).

Start of procedure

1. Locate the DAP that enables ICON to access the IDB to store data for the applicable data domain.
2. On the `Options` tab, create a new configuration section, named `gim-etl`, and configure the options that are required for Genesys Info Mart:
   - `role`—Specifies which data domain Genesys Info Mart will extract through the DAP:
     - For Configuration details, `role=ICON_CFG`
     - For Voice details, `role=ICON_CORE`
     - For Multimedia details, `role=ICON_MM`
     - For Outbound Contact details, `role=ICON_OCS`
   - `default-schema`—(Optional) Specifies the database schema or owner name, if it is different from the database user ID.
   - `geo-location`—(Optional) Specifies the location of the IDB.
   - `jdbc-*` options—Specify JDBC connection parameters. You must configure the `jdbc-url` option or other `jdbc-*` options if one of the following conditions applies:
     - DB Server is not running on the same host as the database management system (DBMS) server.
     - The port of the DBMS listener is not the default (`1521` for Oracle, `1433` for Microsoft SQL Server, or `5432` for PostgreSQL).

— For Oracle, the SID name is not the same as the TNS name.
— You want to enable SSL over the JDBC connection.

For information about how to create the `[gim-etl]` section and options, see Step 7 on page 229.

For more information about the `[gim-etl]` options, including requirements and restrictions for various roles, as well as information about the `jdbc-*` options, see Table 11.

End of procedure

Next Steps

- Verify that the `CLASSPATH` environment variable provides the correct path to the JDBC driver. For more information about modifying the `CLASSPATH` environment variable, see "Installing JDBC Drivers" on page 325.
- Create the Info Mart DAP. See "Info Mart DAP" on page 218.

Extraction DAP Options

Table 11 describes the required configuration options that you can configure in the `[gim-etl]` section on the `Options` tab of each extraction DAP `Application` object.

Note: Changes to extraction DAP configuration options take effect at the next run of the extraction job for that particular data domain; no restart of Genesys Info Mart Server is required.

Table 11: Extraction DAP Configuration Options

Option	Description
default-schema	Specifies the database schema or owner name, if it is different from the database user ID. If you do not specify this option, the database tables and other database objects are assumed to be owned by the user name that you specify on the `DB Info` tab. **Note for PostgreSQL:** Genesys recommends that you use lower case for all schema names. If you cannot do so, specify the value in case-sensitive mode by surrounding the value with a set of double quotes. Default Value: No default value Valid Values: Any valid database owner or schema name Changes Take Effect: At the next run of the extraction job for the particular data domain

Table 11: Extraction DAP Configuration Options (Continued)

Option	Description
geo-location	Specifies the location of the IDB. The option value is a string, which Genesys Info Mart compares against the value of the `geo-location` option in the Info Mart DAP. If the values are the same, the IDB is considered to be local; if the values are not the same, the IDB is considered to be remote. For an example, see the description of the `geo-location` option on page 221. In an HA environment, Genesys Info Mart uses geolocation as a tie-breaker to determine the best IDB from which to extract data: If data-quality criteria do not identify one IDB as the best, Genesys Info Mart gives preference to the local IDB. For more information about how Genesys Info Mart selects the IDB from which to extract, see "Criteria for Best IDB" on page 407. Default Value: "" Valid Values: Any string
jdbc-host	Specifies the host on which the DBMS is running. Configure this option if the DBMS server for the IDB is not on the same host as the DB Server. **Note:** If the `jdbc-url` option is defined, the host name that is specified in the URL overwrites the value of `jdbc-host`. Default Value: No default value Valid Values: Any valid host name Changes Take Effect: At the next run of the extraction job for the particular data domain
jdbc-port	Specifies the port of the DBMS listener. Configure this option if the DBMS uses a nondefault listener port. **Note:** If the `jdbc-url` option is defined, the port that is specified in the URL overwrites the value of `jdbc-port`. Default Value: `1521` (for Oracle), `1433` (for Microsoft SQL Server), or `5432` (for PostgreSQL) Valid Values: Any valid port number Changes Take Effect: At the next run of the extraction job for the particular data domain

Table 11: Extraction DAP Configuration Options (Continued)

Option	Description
jdbc-SID	(For Oracle only) Specifies the SID name, which Genesys Info Mart uses as a connection parameter to connect to Oracle databases. DB Server uses the Transparent Network Substrate (TNS) name, which is specified in the `DBMS Name` field on the `DB Info` tab of the DAP `Application` object. Configure this option if the SID name is not the same as the TNS name. **Note:** If the `jdbc-url` option is defined, the SID name that is specified in the URL overwrites the value of `jdbc-SID`. Default Value: No default value Valid Values: Any valid SID name Changes Take Effect: At the next run of the extraction job for the particular data domain

jdbc-url	Specifies all of the connection parameters for JDBC connection to the IDB. For the required syntax for this option, consult the vendor documentation for your JDBC driver. For example:
	• Oracle requires the following format to specify connection parameters for the Oracle thin client:
	`jdbc:oracle:thin:@(DESCRIPTION = (ADDRESS = (PROTOCOL = TCP)(HOST = <database_host_name>)(PORT = 1521))(CONNECT_DATA = (SERVER = DEDICATED)(SID = <Oracle_SID>)))`
	• Microsoft SQL Server requires the following format to specify connection parameters for the Microsoft JDBC driver:
	`jdbc:sqlserver://<dbserver>;DatabaseName=<dbname>;SelectMethod=cursor`
	• PostgreSQL requires the following format to specify connection parameters for the PostgreSQL JDBC driver:
	`jdbc:postgresql://<database_host_name>:<port>/<dbname>`
	Notes:
	• If this option is defined, it overwrites the values of any other `jdbc-*` options that might be configured for this DAP.
	• To enable SSL over the JDBC connection, you must specify additional parameters, as described in your RDBMS vendor documentation. For example, the following are sample URLs for Oracle and Microsoft SQL Server, respectively:
	◆ `jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=tcps)(HOST=servername)(PORT=2484))(CONNECT_DATA=(SERVICE_NAME=servicename)))`
	◆ `jdbc:sqlserver://host:1433;databaseName=NAME;integratedSecurity=true;encrypt=true;trustServerCertificate=true`
	/continued...

Table 11: Extraction DAP Configuration Options (Continued)

Option	Description
jdbc-url (continued)	• To enable IPv6 over the JDBC connection, you must specify additional parameters, as described in your RDBMS vendor documentation. For example, the following are sample URLs for Oracle, Microsoft SQL Server, and PostgreSQL, respectively: ◆ `jdbc:oracle:thin:@(DESCRIPTION = (ADDRESS = (PROTOCOL = TCP)(HOST = [<database_ipv6_address>])(PORT = 1521))(CONNECT_DATA = (SERVER = DEDICATED)(SERVICE_NAME = <oracle_SID>)))` ◆ `jdbc:sqlserver://;serverName=<database_ipv6_address>;DatabaseName=<dbname>;SelectMethod=cursor` ◆ `jdbc:postgresql://[<database_ipv6_address>]:<port>/<dbname>` Default Value: No default value Valid Values: Any valid JDBC URL Changes Take Effect: At the next run of the extraction job for the particular data domain
role	Specifies what type of data is stored in the source database that Genesys Info Mart accesses through this DAP. Default Value: No default value Valid Values: Any of the following: `ICON_CFG` Specifies the connection information for an IDB from which Genesys Info Mart will extract ICON Configuration details. This role must be associated with one and only one DAP, or with redundant DAPs that constitute an HA set. `ICON_CORE` Specifies the connection information for an IDB from which Genesys Info Mart will extract ICON Voice details, including agent activity details. This role is optional, but you need at least one DAP with a `role` option of `ICON_CORE` or `ICON_MM`. `ICON_MM` Specifies the connection information for an IDB from which Genesys Info Mart will extract ICON Multimedia details, including agent activity details. This role is optional, but you need at least one DAP with a `role` option of `ICON_CORE` or `ICON_MM`. `ICON_OCS` Specifies the connection information for an IDB from which Genesys Info Mart will extract ICON OCS details. **Note:** You cannot reconfigure the DAP role from `ICON_CORE` to `ICON_MM`, or vice versa, once data extraction occurs. Changes Take Effect: The next time that extraction jobs are launched

Info Mart DAP

To process and store data, Genesys Info Mart Server requires access to its target database. You can create and configure a dedicated JDBC DAP to enable the Genesys Info Mart Server to access the Info Mart database, or you can create and configure a non-JDBC DAP that you can also use for the Genesys Info Mart Administration Console to access the same database.

Use the following procedures, as applicable, to create and configure the Info Mart DAP that provides access to the target database:

- Procedure: Configuring a dedicated Info Mart DAP, on page 218
- Procedure: Configuring a shared Info Mart and Administration Console DAP, on page 219

Procedure: Configuring a dedicated Info Mart DAP

Purpose: To configure a new, dedicated DAP that enables Genesys Info Mart Server to access the Info Mart database.

Prerequisites

- Procedure: Preparing the Info Mart database, on page 191

Start of procedure

1. Create a JDBC DAP, as described in Procedure: Configuring JDBC DAPs, on page 225.
2. Configure the required options in the `[gim-etl]` section on the `Options` tab:
   - `role`—The required value is `INFO_MART`.
   - `default-schema`—(Optional) Specifies the database schema or owner name, if it is different from the database user ID.
   - `geo-location`—(Optional) Specifies the location of the Info Mart database.

   For information about how to create the `[gim-etl]` section and options, see Step 8 on page 227.

   For more information about the options, see Table 12 on page 220.

End of procedure

Next Steps

- Create the Administration Console DAP. See Procedure: Configuring a dedicated Administration Console DAP, on page 224.

Procedure: Configuring a shared Info Mart and Administration Console DAP

Purpose: To configure a non-JDBC DAP that enables both Genesys Info Mart and the Genesys Info Mart Administration Console to access the Info Mart database.

Prerequisites

- Procedure: Preparing the Info Mart database, on page 191.
- Install DB Server to handle database requests that the Genesys Info Mart Administration Console submits through this DAP.

Start of procedure

1. Create a non-JDBC DAP for the Info Mart database, as described in Procedure: Configuring non-JDBC DAPs, on page 228.
2. On the `Options` tab, create a new configuration section, named `gim-etl`, and configure the required options:
   - `role`—Specifies the connection information for Genesys Info Mart and the Genesys Info Mart Administration Console, respectively:
     - For Genesys Info Mart, the required value is `INFO_MART`.
     - For the Genesys Info Mart Administration Console, the required value is `ADMIN_CONSOLE`.

     For the shared DAP, the required value is `ADMIN_CONSOLE, INFO_MART`.
   - `default-schema`—(Optional) Specifies the database schema or owner name, if it is different from the database user ID.
   - `geo-location`—(Optional) Specifies the location of the Info Mart database.
   - `jdbc-*` options—Specify JDBC connection parameters for Genesys Info Mart. You must configure the `jdbc-url` option or other `jdbc-*` options, if one of the following conditions applies:
     - DB Server is not running on the same host as the DBMS server for the Info Mart database.
     - The port of the DBMS listener is not the default (`1521` for Oracle, `1433` for Microsoft SQL Server, or `5432` for PostgreSQL).

— For Oracle, the SID name is not the same as the TNS name.
— You want to enable SSL over the JDBC connection.

For information about how to create the `[gim-etl]` section and options, see Step 7 on page 229.

For more information about all of the configuration options for the shared Info Mart DAP, see Table 12 on page 220.

End of procedure

Next Steps

- Verify that the `CLASSPATH` environment variable provides the correct path to the JDBC driver. For more information about modifying the `CLASSPATH` environment variable, see "Installing JDBC Drivers" on page 325.
- Configure the Genesys Info Mart `Application` object. For more information, see Chapter 12 on page 231.

Info Mart DAP Options

Table 12 describes the required options that you must configure in the `[gim-etl]` section on the `Options` tab of the Info Mart DAP `Application` object.

Note: Changes to configuration options for the Info Mart DAP take effect on restart of Genesys Info Mart Server.

Table 12: Info Mart DAP Configuration Options

Option	Description
default-schema	Specifies the database schema or owner name, if it is different from the database user ID. If you do not specify this option, the database tables and other database objects are assumed to be owned by the user name that you specify on the `DB Info` tab. **Note for PostgreSQL:** Genesys recommends that you use lower case for all schema names. If you cannot do so, specify the value in case-sensitive mode by surrounding the value with a set of double quotes. Default Value: No default value Valid Values: Any valid database schema or owner name

Table 12: Info Mart DAP Configuration Options (Continued)

Option	Description
role	Specifies the connection information for the Info Mart database. This role is required; it can be associated with only one DAP. Default Value: No default value Valid Values: • `INFO_MART`—Specifies connection information for the Genesys Info Mart Server to access the Info Mart database to read and write data. Use this value for a dedicated JDBC DAP. • `ADMIN_CONSOLE, INFO_MART`—Specifies connection information for the Genesys Info Mart Administration Console and the Genesys Info Mart Server, respectively, to access the Info Mart database. Use this value for a shared non-JDBC DAP. For more information about the `ADMIN_CONSOLE` role, see Table 13 on page 225.
geo-location	Specifies the location of the Info Mart database. The option value is a string, which Genesys Info Mart compares against the value of the `geo-location` option in redundant extraction DAPs. If the values are the same, the IDB for which the extraction DAP provides the connection information is considered to be local; if the values are not the same, the IDB is considered to be remote. For example, if the value of the `geo-location` option on the Info Mart DAP is `Site 1`, the value of the option on the extraction DAP for IDB-1 is `Site 1`, and the value on the extraction DAP for IDB-2 is `Site 2`, Genesys Info Mart will consider IDB-1 to be local and IDB-2 to be remote, regardless of where the hosts are physically located. In an HA environment, Genesys Info Mart uses geolocation as a tie-breaker to determine the best IDB from which to extract data: If data-quality criteria do not identify one IDB as the best, Genesys Info Mart gives preference to the local IDB. For more information about how Genesys Info Mart selects the IDB from which to extract, see "Criteria for Best IDB" on page 407. Default Value: "" Valid Values: Any string
jdbc-host	Specifies the host on which the DBMS is running. Configure this option if the DBMS server for the Info Mart database is not on the same host as the DB Server. **Note:** If the `jdbc-url` option is defined, the host name that is specified in the URL overwrites the value of `jdbc-host`. Default Value: No default value Valid Values: Any valid host name

Table 12: Info Mart DAP Configuration Options (Continued)

Option	Description
jdbc-port	Specifies the port of the DBMS listener. Configure this option if the DBMS uses a non-default listener port. **Note:** If the `jdbc-url` option is defined, the port that is specified in the URL overwrites the value of `jdbc-port`. Default Value: `1521` (for Oracle),`1433` (for Microsoft SQL Server), or `5432` (for PostgreSQL) Valid Values: Any valid port number
jdbc-SID	(For Oracle only) Specifies the SID name, which Genesys Info Mart uses as a connection parameter to connect to Oracle databases. DB Server uses the Transparent Network Substrate (TNS) name, which is specified in the `DBMS Name` field on the `DB Info` tab of the DAP `Application` object. Configure this option if the SID name is not the same as the TNS name. **Note:** If the `jdbc-url` option is defined, the SID name that is specified in the URL overwrites the value of `jdbc-SID`. Default Value: No default value Valid Values: Any valid SID name

jdbc-url	Specifies all of the connection parameters for JDBC connection to the Info Mart database. For the required syntax for this option, consult the vendor documentation for your JDBC driver. For example:
	• Oracle requires the following format to specify connection parameters for the Oracle thin client:
	`jdbc:oracle:thin:@(DESCRIPTION = (ADDRESS = (PROTOCOL = TCP)(HOST = <database_host_name>)(PORT = 1521))(CONNECT_DATA = (SERVER = DEDICATED)(SID = <Oracle_SID>)))`
	• Microsoft SQL Server requires the following format to specify connection parameters for the Microsoft JDBC driver:
	`jdbc:sqlserver://<dbserver>;DatabaseName=<dbname>;SelectMethod=cursor`
	• PostgreSQL requires the following format to specify connection parameters for the PostgreSQL JDBC driver:
	`jdbc:postgresql://<database_host_name>:<port>/<dbname>`
	Notes:
	• If this option is defined, it overwrites the values of any other `jdbc-*` options that might be configured for this DAP.
	/continued...

Table 12: Info Mart DAP Configuration Options (Continued)

Option	Description
jdbc-url (continued)	• To enable SSL over the JDBC connection, you must specify additional parameters, as described in your RDBMS vendor documentation. For example, the following are sample URLs for Oracle and Microsoft SQL Server, respectively: ◆ `jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS=(PROTOCOL=tcps)(HOST=servername)(PORT=2484))(CONNECT_DATA=(SERVICE_NAME=servicename)))` ◆ `jdbc:sqlserver://host:1433;databaseName=NAME;integratedSecurity=true;encrypt=true;trustServerCertificate=true` • To enable IPv6 over the JDBC connection, you must specify additional parameters, as described in your RDBMS vendor documentation. For example, the following are sample URLs for Oracle, Microsoft SQL Server, and PostgreSQL, respectively: ◆ `jdbc:oracle:thin:@(DESCRIPTION = (ADDRESS = (PROTOCOL = TCP)(HOST = [<database_ipv6_address>])(PORT = 1521))(CONNECT_DATA = (SERVER = DEDICATED)(SERVICE_NAME = <oracle_SID>)))` ◆ `jdbc:sqlserver://;serverName=<database_ipv6_address>;DatabaseName=<dbname>;SelectMethod=cursor` ◆ `jdbc:postgresql://[<database_ipv6_address>]:<port>/<dbname>` Default Value: No default value Valid Values: Any valid host name

Administration Console DAP

To monitor and manage ETL jobs, the Genesys Info Mart Administration Console requires access to the Info Mart database. The Genesys Info Mart Administration Console connects to the Info Mart database through a non-JDBC DAP that is associated with a DB Server, and it does not use any of the JDBC options.

You can create and configure a dedicated non-JDBC DAP to enable the Genesys Info Mart Administration Console to access the Info Mart database, or you can configure a non-JDBC DAP that enables both Genesys Info Mart and the Genesys Info Mart Administration Console to access the Info Mart database.

- To configure a shared DAP for Genesys Info Mart and the Genesys Info Mart Administration Console, see Procedure: Configuring a shared Info Mart and Administration Console DAP, on page 219.

To configure a dedicated Administration Console DAP, use the following procedure.

Note: Changes to configuration options for the Administration Console DAP take effect on restart of the Genesys Info Mart Administration Console.

Procedure: Configuring a dedicated Administration Console DAP

Purpose: To configure a dedicated, non-JDBC DAP that enables the Genesys Info Mart Administration Console to access the Info Mart database.

Prerequisites

- Procedure: Preparing the Info Mart database, on page 191.
- Install DB Server to handle database requests that are submitted through this DAP.

Start of procedure

1. Create a non-JDBC DAP for the Info Mart database, as described in Procedure: Configuring non-JDBC DAPs, on page 228.
2. On the `Options` tab, create a new configuration section, named `gim-etl`, and configure the `role` option. Set the value of this option to `ADMIN_CONSOLE`.

   For more information about the `role` option for the Administration Console DAP, see Table 13 on page 225.

   For information about how to create the `[gim-etl]` section and options, see Step 7 on page 229.

End of procedure

Next Steps

- Verify that the `CLASSPATH` environment variable provides the correct path to the JDBC driver. For more information about modifying the `CLASSPATH` environment variable, see "Installing JDBC Drivers" on page 325.
- Configure the Genesys Info Mart `Application` object. For more information, see Chapter 12 on page 231.

Table 13: Configuration Options for the Administration Console DAP

Option	Description
role	Specifies the connection information for the Info Mart database. Default Value: No default value Valid Values: `ADMIN_CONSOLE` Specifies the read-only connection information to the Info Mart database from which the Genesys Info Mart Administration Console will retrieve ETL job status, job execution history, and job schedule data. Changes Take Effect: On restart of the Genesys Info Mart Administration Console **Note:** This DAP must have a DB Server associated with it.

DAP Configuration Supporting Procedures

This section describes how to configure the two standard types of DAPs. JDBC DAPs are configured by using JDBC DAP options, while non-JDBC DAPs do not use Java Database Connectivity.

Configuring JDBC DAPs

Procedure: Configuring JDBC DAPs

Purpose: To configure a JDBC DAP to use for access to Genesys Info Mart source and target databases.

This is a supporting procedure for any DAP configuration that involves a JDBC DAP.

Prerequisites

- Before you can configure the DAPs that are required for your topology, you must import the 8.x DAP template into your environment. For information about how to create new `Application` objects, refer to the *Framework 8.x Deployment Guide*.

Start of procedure

1. Right-click the Applications folder, and select New > Application. The Browse dialog box appears.
2. From the list, select the Database_Access_Point_800 template, and then click OK. The New Application Properties dialog box appears.

General Tab

3. On the General tab:
   a. In the Name text box, enter the name of this DAP (for example, GIM80_InfoMart_DAP).
   b. Select the JDBC Connection check box.
   c. Select the State Enabled check box, if it is not already selected.

Tenants Tab

4. No configuration of tenants is required. Proceed to the next tab.

> **Note:** If you do not have a multi-tenant environment, the Tenants tab does not appear.

DB Info Tab

5. On the DB Info tab:
   a. Verify that the DBMS Name text box is disabled. If it is not, return to the General tab, and select the JDBC Connection check box.
   b. From the DBMS type drop-down list, select the DBMS type. This type will correspond to the DBMS type of the database to which this DAP will connect.
   c. In the Database Name text box, enter the exact name of the database to which this DAP will connect. For example, on Oracle, this is the Transparent Network Substrate (TNS) name.
   d. In the User Name and Password text boxes, enter the user name and password of the user who has read and/or write access to the database tables.

JDBC Info Tab

6. On the JDBC Info tab:
   - From the Role drop-down list, select Role Main.

> **Note:** Genesys Info Mart ignores all other fields on the JDBC Info tab.

Server Info Tab

7. On the Server Info tab:
   a. From the Host drop-down list, select the host name of the database server. If the host name does not appear in the list, you must add it.

b. Specify the communication port for the database server.

i. With the `default` port active, click the `Edit Port` button. The `Properties` dialog box appears.

ii. In the `Communication Port` text box, enter the port number for the database server, and then click `OK`.

For Microsoft SQL Server, the default port is `1433`; for Oracle, it is `1521`; for PostgreSQL, it is `5432`.

Options Tab

8. On the `Options` tab:

a. Click the `Create New Section/Option` icon. The `Add Section` dialog box appears.

b. In the `Section Name` text box, enter `gim-etl`, and then click `OK`.

c. Double-click the `gim-etl` section name.

d. Specify the default schema, if required:

Note: If the schema or owner is different from the user ID, you must specify the `default-schema` option. If you do not specify this option, the schema that will be used is the user ID that you specified on the `DB Info` tab of the DAP.

i. Click the `Create New Section/Option` icon. The `Edit Option` dialog box appears.

ii. In the `Option Name` text box, enter `default-schema`.

iii. In the `Option Value` text box, enter the name of the schema owner, and then click `OK`.

For PostgreSQL, specify the name in lower case or in case-sensitive mode (in other words, using double quotes).

This option specifies an owner that is different from the one that is associated with the specified user ID.

e. Specify the DAP role:

i. Click the `Create New Section/Option` icon again to add another option. The `Edit Option` dialog box appears.

ii. In the `Option Name` text box, enter `role`.

iii. In the `Option Value` text box, enter the `role` that corresponds to the particular database. For information about which role you should specify, see Table 11 on page 214.

Note: You can configure one or more roles for a single DAP by specifying all of the roles, separated by commas—for example: `ICON_CFG, ICON_CORE, ICON_OCS`.

f. (Optional) Specify SSL connection parameters:
   i. Create another new option and name it `jdbc-url`.
   ii. In the `Option Value` text box, enter the parameters that your RDBMS requires the JDBC user to specify in order to use SSL. For information about the specific parameters that your RDBMS requires, see your RDBMS vendor documentation.

9. Click `Apply` to save your changes, and then click `OK` to close the `New Application Properties` dialog box.
10. Repeat Step 1 through 9 for each JDBC DAP that you require.

End of procedure

Next Steps

- Return to the main procedure for the DAP that you have been creating.

Configuring Non-JDBC DAPs

Procedure: Configuring non-JDBC DAPs

Purpose: To configure a non-JDBC DAP so that Genesys Info Mart clients can access target databases.

This is a supporting procedure for any DAP configuration that involves a non-JDBC DAP.

Prerequisites

- Before you can configure the DAPs that are required for your topology, you must first import the 8.x DAP template into your environment. For information on creating new `Application` objects, refer to the *Framework 8.x Deployment Guide.*
- Before you can configure a non-JDBC DAP, you must install the DB Server that handles database requests that are submitted through this DAP. For information about how to configure and install DB Server, refer to the *Framework 8.x Deployment Guide.*

Start of procedure

1. Right-click the `Applications` folder, and select `New > Application`. The `Browse` dialog box appears.
2. From the list, select the `Database_Access_Point_800` template, and then click `OK`. The `New Application Properties` dialog box appears.

General Tab

3. On the General tab:
   a. In the Name text box, enter the name of this DAP (for example, GIM81_AdminConsole_DAP).
   b. Click the Browse button that is next to the DB Server text box, and then select the DB Server that has been configured to connect to the Info Mart database.
   c. Do *not* select the JDBC Connection check box.
   d. Select the State Enabled check box, if it is not already selected.

Tenants Tab

4. No configuration of tenants is required. Proceed to the next tab.

Note: If you do not have a multi-tenant environment, the Tenants tab does not appear.

DB Info Tab

5. On the DB Info tab:
   a. In the DBMS Name text box, enter the name that is particular to the DBMS type:
      - Oracle— Enter the TNS name that the database client uses to connect to the database.
      - Microsoft SQL Server—Enter the name of the SQL Server in which the database resides.
      - PostgreSQL—Enter the name of the host on which PostgreSQL is running.
   b. From the DBMS type drop-down list, select the DBMS type. This type will correspond to the DBMS type of the database to which this DAP will connect.
   c. In the Database Name text box, enter the exact name of the database to which this DAP will connect.
   d. In the User Name and Password text boxes, enter the user name and password of the user who has read and/or write access to the database tables.

JDBC Info Tab

6. On the JDBC Info tab, verify that all of the fields are disabled. If they are not disabled, return to the General tab, and clear the JDBC Connection check box.

Options Tab

7. On the Options tab:
   a. Click the Create New Section/Option icon. The Add Section dialog box appears.
   b. In the Section Name text box, enter gim-etl, and then click OK.
   c. Double-click the gim-etl section name.
   d. Specify the DAP role:
      i. Click the Create New Section/Option icon. The Edit Option dialog box appears.
      ii. In the Option Name text box, enter role.

iii. In the Option Value text box, enter the required role.

For the Genesys Info Mart Administration Console DAP, enter ADMIN_CONSOLE. See Table 13 on page 225.

e. Specify the default schema, if required:

Note: If the schema or owner is different from the user ID, you must specify the default-schema option. If you do not specify this option, the schema that will be used is the user ID that you specified on the DB Info tab of the DAP.

i. Click the Create New Section/Option icon. The Edit Option dialog box appears.

ii. In the Option Name text box, enter default-schema.

iii. In the Option Value text box, enter the name of the schema owner, and then click OK.

For PostgreSQL, specify the name in lower case or in case-sensitive mode (in other words, using double quotes).

This option specifies an owner that is different from the one that is associated with the specified user ID.

f. Specify JDBC connection parameters:

i. Create another new option and name it jdbc-url, or create additional, separate jdbc-* options, as described in Table 11 on page 214 or Table 12 on page 220.

ii. In the Option Value text box, enter the parameters that your RDBMS requires.

8. Click Apply to save your changes, and then click OK to close the New Application Properties dialog box.

End of procedure

Next Steps

- Return to the main procedure for the DAP that you have been creating.

Chapter

12 Configuring the Genesys Info Mart Application

This chapter describes how to configure Genesys Info Mart, which you must do before you install it. It contains the following sections:

Overview

Before You Proceed

Before you can configure Genesys Info Mart, make sure that you have:

- Prepared your data-source applications and their databases.
- Prepared your Genesys Info Mart target databases.
- Configured the database access points (DAPs) that are required to access the source and target databases in your topology.

If you have not already done so, review the preceding chapters before you continue.

Task Flow for Configuring Genesys Info Mart

The following table summarizes the task flow to configure an `Application` object and related configuration objects for Genesys Info Mart.

Task Summary: Configuring Genesys Info Mart

Objective	Related Procedures and Actions
Create a server application to support detailed reporting about interactions of various media types.	Create an `Application` object for Genesys Info Mart Server: 1. Import the `Genesys_Info_Mart_ETL_800.apd` template. For more information, see Procedure: Importing the application template. 2. Create the `Application` object with basic configuration. For more information, see Procedure: Creating the Genesys Info Mart application, on page 235.
Customize the configuration of the Genesys Info Mart `Application` object to support the functionality that you require.	1. Review the configuration options and decide on the settings that are appropriate for your environment. See "Options Summary Charts" on page 240 and Chapter 14 on page 271. 2. Modify settings of configuration options, as required for your environment. For more information, see Procedure: Setting the Genesys Info Mart application options, on page 239. (Optional) When you configure the Genesys Info Mart application, create custom `[date-time-*]` configuration sections to support multiple calendar dimensions. For more information, see Procedure: Setting the Genesys Info Mart application options, Step 6 on page 240.
Configure the other objects that are required to support detailed reporting about interactions of various media types.	Configure necessary options in individual configuration objects, for use by Genesys Info Mart Server. For more information, see: 1. Procedure: Setting the Annex tab to display, on page 254 2. (For Voice and Multimedia details) Procedure: Configuring the switch for ICON and Genesys Info Mart reporting, on page 255 3. (For Multimedia details) Procedure: Setting Media Type Business Attribute object options for Genesys Info Mart reporting, on page 260 4. (For Voice and Multimedia details) Procedure: Configuring a DN for ICON and Genesys Info Mart reporting, on page 261 5. (For Multimedia details) Procedure: Setting Script object options for Genesys Info Mart reporting, on page 265 6. (For Outbound Contact details) "Configuring Field Objects" on page 266

If you need to enable high availability (HA) in either a new or an existing Genesys Info Mart environment, see "Enabling High Availability" on page 152.

Importing the Application Template

Before you can configure an Application object for Genesys Info Mart, you must import its application template. The application template provides a majority of the configuration options, as well as the default values for them.

The Genesys Info Mart installation package includes an application template, which is named Genesys_Info_Mart_ETL_<release>.apd.

Note: If you plan on deploying Genesys Interactive Insights (GI2) reports or the Reporting and Analytics Aggregates (RAA) package, review the information about the template for the Genesys Info Mart application and recommendations on application deployment in the *Reporting and Analytics Aggregates 8.1 Deployment Guide*.

Procedure: Importing the application template

Purpose: To import an application template that you can use to create as many Application objects of the same type as you need.

Note: You must import the application template only once, no matter how many Genesys Info Mart Application objects you create.

Prerequisites

- Review "Before You Proceed" on page 231.

Start of procedure

1. Insert the Genesys Info Mart CD into your CD-ROM drive.
2. Start Genesys Configuration Manager.
3. Navigate to Configuration > Environment.
4. Right-click the Application Templates folder, and select Import application template.
5. Navigate to the Templates directory on the Genesys Info Mart CD or in your FTP package, and select the following template to import:

   Genesys_Info_Mart_ETL_<release>.apd
6. Click Open.

   The New Application Template Properties dialog box appears (Figure 14 on page 234).

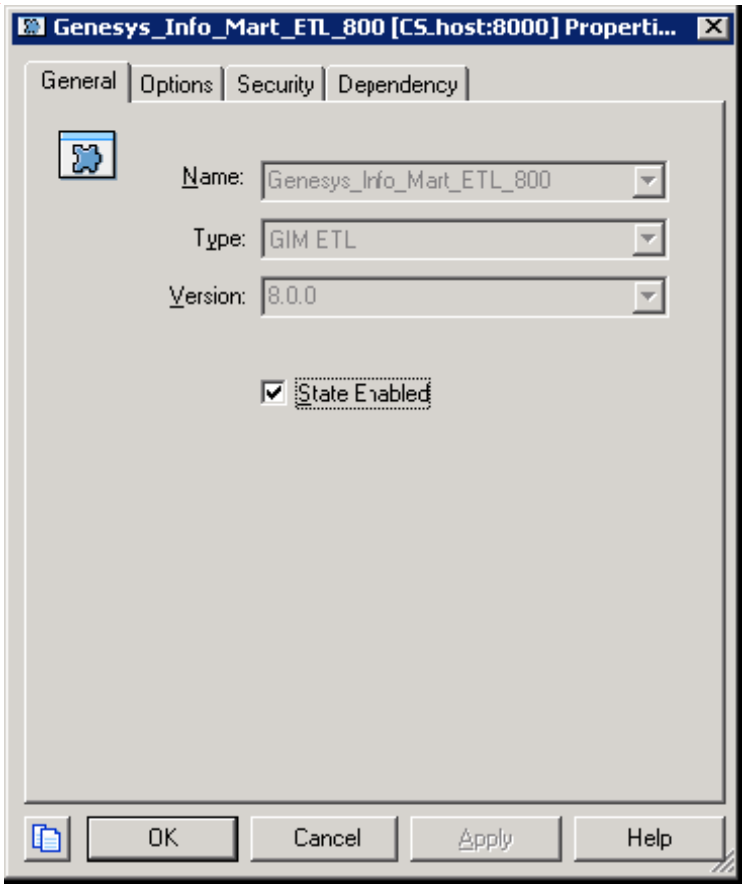

Figure 14: New Application Template Properties Dialog Box

7. Either enter a new name for the Genesys Info Mart template in the `Name` text box, or accept the default name.
8. Click `OK`.

End of procedure

Next Steps

- If you plan to deploy GI2 reports or the RAA package, modify the Genesys Info Mart application template to support aggregation, as described in the *Reporting and Analytics Aggregates 8.1 Deployment Guide*.
- Otherwise, see “Configuring the Genesys Info Mart Application” on page 235.

Configuring the Genesys Info Mart Application

After you import the application template, you can create and configure an `Application` object for Genesys Info Mart by using Configuration Manager.

Procedure: Creating the Genesys Info Mart application

Purpose: To create an application with basic configuration for your environment, by using the Genesys Info Mart 8.1 application template that you just imported.

Prerequisites

- Procedure: Importing the Application Template, on page 233.
- In an environment in which you plan to deploy GI2 reports or the RAA package, you have completed any required modifications to the Genesys Info Mart application template, as described in the *Reporting and Analytics Aggregates 8.1 Deployment Guide*.

Start of procedure

1. In the Genesys Configuration Manager, right-click the `Applications` folder, and then select `New > Application`.

   The `Application Templates Browse` dialog box appears.

2. Select the Genesys Info Mart 8.1 template that you just created, and then click `OK`.

   The `New Application Properties` dialog box appears (Figure 15 on page 236).

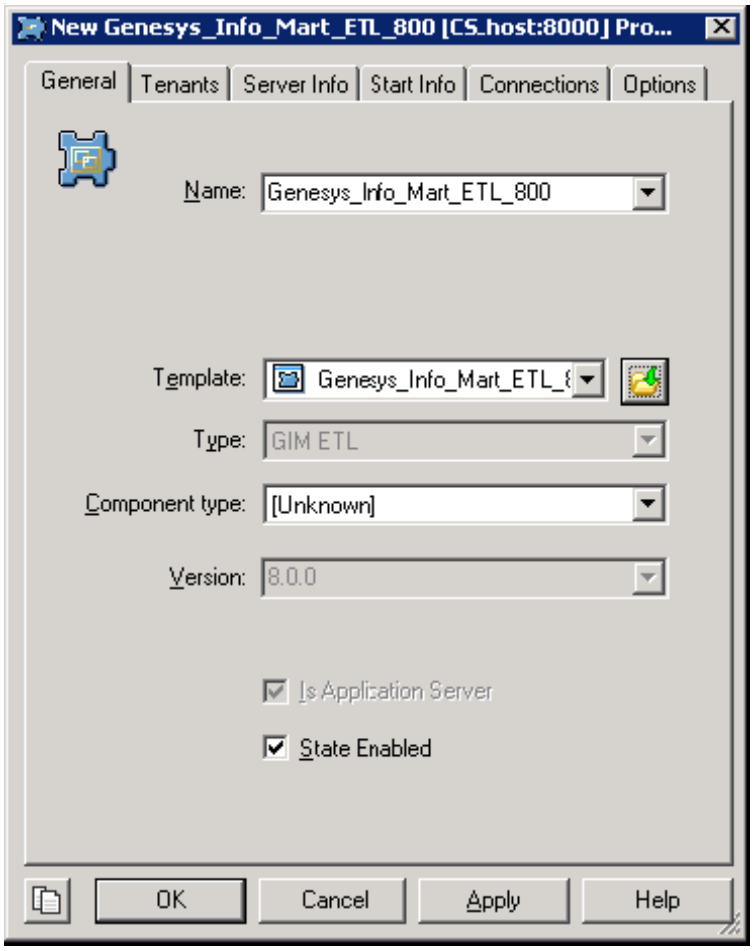

Figure 15: New Application Properties Dialog Box

General Tab

3. On the `General` tab:
   a. In the `Name` text box, enter a name for your Genesys Info Mart application, or select one from the drop-down list.
   b. Make sure that the `State Enabled` check box is selected.

Tenants Tab

4. No configuration of tenants is required. Proceed to the next tab.

> **Note:** If you do not have a multi-tenant environment, the `Tenants` tab does not appear.

Server Info Tab

5. On the `Server Info` tab:
   a. In the `Host` text box, select the host on which you will install the Genesys Info Mart Server, or click `Browse` to navigate to the host location. If the host does not appear in the list, you must add it.

b. In the Communications Port text box, enter the port number that corresponds to your host. The Genesys Info Mart Manager and Genesys Info Mart Administration Console use this port to communicate with the Genesys Info Mart Server.

Start Info Tab

6. On the Start Info tab:
   a. Enter *any* value in each of the following text boxes:
      i. Working Directory
      ii. Command Line
      iii. Command Line Arguments

> **Note:** The values that you enter are merely placeholders; they are updated with actual values during the Genesys Info Mart installation process.

   b. If your deployment will include primary and standby instances of the Genesys Info Mart application, ensure that the Auto-Restart check box is unchecked on both Application objects (in other words, automatic restart is not selected). If your deployment will include only a single instance of the Genesys Info Mart application, you may choose to select Auto-Restart.

Connections Tab

7. On the Connections tab, add a connection to:
   a. Each Interaction Concentrator (ICON) Application that you configured to capture data-source information for Genesys Info Mart.
   b. Each of the DAPs that you configured in "Configuring DAPs" on page 207.

> **Note:** If your deployment uses the Genesys Info Mart Administration Console, be sure to add a connection to the Administration Console database access point (DAP).

   c. The Message Server Application object—Genesys Info Mart uses Message Server to send log messages to the Genesys Central Logger.

      For information about how to secure the connection to Message Server, see "Enabling Secure Connection Features" on page 154.

d. (Optional) The Configuration Server application (named `confserv`)—You must configure an overt connection to Configuration Server only if you want to:

i. Use Advanced Disconnect Detection Protocol (ADDP) for the connection.

To enable ADDP for the connection, specify `addp` as the `Connection Protocol`, and set the values for the `Local Timeout`, `Remote Timeout`, and `Trace Mode` properties. For more information, see the *Framework 8.x Deployment Guide*.

ii. Secure the connection. For more information, see "Enabling Secure Connection Features" on page 154.

Note: In an environment with distributed data centers, configure this optional connection to Configuration Server Proxy instead of to Configuration Server.

8. Click `Apply` to save your changes, and then click `OK` to close the `New Applications Properties` dialog box.

End of procedure

Next Steps

- Verify that customized and required options are set correctly. See the information and instructions in "Configuring Options for Genesys Info Mart".

Configuring Options for Genesys Info Mart

This section describes how to customize your Genesys Info Mart configuration settings to best suit your environment. Refer to this section while you install your Genesys Info Mart 8.1 application, or refer to it after the initial configuration to perform additional customization.

This section includes the following procedure that helps you set configuration options in the `Genesys Info Mart Application` object:

- Procedure: Setting the Genesys Info Mart application options, on page 239

This section includes tables that summarize the Genesys Info Mart options by area of functionality:

- Table 14, "Genesys Info Mart Data-Related Options," on page 241
- Table 15, "Genesys Info Mart Operations-Related Options," on page 248

This section also includes a table that lists configuration options that are provided with Genesys Info Mart for GI2 reports and the RAA package:

- Table 16, "Genesys Info Mart Options for GI2 and RAA," on page 252

To customize data processing for certain environments and various types of ICON details, you might also have to set options on other configuration objects, such as DNs. For more information, see Chapter 13 on page 253.

Procedure: Setting the Genesys Info Mart application options

Purpose: To configure the settings for the Genesys Info Mart 8.1 application to better suit your environment.

You specify configuration options on the Options tab of the Genesys Info Mart Application object.

Options are specific to your application and release. They appear on the Options tab in sections that are devoted to specific functions. Refer to:

- "Options Summary Charts" on page 240 for information on general functional grouping.
- Chapter 14 on page 271 for a complete listing and explanation of all of the configuration options.

Prerequisites

- Procedure: Configuring the Genesys Info Mart Application, on page 235

Start of procedure

1. In the Genesys Configuration Manager, open the Application object that you configured for your Genesys Info Mart.

   The Application Properties dialog box appears.
2. Select the Options tab.
3. If you plan to use GI2 or RAA, review the options that are listed in Table 16 on page 252, and change the default settings as required for your environment.

   If you do not plan to use GI2 or RAA, ensure that these aggregation-related options retain their default values.

   Note: The options that are listed in Table 16 are necessary to enable aggregation in your Genesys Info Mart deployment, but not sufficient for GI2 reports. For more information about how to configure Genesys Info Mart to support aggregation, see the *Reporting and Analytics Aggregates 8.1 Deployment Guide.*

4. Configure the options that control the extract, transform, and load (ETL) process.

   Refer to Table 15 on page 248 for the list of relevant options and their default values.

5. Configure the options that affect the content and quality of data that is stored in the Info Mart database.

   Refer to Table 14 on page 241 for the list of relevant options and their default values.

6. (Optional) Customize the calendar dimensions:
   - To modify the default calendar to support the ISO 8601 standard for week numbering, modify the values of the `simple-week-numbering`, `first-day-of-week`, and `min-days-in-first-week` options in the `[date-time]` section (see page 271).
   - To change the default calendar (which is a Gregorian one) to a fiscal calendar, add the `fiscal-year-week-pattern` and `fiscal-year-start` options (see page 274) to the `[date-time]` section on the `Options` tab, and set the values as required.
   - To provide multiple, customized calendars, create additional `[date-time-*]` sections with the same options as in the `[date-time]` section, and configure the options as required for your reporting purposes.

7. After you set all of the desired options for your environment, click `Apply` to save your changes, and then click `OK` to close the `Applications Properties` dialog box.

End of procedure

Next Steps

- If you are creating additional calendars for reporting purposes, modify the `make_gim.sql` or `make_gim_partitioned.sql` script to create the custom calendar tables in the Info Mart database. For more information, see the prerequisites for the database-creation procedure on page 191.
- Otherwise, proceed to configuring ICON- and Genesys Info Mart–related options on other objects. See Chapter 13 on page 253.

Options Summary Charts

This section groups the Genesys Info Mart configuration options, as follows:

- Data-Processing Configuration Options, page 241
- Operations-Related Configuration Options, page 248
- Aggregation-Related Configuration Options, page 252

Data-Processing Configuration Options

Table 14 on page 241 lists the configuration options that affect the content and quality of data that is stored in the Genesys Info Mart database. The table groups the options by functional area and, within each functional area, by type of data. For each functional area, the options that apply in all environments are listed first.

Review the options to identify the ones that are relevant to the reports that you are required to provide on the data types that are relevant to your environment.

For example, suppose that you would like to report on interactions and agent activity at the DN and ACD queue levels, using Genesys Info Mart pre-aggregated data, in a contact center that processes inbound voice interactions. In this case, select the necessary options from the following sections in Table 14: "Voice media interactions data" on page 241, "Agent activity data not specific to media type" on page 243, "Voice agent activity data" on page 243, and "Queue activity data for voice (ACD queue or Virtual Queue)" on page 244.

Refer to Chapter 14 on page 271, as required, to understand how a particular configuration option works and what functionality it enables. Finally, review the valid values, which are documented as part of the option descriptions, and determine the values that are appropriate to your environment.

Carefully consider the values that you set for data-related options. Changing these settings after Genesys Info Mart has started collecting data compromises data consistency, especially for long-lived multimedia interactions. For example, if an agent works on an e-mail reply over a period of time and periodically saves drafts to an Interaction Workbin, there might be mixed results for the workbin activity that is associated with the e-mail interaction if the value of the `populate-workbin-as-hold` option (see page 307) changes: Some of the workbin time might be considered to be mediation (which Genesys Info Mart might or might not represent in the dimensional model, depending on other configuration), and some might be reported as hold.

Table 14: Genesys Info Mart Data-Related Options

Configuration Object	Section Name	Option Name and Default Value	Comments
Voice media interactions data			
Genesys Info Mart Application	`[gim-etl]`	`short-abandoned-threshold=10`	In the `Application` object, configure the option on the `Options` tab.
	`[gim-transformation]`	`default-ivr-to-self-service=false`	
(Network) Switch	`[gim-etl]`	`network-switch=FALSE`	In the `Switch` object, configure the option on the `Annex` tab.

Table 14: Genesys Info Mart Data-Related Options (Continued)

Configuration Object	Section Name	Option Name and Default Value	Comments
Multimedia interactions data			
Genesys Info Mart Application	[gim-etl]	max-parties-per-call=100 (option discontinued in release 8.1.4) max-thread-duration-after-inactive-in-days=31 (30 in release 8.1.1)	In the Application object, configure the option on the Options tab.
	[gim-etl-media-<media type>] where <media type> exactly matches the name of the applicable Media Type Business Attribute (for example, email).	short-abandoned-threshold=10	Each [gim-etl-media-<media type>] section contains options that apply for the specific media type only. This setting can be overridden at the tenant level. In the Application object, configure the option on the Options tab.
Tenant Media Type Business Attribute	[gim-etl-media]	short-abandoned-threshold=10	For the particular Media Type attribute, which is configured for a particular tenant, this setting overrides the same option that is configured at the application level. In the Media Type Business Attribute object for the tenant, configure the option on the Annex tab.

Table 14: Genesys Info Mart Data-Related Options (Continued)

Configuration Object	Section Name	Option Name and Default Value	Comments
Agent activity data not specific to media type			
Genesys Info Mart Application	[gim-etl-populate]	populate-sm-resource-session-facts=true populate-sm-resource-state-facts=true populate-sm-resource-state-reason-facts=true **Note:** These options were discontinued in release 8.1.1.	In the Application object, configure the options on the Options tab.
	[gim-etl]	max-session-duration-in-hours=24 max-state-duration=14400 sm-resource-state-priority=ACW, NOT_READY, BUSY, READY	
	[gim-transformation]	ignored-reason-codes=INTERACTION_WORKSPACE	
Switch	[gim-etl]	factor-dnd-into-sm-resource-states=FALSE (for voice-handling and SIP switches) factor-dnd-into-sm-resource-states=TRUE (for multimedia-handling switches)	In the Switch object, configure the option on the Annex tab.
Voice agent activity data			
Genesys Info Mart Application	[gim-etl-populate]	populate-sm-voice-resource-activity=true	In the Application object, configure the option on the Options tab.
Multimedia agent activity data			
Genesys Info Mart Application	[gim-etl-populate]	populate-sm-email-resource-activity=true populate-sm-chat-resource-activity=true	In the Application object, configure the options on the Options tab.

Table 14: Genesys Info Mart Data-Related Options (Continued)

Configuration Object	Section Name	Option Name and Default Value	Comments
Queue activity data for voice (ACD queue or Virtual Queue)			
Genesys Info Mart Application	[gim-etl]	q-answer-threshold-voice=60 q-short-abandoned-threshold-voice=10	The q-answer-threshold-voice setting can be overridden on the Switch or DN object level. In the Application object, configure the options on the Options tab.
	[gim-transformation]	msf-target-route-thru-queue=false	In the Application object, configure the option on the Options tab.
DN (Virtual Queue or ACD Queue), Switch	[gim-etl]	q-answer-threshold-voice=60	Overrides the same option that is configured at the application level. In the Switch or DN object, configure the options on the Annex tab.
DN (Virtual Queue or ACD Queue)	[gim-etl]	link-msf-userdata=false	In the DN object, configure the option on the Annex tab.

Table 14: Genesys Info Mart Data-Related Options (Continued)

Configuration Object	Section Name	Option Name and Default Value	Comments
Queue activity data for multimedia (Virtual Queue, Interaction Queue, or Interaction Workbin)			
Genesys Info Mart Application	[gim-etl-media-<media type>] where <media type> exactly matches the name of the applicable Media Type Business Attribute (for example, email)	q-answer-threshold=60 q-short-abandoned-threshold=10	Each [gim-etl-media-<media type>] section contains options that apply for the specific media type only. These settings at the application level can be overridden at the tenant level. In addition, for each specific media type, the q-answer-threshold setting at the application level can be overridden for Virtual Queues at the Switch or DN object levels or, for Interaction Queues and Workbins, at the Script level. In the Application object, configure the options on the Options tab.
Genesys Info Mart Application	[gim-etl-populate]	populate-mm-ixnqueue-facts=false populate-mm-workbin-facts=true populate-workbin-as-hold=false	The populate-mm-ixnqueue-facts and populate-mm-workbin-facts settings can be overridden on the Script object level. In the Application object, configure the options on the Options tab.

Table 14: Genesys Info Mart Data-Related Options (Continued)

Configuration Object	Section Name	Option Name and Default Value	Comments
Genesys Info Mart Application	[gim-transformation]	adjust-vq-time-by-strategy-time=false canceled-queues=iWD_Canceled completed-queues=iWD_Completed	In the Application object, configure the option on the Options tab.
Tenant Media Type Business Attribute	[gim-etl-media]	q-answer-threshold=60 q-short-abandoned-threshold=10	For the particular Media Type attribute, which is configured for a particular tenant, these settings override the same options that are configured at the application level. In addition, for each specific media type, the q-answer-threshold setting at the tenant level can be overridden on the DN or Script object levels. In the Media Type Business Attribute object for the tenant, configure the options on the Annex tab.

Table 14: Genesys Info Mart Data-Related Options (Continued)

Configuration Object	Section Name	Option Name and Default Value	Comments
DN (Virtual Queue), Script (Interaction Queue or Interaction Workbin), Switch	[gim-etl-media-<media type>] where <media type> exactly matches the name of the applicable Media Type Business Attribute (for example, email)	q-answer-threshold=60	Each [gim-etl-media-<media type>] section contains options that apply for the specific media type only. This setting overrides the same option that is configured at the application or tenant level. In the Switch or DN object for a Virtual Queue or in the Script object for an Interaction Queue or a Workbin, configure the option on the Annex tab.
Script (Interaction Queue or Interaction Workbin)	[gim-etl-populate]	populate-mm-ixnqueue-facts=false populate-mm-workbin-facts=true	Overrides the same option that is configured at the application level. In the Script object for an Interaction Queue or a Workbin, configure the options on the Annex tab.
DN (Virtual Queue) or Script (Interaction Queue or Interaction Workbin)	[gim-etl]	link-msf-userdata=false	In the DN or Script object, configure the option on the Annex tab.
Outbound Contact data			
Genesys Info Mart Application	[gim-etl]	max-camp-group-session-duration-in-hours=168 max-camp-group-state-duration-in-hours=168 max-chain-processing-duration-in-hours=8	In the Application object, configure the options on the Options tab.

Operations-Related Configuration Options

Table 15 on page 248 lists the configuration options, grouped by functional area, that affect the operations of the Genesys Info Mart components. These options control the ETL process, and most of them apply in any environment.

Review the options to identify the ones that are relevant to the Genesys Info Mart operation in your environment. In particular, decide:

- For new deployments of Genesys Info Mart release 8.1.103.07 or later, from which date you want Genesys Info Mart to start extracting data after the Info Mart database is initialized. (See "Startup".)
- How you want Genesys Info Mart to schedule launching of ETL jobs. (See "Scheduling" on page 249.)
- How you want Genesys Info Mart to process error conditions during data transformation. (See "Transformation error handling" on page 249.)
- What logging level is sufficient for Genesys Info Mart Server. (See "Logging" on page 250.)
- How you want to adjust ETL behavior to meet requirements for data quality and data availability, given the characteristics of your environment and interaction flows. (See "Miscellaneous" on page 250.)
- Starting with release 8.1.2, which, if any, of the redundant IDBs in an HA set you want to identify as preferred for the purposes of extraction. (See "Miscellaneous" on page 250.)
- How you want Genesys Info Mart to optimize Genesys Info Mart Server performance. (See "Performance tuning" on page 251.)
- What data-retention policies and other purge-related strategies you want to implement. (See "Database maintenance" on page 251.)
- What calendar dimensions you need for your reports. (See "Calendar maintenance" on page 251.)

Refer to Chapter 14 on page 271, as required, to understand how a particular configuration option works and what functionality it enables. Finally, review the valid values, which are documented as part of the option descriptions, and determine the values that are appropriate to your environment.

Table 15: Genesys Info Mart Operations-Related Options

Configuration Object	Section Name	Option Name and Default Value	Comments
Startup			
Genesys Info Mart Application	`[gim-etl]`	`etl-start-date`=Initialization date minus `30` days	In the `Application` object, configure the options on the `Options` tab.

Table 15: Genesys Info Mart Operations-Related Options (Continued)

Configuration Object	Section Name	Option Name and Default Value	Comments
Scheduling			
Genesys Info Mart Application	[schedule]	timezone=GMT run-scheduler=false etl-frequency=1 etl-start-time=06:00 etl-end-time=22:00 run-aggregates=false aggregate-schedule=0 1 aggregate-duration=05:00 run-maintain=true maintain-start-time=03:00 run-update-stats=false update-stats-schedule=0/10 *	In the Application object, configure the options on the Options tab.
Transformation error handling			
Genesys Info Mart Application	[error-policy]	error-policy-call-mergecall-missing=resume error-policy-campaign-group-missing=exception error-policy-islink-dangling=resume error-policy-islink-multiple-targets=resume error-policy-islink-multiple-sources=resume error-policy-islink-multiple-vertices=resume error-policy-islink-source-party-missing=resume error-policy-ipurpose-numberformat=resume error-policy-party-created-missing=resume error-policy-party-created-duplicated=resume error-policy-party-parent-missing=resume error-policy-irf-exception=log_db_resume error-policy-irf-exception-resumable=Exception	In the Application object, configure the options on the Options tab.

Table 15: Genesys Info Mart Operations-Related Options (Continued)

Configuration Object	Section Name	Option Name and Default Value	Comments
Logging			
Genesys Info Mart Application	[log4j]	log4j.appender.ConsoleLogger.Threshold=info logging-level=info log-file-name=gim_etl.log max-log-file-size=50MB max-backup-index=10 console-pattern-layout=%d{ISO8601} %-5p %-12t %m%n file-pattern-layout=%d{ISO8601} %-5p %-12t %m%n	In the Application object, configure the options on the Options tab.
	[log]	standard=network verbose=standard	
Miscellaneous			
Genesys Info Mart Application	[gim-etl]	delayed-data-threshold=900 extract-data-stuck-threshold=28860 extract-last-second=false (option discontinued in release 8.1.4) max-call-duration=3600 max-time-deviation=30 memory-threshold=0 merge-failed-is-link-timeout=0 user-event-data-timeout=3600 **Note:** The primary purpose of these options is to control aspects of Genesys Info Mart behavior that affect data quality or data availability. However, these options also have an impact on operating performance (for example, because they influence the size of merge tables or staging tables).	In the Application object, configure the options on the Options tab.
	[gim-transformation]	pipeline-timeout-in-hours=1	
Info Mart and extraction DAPs	[gim-etl]	geo-location=""	In the DAP object, configure the option on the Options tab.

Table 15: Genesys Info Mart Operations-Related Options (Continued)

Configuration Object	Section Name	Option Name and Default Value	Comments
Performance tuning			
Genesys Info Mart Application	[gim-etl]	extract-data-chunk-size=900 extract-data-cfg-facts-chunk-size=90000 extract-data-max-conn=128 extract-data-thread-pool-size=32 merge-chunk-size=200000 max-chunks-per-job=10	In the Application object, configure the options on the Options tab.
	[gim-transformation]	irf-io-parallelism=4 ud-io-parallelism=5	
Database maintenance			
Genesys Info Mart Application	[gim-etl]	days-to-keep-gidb-facts=14 days-to-keep-gim-facts=400 days-to-keep-active-facts=30 (600 in release 8.1.0) days-to-keep-deleted-annex=2 days-to-keep-discards-and-job-history=600 purge-transaction-size=100000 purge-thread-pool-size=32 For partitioned databases only: partitioning-interval-size-gidb=86400 partitioning-interval-size-gim=86400 (604800 in release 8.1.0) partitioning-ahead-range=14	In the Application object, configure the options on the Options tab.
Calendar maintenance			
Genesys Info Mart Application	[date-time] For custom calendars: [date-time-*]	date-time-table-name=DATE_TIME date-time-tz=GMT date-time-start-year=2012 (2010 in release 8.1.0) date-time-min-days-ahead=183 date-time-max-days-ahead=366 simple-week-numbering=true first-day-of-week=1 min-days-in-first-week=1 fiscal-year-week-pattern=none fiscal-year-start=	In the Application object, configure the options on the Options tab.

Aggregation-Related Configuration Options

Table 16 lists the configuration options that are essential to enable aggregation in your deployment, if GI2 or RAA is installed.

Note: There are additional aggregation-related options that are set on the Genesys Info Mart `Application` object, to control the functioning of the aggregation engine and population of the aggregate tables. GI2 has specific requirements for certain option values.

In addition to specifying the aggregation engine class name (in a configuration option), you must also modify the `gim_etl_paths` file to specify the path to the correct `.jar` file.

For more information about all of the aggregation-related options and installation steps, see the *Reporting and Analytics Aggregates 8.1 Deployment Guide.*

Table 16: Genesys Info Mart Options for GI2 and RAA

Section Name	Option Name and Recommended Value	Comments
`[schedule]`	`run-aggregates=true` `aggregate-schedule=`<as specified in the aggregation-specific template> `aggregate-duration=`<as specified in the aggregation-specific template>	
`[gim-etl]`	`aggregation-engine-class-name= "GIMAgg.GimInterfaceImpl.AggregationImpl"`	

Chapter

13 Configuring Supporting Objects

Both Interaction Concentrator (ICON) and Genesys Info Mart require certain settings on other Genesys configuration objects, such as `DN` and `Field` objects, in order to support specific functionality. This chapter describes how to configure supporting objects to meet ICON and Genesys Info Mart requirements. It contains the following sections:

Introduction

Configuration settings on the following configuration objects affect ICON processing in ways that are significant for Genesys Info Mart:

- `Switch` (for Voice or Multimedia details)
- `DN` (for Voice or Multimedia details)
- `Field` (for Outbound Contact details)

Configuration settings on the following configuration objects affect Genesys Info Mart data processing:

- `Switch` (for Voice or Multimedia details)
- `Media Type Business Attribute` (for Multimedia details)
- `DN` (for Voice or Multimedia details)
- `Script` (for Multimedia details)

- Field (for Outbound Contact details)

Use the instructions in this chapter to configure the appropriate settings for your environment.

Annex Tab

In all cases, the options that ICON or Genesys Info Mart require are specified on the Annex tab of the applicable configuration objects. By default, the Annex tab for any configuration object is not displayed in Configuration Manager. The following procedure describes how to set the Annex tab to display, so that you can configure the required options on the various objects.

Procedure: Setting the Annex tab to display

Purpose: To display the Annex tab in the Properties dialog box for an object in Configuration Manager.

Start of procedure

1. Select View > Options.
2. Select the Show Annex tab in Object Properties check box.
3. Click OK.

End of procedure

Configuring Switch Objects

Settings on the Switch object affect how:

- ICON gathers and reports data about voice or multimedia interactions, including interaction, user-data, resource, and agent activity details.
- Genesys Info Mart processes network activity, for all types of ICON details.
- Genesys Info Mart reports do-not-disturb (DND) activity.

The following procedure and Table 17 on page 256 describe the required configuration settings.

Procedure: Configuring the switch for ICON and Genesys Info Mart reporting

For Multimedia details, the only ICON-related option that applies relates to virtual queues.

Prerequisites:

- The `Annex` tab must be displayed for configuration objects in Configuration Manager. For more information, see Procedure: Setting the Annex tab to display, on page 254.

Start of procedure

1. Open the `Annex` tab of the `Switch` object that handles voice or multimedia interactions.

Configure ICON-Related Options

2. Create a new section that is named `gts`, if it does not exist already on the `Annex` tab.
3. Open the `gts` section.
4. Configure the required ICON-related options:
   - For Voice details, configure all required options that are described under "ICON-related options—gts section" in Table 17 on page 256.
   - For Multimedia details, if your deployment includes `Virtual Queue` DNs, verify that the `support-dn-type-5` option is set to the default value of `1` (see page 258). By default, ICON will monitor and store data for all virtual queues that belong to this switch.

   For more information, see the chapter about configuration options in the *Interaction Concentrator 8.x Deployment Guide*.

Configure Genesys Info Mart–Related Options

5. Create a new section that is named `gim-etl`, if it does not exist already on the `Annex` tab.
6. Open the `gim-etl` section.
7. Configure the required Genesys Info Mart–related options:
   - In a network deployment, configure the `network-switch` option.
   - For Voice or Multimedia details, configure the `factor-dnd-into-sm-resource-states` option.

     For more information, see the option descriptions under "Genesys Info Mart–related options—gim-etl section" in Table 17 on page 256.
   - (Optional) For Voice details, specify a value for the queue answer threshold to override the application-level setting. For more information, see the description of the `q-answer-threshold-voice` option on page 299.

8. (Optional) For Multimedia details, specify media-specific values for queue answer thresholds to override application- or tenant-level settings, as required. To specify a media-specific threshold at the switch level, you must first create the applicable `[gim-etl-media-<media type>]` section on the `Annex` tab. For more information, see the description of the `q-answer-threshold` option on page 302.

Table 17: Switch Options

Area of Functionality	Option Name	Recommended Value	Description
ICON-related options—gts section			
Agent state and login session	gls-associations-rule	`0` **Note:** `0` is the only valid value for Genesys Info Mart.	ICON creates a single login session for two DNs that belong to the same place when an agent logs in at one of these DNs.
	gls-flag-on-disconnect	`0` (default)	When reconnecting to T-Server, ICON compares the agent state from its memory with the state from `EventRegistered`. If the in-memory state does not match the currently reported agent state, ICON updates the agent state in both its internal memory and IDB. When disconnecting from T-Server, ICON performs no actions that are specific to agent states.
	gls-use-ts-id	1 (default) **Note:** 1 is the only valid value for Genesys Info Mart.	ICON uses the login session ID that is generated by T-Server (GUID).

Table 17: Switch Options (Continued)

Area of Functionality	Option Name	Recommended Value	Description
Agent metrics	gls-acw-first	To ensure that ICON associates ACW with the first voice interaction, do one of the following: • At the switch level, set this option value to `1`. • Retain the default value of `-1` at the switch level, and set the `gls-acw-first` option to `true` at the ICON application level (see page 171).	Specifies, for this switch, which interaction ICON associates with ACW. For accurate measurement of ACW-related metrics, Genesys Info Mart requires ICON to associate ACW metrics with the voice interaction that immediately precedes the *start* of the ACW (the first voice interaction). Subsequent voice interactions are considered to be related to ACW processing. When the agent logs out, changes his or her state to `Ready`, or goes `NotReady` for any reason other than to perform ACW, ICON reports the end of the current ACW state. This option overrides an explicit setting of the `gls-acw-first` configuration option (described on page 171) at the ICON application level.

Table 17: Switch Options (Continued)

Area of Functionality	Option Name	Recommended Value	Description
Agent metrics	gls-enable-acw-busy	No recommended value, but the setting might affect reporting results: • `0`—ICON does not interrupt `ACW` and `NotReady` agent states. • `1` (default)—ICON interrupts `ACW` and `NotReady` agent states. **Note:** If you set the option value to `0` (false) because you want ICON to report uninterrupted `ACW`, ICON will also report uninterrupted `NotReady` states. Other states that might occur during `NotReady` (for example, `Busy`) will not be reported in IDB or the Info Mart database. For more information, see the section about populating agent activity data in the *Genesys Info Mart 8.1 User's Guide*.	Specifies whether ICON interrupts `ACW` and `NotReady` agent states when an agent places or receives another interaction during the period of time that `ACW` or `NotReady` agent states were invoked. For more information about uninterrupted `ACW` and `NotReady` states, see the chapter about agent states and login sessions in the *Interaction Concentrator 8.x User's Guide.* **Note:** This option is not valid for Session Initiation Protocol (SIP)–compliant switches that handle interactions other than voice interactions.
Virtual queue	support-dn-type-5	`1` (default)	ICON processes events that pertain to DNs of the `Virtual Queue` type that belong to this switch.
Parallel ACD queues	third-party-queue-in-divert	• `0` (default) • `1` (for switches, such as Avaya Communication Manager release 7.6 and later, that enable T-Server to supply `AttributeThirdPartyQueue` in `EventDiverted`)	Controls whether ICON uses the `AttributeThirdPartyQueue` attribute in `EventDiverted` to more accurately represent the party from which a call is distributed, in certain multi-queue scenarios. The option was introduced in Interaction Concentrator release 8.1.000.34.

Table 17: Switch Options (Continued)

Area of Functionality	Option Name	Recommended Value	Description
Genesys Info Mart–related options—gim-etl section			
Scenario recognition	network-switch	• FALSE (default) • TRUE	Specifies whether the switch is a network switch. Changes take effect at the next run of Job_ExtractICON.
Agent metrics	factor-dnd-into-sm-resource-states	No recommended value, but the setting affects reporting results. • FALSE (default for voice and SIP switches)—Genesys Info Mart does not factor DND into summarized resource states and reasons. • TRUE (default for eServices/Multimedia switches)—DND is factored into summarized resource states and reasons.	Specifies whether DND (do-not-disturb) status for a resource on a given switch is factored into resource states and reasons in SM_RES_STATE_FACT and SM_RES_STATE_REASON_FACT tables. Changes take effect on the next ETL cycle. The new option value is not applied to previously loaded facts.

End of procedure

Next Steps

- (Optional) For Multimedia details, Procedure: Setting Media Type Business Attribute object options for Genesys Info Mart reporting, on page 260
- For Voice or Multimedia details, Procedure: Configuring a DN for ICON and Genesys Info Mart reporting, on page 261
- For Outbound Contact details, “Configuring Field Objects” on page 266

Configuring Media Type Business Attribute Objects

Settings on Media Type Business Attribute objects affect Genesys Info Mart reporting of queue metrics for multimedia interactions for a tenant.

Procedure: Setting Media Type Business Attribute object options for Genesys Info Mart reporting

Purpose: To override certain application option settings to implement media-specific thresholds for multimedia interactions for a specific tenant.

Note: For the options that you configure in a `Media Type Business Attribute` object, changes take effect on the next ETL cycle. The new option value is not applied to previously loaded facts.

Prerequisites

- You have the required access privileges to modify properties for `Media Type Business Attribute` objects in the Genesys Configuration Layer.
- The `Annex` tab must be displayed for configuration objects in Configuration Manager. For more information, see Procedure: Setting the Annex tab to display, on page 254

Start of procedure

1. For the media type for which you want to configure tenant-specific thresholds, display the properties for the `Media Type Business Attribute` object for that tenant.
2. Add a new section, named `gim-etl-media`, on the `Annex` tab.
3. Open the new section, and add one or more of the following options, with values that are suitable for that tenant, for the media type:
   - `q-answer-threshold` (see page 302)
   - `q-short-abandoned-threshold` (see page 302)
   - `short-abandoned-threshold` (see page 303)

   Note: The values that are set for these options in a specific `Media Type Business Attribute` object will override values that are set in the Genesys Info Mart `Application` object.

   The value of the `q-answer-threshold` option that is set at the tenant level can be overridden in the `Switch` or individual `DN` objects (for Virtual Queues) or in individual `Script` objects (for Interaction Queues or Interaction Workbins).

End of procedure

Next Steps

- For Voice or Multimedia details, Procedure: Configuring a DN for ICON and Genesys Info Mart reporting
- (Optional) For Multimedia details, Procedure: Setting Script object options for Genesys Info Mart reporting, on page 265
- For Outbound Contact details, "Configuring Field Objects" on page 266

Configuring DN Objects

Settings on DN objects affect:

- How ICON gathers and reports data on IVR and virtual-queue usage for voice or multimedia interactions.
- Genesys Info Mart reporting of queue metrics for voice or multimedia interactions: Options on Virtual Queue or ACD Queue DNs enable you to override certain application-level—or, for multimedia interactions, tenant-level—option settings in a specific DN object, as appropriate.
- Starting with release 8.1.2, whether Genesys Info Mart stores user data for interactions that are in mediation.

The following procedure and Table 18 on page 263 describe the required configuration settings.

Procedure: Configuring a DN for ICON and Genesys Info Mart reporting

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

Prerequisites:

- The Annex tab must be displayed for configuration objects in Configuration Manager. For more information, see Procedure: Setting the Annex tab to display, on page 254.

Start of procedure

1. If you are deploying Genesys Info Mart to report on both ICON Voice details and ICON Multimedia details, make sure that any DN objects for virtual queues that are dedicated to voice calls are configured under the

Switch object that is configured for your traditional telephony switch. Otherwise, the ICON application that is dedicated to handling Voice details cannot track activity on virtual queues.

Configure ICON-Related Options

2. Open the Annex tab of the DN object that handles voice or multimedia interactions.
3. Create a new section that is named gts, if it does not exist already on the Annex tab.
4. Open the gts section.
5. Configure the required ICON-related options:
   - For Voice details, configure all required options that are described in Table 18 on page 263.
   - For Multimedia details, if your deployment includes virtual queues, configure only the monitor option (see page 264).

   For more information, see the chapter about configuration options in the *Interaction Concentrator 8.x Deployment Guide*.

Configure Genesys Info Mart–Related Options

6. If you want to use DN-specific settings for queue metrics, in Configuration Manager, display the properties for the DN object of type Virtual Queue or ACD Queue. Then:
   - To configure thresholds for voice interactions:
     i. On a Virtual Queue or ACD Queue DN for voice interactions, add a new section, named gim-etl, on the Annex tab.
     ii. Open the new section, and add the q-answer-threshold-voice option with the desired value (see page 299).

     Note: The value that is set for this option in a specific DN object of the Virtual Queue or ACD Queue type will override a value that is set in the Genesys Info Mart Application object.

   - To configure media-specific thresholds for multimedia interactions:
     i. For each media type for which you want to configure a custom threshold, add a new section, named gim-etl-media-<media type>, on the Annex tab of a Virtual Queue DN for multimedia interactions.

     The <media type> that you specify in the section name must match the name of the Media Type Business Attribute exactly, including case (for example, email).

ii. Open the new section, and add the `q-answer-threshold` option with the desired value (see page 302).

Note: The value that is set for this option in a specific `DN` object of the `Virtual Queue` type will override a value that is set in the Genesys Info Mart `Application` object or in the tenant-specific `Media Type Business Attribute` object.

7. (For release 8.1.2 and later) For every `Virtual Queue` or `ACD Queue` DN for which you want mediation segment facts (MSFs) to store associated user data:
   a. Create or open the `[gim-etl]` section on the `Annex` tab.
   b. Add the `link-msf-userdata` option, and set the value to `true`. (The default value is `false`.)

Table 18: DN Options

Area of Functionality	Option Name	Recommended Value	Description
ICON-related options—gts section			
Scenario recognition	ivr	`0` (default)	ICON does not recognize this DN as an IVR resource unless the DN configuration satisfies one of the following criteria: • DN has a type of `Voice Treatment Port` in Configuration Database. • DN has a type of `ACD Position` or `Extension`, and it is specified as an `Associated DN` in the properties of the IVR port at ICON startup time. **Note:** See the `gls-ivr` option description on page 175 to exclude data about agent activity that is associated with this IVR device from IDB storage.

Table 18: DN Options (Continued)

Area of Functionality	Option Name	Recommended Value	Description
Network Routing Solution	emulate-event-queued	• 0—EventQueued is not emulated. • 1—EventQueued is emulated. **Note:** Network Routing Solution uses Service Number DNs. There is no EventQueued-related option at the switch level for Service Numbers. If the Network T-Server does not generate EventQueued, set this option to 1 on the Service Number DNs.	Enables the emulation of EventQueued for this particular DN. **Note:** Generation of EventQueued for a Routing Point, Routing Queue, and External Routing Point depends on the particular T-Server and its switch. ICON requires this event for correct party representation in any environment.
Virtual queue	monitor	1 (default)	ICON processes virtual queue–related events for this DN. **Note:** This option is meaningful only when the support-dn-type-5 option (see page 258) is set to 1 (default) in the corresponding Switch object.

End of procedure

Next Steps

- (Optional) For Multimedia details, Procedure: Setting Script object options for Genesys Info Mart reporting, on page 265.
- For Outbound Contact details, "Configuring Field Objects" on page 266.
- If you have not already done so, configure or complete the preparation of Interaction Concentrator and Genesys Info Mart. For more information, see Chapter 9 on page 159 and Chapter 12 on page 231.

Configuring Script Objects

Settings on Script objects affect Genesys Info Mart reporting of Interaction Queue or Interaction Workbin metrics for multimedia interactions.

Procedure: Setting Script object options for Genesys Info Mart reporting

Purpose: To override certain application-level or tenant-level option settings in a specific `Script` object, as appropriate, or to enable storage of user data for multimedia interactions that are in Interaction Queues or Interaction Workbins.

Note: For the options that you configure in a `Script` object, changes take effect on the next ETL cycle. The new option value is not applied to previously loaded facts.

Prerequisites

- You have the required access privileges to modify properties for `Script` objects in the Genesys Configuration Layer.
- The `Annex` tab must be displayed for configuration objects in Configuration Manager. For more information, see Procedure: Setting the Annex tab to display, on page 254.

Start of procedure

1. In Configuration Manager, display the properties for the `Script` object that corresponds to the Interaction Queue or Interaction Workbin.
2. For each media type for which you want to configure a custom threshold:
   a. Add a new section, named `gim-etl-media-<media type>`, on the `Annex` tab.

      The `<media type>` that you specify in the section name must match the name of the `Media Type Business Attribute` exactly, including case (for example, `email`).
   b. Open the new section, and add the `q-answer-threshold` option with the desired value (see page 302).

   Note: The value that is set for this option in a specific `Script` object will override a value that is set in the Genesys Info Mart `Application` object or in the tenant-specific `Media Type Business Attribute` object.
3. Add a new section, named `gim-etl-populate`, on the `Annex` tab.

4. Add either of the following options, and specify the value that you want for this particular Interaction Queue or Interaction Workbin (see page 304):
   - `populate-mm-ixnqueue-facts`
   - `populate-mm-workbin-facts`

   Note: The value that is set for options in a specific `Script` object will override a value that is set for corresponding options in the `gim-etl-populate` section of the Genesys Info Mart `Application` object.

5. (For release 8.1.2 and later) For every Interaction Queue or Interaction Workbin for which you want MSFs to store associated user data:
   a. Open the `[gim-etl]` section on the `Annex` tab.
   b. Add the `link-msf-userdata` option, and set the value to `true`. (The default value is `false`.)

End of procedure

Next Steps

- For Outbound Contact details, "Configuring Field Objects".
- If you have not already done so, configure or complete preparation of Interaction Concentrator and Genesys Info Mart. For more information, see Chapter 9 on page 159 and Chapter 12 on page 231.

Configuring Field Objects

To enable Genesys Info Mart to report on Outbound Contact activity, you must:

1. Configure `Field` objects so that Outbound Contact Server (OCS) will send the required data and ICON will store it. For more information, see Procedure: Configuring the storage of OCS record field data.
2. Map `Field` objects to columns in the Info Mart database. For more information, see Procedure: Configuring the mapping of OCS record fields, on page 268.

For more information about Outbound Contact data in Genesys Info Mart, see Chapter 7 on page 127.

Procedure: Configuring the storage of OCS record field data

Purpose: To enable the `Field` object settings that are required for ICON to store mandatory field data in IDB.

Prerequisites

- You have the required access privileges to modify properties for `Field` objects in the Genesys Configuration Layer.
- The `Annex` tab must be displayed for configuration objects in Configuration Manager. For more information, see Procedure: Setting the Annex tab to display, on page 254.

Start of procedure

1. Review the configuration options in Table 19 on page 267 that you should configure in the `default` section on the `Annex` tab of the `Record Field` object.
2. For each mandatory field that is listed in Table 8 on page 129, add the `icon_attribute` option to the `default` section on the `Annex` tab of the `Field` object. For information about recommended values and restrictions, see Table 19.
3. For each nonmandatory (custom) field that you want Genesys Info Mart to store in its database, add the `icon_attribute` option to the `default` section on the `Annex` tab of the `Field` object. For information about recommended values and restrictions, see Table 19.

> **Note:** You do not have to configure the `send_attribute` option for any fields.

Table 19: Field Object—Annex Options—default Section

Option Name	Recommended Value	Description
icon_attribute	• For nonsensitive data, set this option to `1` to store the data in the IDB `GO_CUSTOM_FIELDS` and `GO_FIELDHIST` tables. • For sensitive data, set this option to `2` to store the data in the IDB `GO_SECURE_FIELDS` and `GO_SEC_FIELDHIST` tables.	Instructs ICON as to whether or not to report on this field and which tables to use to store the record field data.

Table 19: Field Object—Annex Options—default Section (Continued)

Option Name	Recommended Value	Description
right_person	(Optional) Specify a value that indicates that the right person is contacted. **Examples:** • TRUE • YES • 1	Specifies the value of the record field when the right person is contacted. If the value of the Right Person Contacted field matches this option value (case-insensitive), Genesys Info Mart sets the RPC_FLAG in its CONTACT_ATTEMPT_FACT table to 1. **Note:** Do not specify this option for more than one field within a calling list.
conversion	(Optional) Specify a value that indicates that a transaction was successful.	Marks the field that indicates that the answered call was a successful transaction. If the field is updated with a value that equals the value that is configured for this option, Genesys Info Mart recognizes it as a conversion indicator. **Note:** This option should not be specified for more than one field within a calling list.

End of procedure

Next Steps

- Procedure: Configuring the mapping of OCS record fields

Procedure: Configuring the mapping of OCS record fields

Purpose: To enable the Field object settings that are required for mapping to the Genesys Info Mart database.

To determine what fields you need to map, see "Worksheet for Mapping OCS Record Fields" on page 432. Configure the mapping only for those nonmandatory Field objects that you want to store in the Info Mart database. Follow this procedure to configure the Annex tab of the corresponding Field object, to indicate the Info Mart table name and column name to which it will be mapped.

Prerequisites

- The Annex tab must be displayed for configuration objects in Configuration Manager. For more information, see Procedure: Setting the Annex tab to display, on page 254.
- Procedure: Configuring the storage of OCS record field data, on page 267.
- Mapping requirements must be finalized, as per "Worksheet for Mapping OCS Record Fields" on page 432.

Start of procedure

1. Review the configuration options in Table 19 on page 267 and Table 20 on page 270.
2. On the Annex tab of the Field object, add a section that is named gim-etl-mapping.

Notes: When you configure options in the gim-etl-mapping section:
- Do not map more than one field to the same Info Mart table and column.
- Do not configure options for extra Info Mart table columns that will not store Field object data.

3. In the gim-etl-mapping section, add an option that is named table-name, and set its value to the name of the Info Mart table in which you want the ETL to store this field (see Table 20 on page 270).
4. In the gim-etl-mapping section, add an option that is named column-name, and set its value to the name of the column in the Info Mart table in which you want the ETL to store this field (see Table 20).
5. In addition, Genesys Info Mart supports the designation of nonmandatory fields to indicate right party contacted and conversion:
   - For a field to indicate that the correct party was contacted, add the right_person option (see page 268) to the default section on the Annex tab of the Field object, and set its value to the field value that you want to indicate that the right party has been contacted (for example, TRUE).
   - For a field to indicate conversion, add the conversion option (see page 268) to the default section on the Annex tab of the Field object, and set its value to the field value that you want to indicate that a conversion has taken place (for example, TRUE).

Table 20: Field Object—Annex Options—gim-etl-mapping Section

Option Name	Recommended Value	Description
table-name	The name of one of the following Info Mart tables: • CONTACT_ATTEMPT_FACT • RECORD_FIELD_GROUP_1 • RECORD_FIELD_GROUP_2	Specifies the name of the Info Mart table in which the ETL should store the value of this field. The option value is case insensitive. **Note:** Configure this option only for the nonmandatory Field objects that you want to store in the Info Mart database.
column-name	The name of a column (in the table that is specified by the table-name option) that is not used in the configuration of any other Field object. The type of the OCS field must match the data type of the column. Refer to "Worksheet for Mapping OCS Record Fields" on page 432 for the list of column names.	Specifies the name of the column in the selected Info Mart table in which the ETL should store the value of this field. The option value is case insensitive. **Note:** Configure this option only for the nonmandatory Field objects that you want to store in the Info Mart database.

End of procedure

Next Steps

- If you have not already done so, configure or complete preparation of Interaction Concentrator and Genesys Info Mart. For more information, see Chapter 9 on page 159 and Chapter 12 on page 231.
- "Preparing the Genesys Info Mart Server Host" on page 323.

Chapter

14 Genesys Info Mart Configuration Options Reference

This chapter describes the configuration options that you can set for effective operation of your Genesys Info Mart application.

For ease of reference, the Genesys Info Mart `Application` object configuration options are listed in alphabetical order by configuration section.

This chapter contains the following sections:

date-time Section

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

`Job_InitializeGIM` populates data in all configured calendars when it initializes the Info Mart database. `Job_MaintainGIM` subsequently maintains the calendars in accordance with options that are specified in the `[date-time]` and custom `[date-time-*]` configuration sections.

The maintenance job automatically adjusts for special requirements such as daylight saving time (DST) and fiscal years that do not start on the same day every year (floating fiscal years).

Note: Consider the settings for the `date-time` options carefully before the calendar dimension tables are populated for the first time. You can subsequently change the values of the `date-time-min-days-ahead` and `date-time-max-days-ahead` options at any time. However, changing any of the other `date-time` options during runtime can introduce inconsistencies into the calendar data and affect reporting results adversely. For example, if you change the timezone option (`date-time-tz`) after Genesys Info Mart has been initialized, your reports might mix the results for different timezones within the same reporting interval.

If you want to change calendar options during runtime, see the procedure about changing calendar options in the *Genesys Info Mart 8.1 Operations Guide,* which provides information about additional steps that are required to maintain reporting consistency.

date-time-max-days-ahead

Default Value: `366`
Valid Values: Any positive integer
Dependencies: None
Changes Take Effect: At the next run of `Job_MaintainGIM`

Specifies, in number of days, how far ahead the calendar dimension table will be populated. The default value specifies that the calendar dimension will be populated up to a year in advance (365 days + 1 day for leap years). Genesys does not recommend that you populate the calendar tables more than a year in advance, in case there are changes to DST or other international time standards that might invalidate the prepopulated data.

Note: Ensure that you populate the calendar far enough ahead to meet the requirements of your reporting intervals.

date-time-min-days-ahead

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

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

date-time-start-year

Default Value: 2012
Valid Values: 1970–2038
Dependencies: None
Changes Take Effect: At the next run of Job_MaintainGIM

Specifies the year that the calendar starts.

When you are setting this option, ensure that you choose a start year that provides sufficient buffer to prevent inconsistencies or unexpected missing dimensions around the start of the calendar. Genesys recommends that you set the value so that the calendar starts at least one year prior to any date that might be encountered in the data.

Be aware that Genesys Info Mart uses GMT for internal time references, and this affects exactly when the calendar starts. For example, if the other [date-time] options that affect the start date are set so that the calendar will start at 00:00 AM on January 1, 2012, and the date-time-tz option is set to Eastern European Time (GMT + 2), the calendar table will be populated with dimensions starting at 02:00 AM on January 1, 2012.

In release 8.1.0, the default value of this option was 2010.

date-time-table-name

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

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

date-time-tz

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

Specifies the time zone for the calendar. You can use any valid time zone that is supported by the version of the Java Runtime Environment (JRE) that runs the Genesys Info Mart Server.

For more information about supported time zones, see the documentation about calendar time zones on the Java developer website or other public resources—for example:

- http://www.java2s.com/Tutorial/Java/0120__Development/GettingallthetimezonesIDs.htm

- http://en.wikipedia.org/wiki/Zone.tab

Note: Particularly in deployments that use Genesys Interactive Insights (GI2) or the separately installed Reporting and Analytics Aggregates (RAA) package, ensure that the time zone is set appropriately for your deployment before you initialize Genesys Info Mart or before aggregation starts.

first-day-of-week

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

Specifies the day of the week that is considered to be the start of the week. For example, 1 (Sunday) is usually the first day of the week in the United States; for countries that use the ISO 8601 standard, 2 (Monday) is the first day of the week.

fiscal-year-start

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

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

If simple-week-numbering=true, every fiscal year starts on the fixed date that is specified by this option.

If simple-week-numbering=false, the fiscal year starts on the first day of the week that contains the date that is specified by this option; however, the actual start date depends on the value of the first-day-of-week option. Genesys Info Mart adjusts automatically for the floating fiscal year. For example, if simple-week-numbering=false, fiscal-year-start=3-1, and first-day-of-week=1, then:

- Fiscal year 2012 starts on February 26.
- Fiscal year 2013 starts on February 24.
- Fiscal year 2014 starts on February 23.

fiscal-year-week-pattern

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

Specifies the pattern for the number of weeks in each month of a fiscal quarter. For example, `544` means 5 weeks in the first month, 4 weeks in the second month, and 4 weeks in the third month of each quarter.

A value of `none` means that the calendar will not be a fiscal one.

min-days-in-first-week

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

Specifies the minimum number of days from the new year that must be in the first week of the year, if simple week numbering is not used and there are no partial weeks in the calendar year. The ISO 8601 standard does not use simple week numbering.

For example, if `simple-week-numbering=false, first-day-of-week=2,` and January 1 of the new year is on a Friday, there are 3 days from the new year in the week that starts on Monday, December 28. Therefore:

- If the value of this option is set to `1`, the calendar will count the first week of the new year as starting on Monday, December 28.
- If the value of this option is set to `4`, the week that starts on Monday, December 28, will be assigned to the previous year, and the calendar will count the first week of the new year as starting on Monday, January 4.

The ISO 8601 definition of the first week in the year is the week that has the first Thursday in it. To conform to the ISO 8601 standard, set `simple-week-numbering=false, first-day-of-week=2,` and `min-days-in-first-week=4`.

simple-week-numbering

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

Specifies whether the calendar year and the week-numbering year coincide. For simple week numbering, Week 1 always begins on the first day of the calendar year (for Gregorian calendars, January 1; for fiscal calendars, the day that is specified in the `fiscal-year-start` option). As a result, the first and last weeks of the year might be partial weeks, because the first week will not necessarily start with the day that is specified by the `first-day-of-week` option (see page 274). To comply with ISO 8601 week numbering, set the value of this option to `false`.

error-policy Section

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

Notes:

- By default, all of the `[error-policy]` options—except for `error-policy-irf-exception`, `error-policy-irf-exception-resumable`, and `error-policy-campaign-group-missing`—are set so as not to generate an exception when the transformation job encounters data inconsistencies. The default settings mean that Genesys Info Mart will attempt to recover from inconsistencies in the source data and continue processing. The implications for data quality depend on the particular call flow and environment. The `STATUS` field in the `INTERACTION_FACT` record indicates the type of error that was encountered.
- The default value for `error-policy-irf-exception` (page 277) is `log_db_resume`. If you set the value of this option to `exception`, Genesys Info Mart will fail the transformation job when it encounters error-policy exceptions. You can set an alarm on the log event that is generated when the job fails.

error-policy-call-mergecall-missing

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

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

- `exception`—Instructs the transformation logic to interrupt transformation of the interaction with an exception, which is handled as specified by the `error-policy-irf-exception` option.
- `resume`—Instructs the transformation logic to ignore references to the missing data and continue with transformation. The transformation job logs the following error message:
  `Interaction(...):call(...): merge call(...) is missing.`

error-policy-campaign-group-missing

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

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

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

error-policy-ipurpose-numberformat

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

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

- exception—Instructs the transformation logic to interrupt transformation of the interaction with an exception, which is handled as specified by the error-policy-irf-exception option.
- resume—Instructs the transformation logic to process the data as if the IPurpose KVP were not attached. In this case, whether the IVR is treated as a handling resource or a mediation resource depends on the value that is configured for the default-ivr-to-self-service option (see page 310).

error-policy-irf-exception

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

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

- log_db_resume (default)—Instructs the transformation logic to discard the problematic interaction thread, write corresponding information into the STG_TRANSFORM_DISCARDS table, and resume processing.
- resume—Instructs the transformation logic to discard the problematic interaction thread and resume processing, without writing corresponding information into the database.
- exception—Instructs the transformation logic to fail the job.

error-policy-irf-exception-resumable

Default Value: Exception
Valid Values: Any correct Java expression

Dependencies: error-policy-irf-exception=log_db_resume or resume
Changes Take Effect: On the next ETL cycle

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

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

- Exception—Genesys Info Mart will behave as described for error-policy-irf-exception=log_db_resume or resume.
- InteractionTransformException—Genesys Info Mart will behave as described for error-policy-irf-exception=log_db_resume or resume.
- NullPointerException—The transformation job will fail.
- IllegalStateException—The transformation job will fail.

error-policy-islink-dangling

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

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

- exception—Instructs the transformation logic to interrupt transformation of the interaction with an exception, which is handled as specified by the error-policy-irf-exception option.
- resume—Instructs the transformation logic to process the interaction as if the missing IS-Link information were for a remote site that is not monitored by ICON. For example, an internal transfer will be transformed as an inbound or outbound interaction.

error-policy-islink-multiple-sources

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

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

- exception—Instructs the transformation logic to interrupt transformation of the interaction with an exception, which is handled as specified by the error-policy-irf-exception option.
- resume—Instructs the transformation logic to choose one of the source records randomly and ignore the other source records.

error-policy-islink-multiple-targets

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

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

- exception—Instructs the transformation logic to interrupt transformation of the interaction with an exception, which is handled as specified by the error-policy-irf-exception option.
- resume—Instructs the transformation logic to choose one of the target records randomly and ignore the other target records.

error-policy-islink-multiple-vertices

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

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

error-policy-islink-source-party-missing

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

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

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

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

error-policy-party-created-duplicated

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

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

- exception—Instructs the transformation logic to interrupt transformation of the interaction with an exception, which is handled as specified by the error-policy-irf-exception option.
- resume—Instructs the transformation logic to treat the first record that it reads as the party created record and to ignore the other party created records.

error-policy-party-created-missing

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

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

- exception—Instructs the transformation logic to interrupt transformation of the interaction with an exception, which is handled as specified by the error-policy-irf-exception option.
- resume—Instructs the transformation logic to construct a party created record, based on assumptions from the first party history record that it reads.

error-policy-party-parent-missing

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

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

- exception—Instructs the transformation logic to interrupt transformation of the interaction with an exception, which is handled as specified by the error-policy-irf-exception option.
- resume—Instructs the transformation logic to ignore the missing data and continue processing.

gim-etl Section

Use this configuration section to set general options.

aggregation-engine-class-name

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

Specifies the class name of the aggregation package, if it is installed. If your deployment uses GI2 or the separately installed RAA package, specify the following value:

`"GIMAgg.GimInterfaceImpl.AggregationImpl"`

For more information, see the *Reporting and Analytics Aggregates 8.1 Deployment Guide*.

days-to-keep-active-facts

Default Value: `30`
Valid Values: Any positive integer
Dependencies: None in releases 8.1.0 and 8.1.1; starting with release 8.1.2, `days-to-keep-gim-facts`
Changes Take Effect: At the next run of `Job_TransformGIM` or `Job_MaintainGIM`

Note: The role of this option has been enhanced in Genesys Info Mart release 8.1.1. The following description applies to behavior starting with release 8.1.1. For the benefit of Genesys Info Mart 8.1.0 customers, Appendix E on page 451 provides information about purge behavior in previous 8.x releases.

Specifies the maximum number of days to retain active multimedia interactions in GIDB and the dimensional model, including certain Staging tables, after which the interactions become eligible for artificial termination.

Artificial Termination of Long-Living Interactions

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

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

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

Purging

Genesys Info Mart does not purge active interactions. However, the relationship of `days-to-keep-active-facts` to `days-to-keep-gidb-facts` and `days-to-keep-gim-facts` affects the purge threshold for GIDB and dimensional model fact data:

- If `days-to-keep-active-facts` is smaller than `days-to-keep-gidb-facts`, Genesys Info Mart terminates the interaction artificially when the `days-to-keep-active-facts` interval expires, then subsequently purges the terminated interaction from GIDB when the `days-to-keep-gidb-facts` interval expires and from the dimensional model when the `days-to-keep-gim-facts` interval expires.
- If `days-to-keep-active-facts` is greater than or equal to `days-to-keep-gidb-facts` but smaller than `days-to-keep-gim-facts`, Genesys Info Mart terminates the interaction artificially when the `days-to-keep-active-facts` interval expires, then purges the terminated interaction from GIDB in the next run of the maintenance job and from the dimensional model when the `days-to-keep-gim-facts` interval expires.
- Starting with release 8.1.2, Genesys Info Mart requires `days-to-keep-active-facts` to be smaller than `days-to-keep-gim-facts`. For earlier releases, Genesys strongly recommends that you configure `days-to-keep-active-facts` to be smaller than `days-to-keep-gim-facts`.

Active interactions can affect the purge threshold for completed facts. Tables 21 and 22, starting on page 285, provide examples of different scenarios for GIDB and the dimensional model, respectively.

Related Information

For a list of the tables in the Info Mart database that contain multimedia interaction data that is affected by this option, see Appendix D on page 445.

For information about the purging rules and guidelines to consider when you are setting the value of this option, see the section about purging rules, starting on page 408.

days-to-keep-deleted-annex

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

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

The default value of `days-to-keep-deleted-annex` is small, because there is likely little reason to retain deleted data for significant periods of time. The major reason that Genesys Info Mart provides `*_ANNEX` data is to support GI2 visibility controls, and GI2 uses only active `*_ANNEX` data for this purpose.

For a list of the *_ANNEX tables that are purged by Job_MaintainGIM, see Appendix D on page 445.

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

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

Specifies the number of days to retain data in the discard tables and audit and history tables.

The discard tables are Staging tables that store operational data that the transformation job was unable to process—for example, voice interaction data with unresolved IS-Links, or Configuration details records with missing configuration objects. The audit and history tables are Control tables that store information about data lineage and about ETL processing activity. Information in the discard, audit, and history tables is useful for troubleshooting.

Records in the discard, audit, and history tables are purged based on the timestamp of the ETL processing event—ETL_TS for the discard (Staging) tables and CREATED_TS for the audit and history (Control) tables.

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

For a list of the discard, audit, and history tables that are purged by Job_MaintainGIM, see Appendix D on page 445.

days-to-keep-gidb-facts

Default Value: 14
Valid Values: Any positive integer
Dependencies: None in releases 8.1.0 and 8.1.1; starting with release 8.1.2, days-to-keep-gim-facts
Changes Take Effect: At the next run of Job_MaintainGIM

Note: The following description applies to purge behavior starting with release 8.1.1. For the benefit of Genesys Info Mart 8.1.0 customers, Appendix E on page 451 provides information about purge behavior in previous 8.x releases.

Specifies the number of days to retain fact data in GIDB. Facts that have a start time that is earlier than the retention period are eligible to be purged.

For multimedia interactions, this option specifies the retention period for completed facts, provided that there are no active facts that have the same (or earlier) start timestamps. If there are active facts with the same (or earlier) start

timestamps, the eligible completed facts will not actually be purged until these active interactions have been terminated and, therefore, also become eligible to be purged, as described in `days-to-keep-active-facts` (see page 281).

For a list of the GIDB tables that the maintenance job purges, see Appendix D on page 445.

For information about the purging rules and guidelines to consider when you are setting the value of this option, see the sections about purging rules and purging strategies, starting on page 408.

days-to-keep-gim-facts

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

Note: The following description applies to purge behavior starting with release 8.1.1. For the benefit of Genesys Info Mart 8.1.0 customers, Appendix E on page 451 provides information about purge behavior in previous 8.x releases.

Specifies the number of days to retain fact data in the dimensional model. Facts that have a start time that is earlier than the retention period are eligible to be purged. `Job_MaintainGIM` does not purge active fact data, dimension data (except for the `*_ANNEX` dimensions—see page 405), or aggregate tables (if aggregation is enabled).

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

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

also become eligible to be purged, as described in `days-to-keep-active-facts` (see page 281).

Note: By definition, all voice interaction data in the dimensional model relates to completed interactions. In release 8.1.1, as Table 22 on page 286 shows, active multimedia interactions can delay the purging of completed interactions if the period of time for which multimedia interactions can remain active in your deployment (`days-to-keep-active-facts`) is greater than `days-to-keep-gim-facts`.

Starting with release 8.1.2, Genesys Info Mart no longer permits this configuration. In 8.1.1 deployments, Genesys strongly recommends against this configuration, particularly in deployments that include RAA.

For a list of the tables in the dimensional model that the maintenance job purges, see Appendix D on page 445.

For information about the purging rules and guidelines to consider when you are setting the value of this option, see the sections about purging rules and purging strategies, starting on page 408.

How Purge Options Work

To illustrate how various `days-to-keep-*` options combine to determine when data will be purged, Tables 21 and 22 on page 285 provide examples of different scenarios for GIDB and the dimensional model, respectively.

Table 21: GIDB Purge Example

Scenario		Start Time of Facts to Be Purged
• `days-to-keep-gidb-facts=14`. • `days-to-keep-active-facts=30`. • `Job_MaintainGIM` is running on June 4, 2012, after the last ETL cycle for the day (in other words, `days-to-keep-gidb-facts` threshold is May 21, 2012; `days-to-keep-active-facts` threshold is May 5, 2012).		
1	A previously active e-mail interaction, which started on May 5, 2012, was artificially terminated on June 4, 2012. Other e-mail interactions, which started on May 6, 2012, and later, remain active.	Voice interactions, agent activity, Outbound Contact facts: May 21, 2012, or earlier Multimedia interactions: May 5, 2012, or earlier
2	The earliest active fact relates to an e-mail interaction that started on May 27, 2012.	All facts: May 21, 2012, or earlier

Table 22: Dimensional Model Purge Example

Scenario		Start Time of Facts to Be Purged
• `days-to-keep-gim-facts=400`. • `days-to-keep-active-facts=30`. • `Job_MaintainGIM` is running on June 4, 2012, after the last ETL cycle for the day (in other words, `days-to-keep-gim-facts` threshold is May 1, 2011; `days-to-keep-active-facts` threshold is May 5, 2012).		
1	The oldest previously active e-mail interaction, which started on May 5, 2012, was artificially terminated on June 4, 2012.	All facts: May 1, 2011, or earlier
• `days-to-keep-gim-facts=400`. • `days-to-keep-active-facts=600`. • `Job_MaintainGIM` is running on June 4, 2012, after the last ETL cycle for the day (in other words, `days-to-keep-gim-facts` threshold is May 1, 2011; `days-to-keep-active-facts` threshold is October 12, 2010). **Note:** This configuration is not supported starting with release 8.1.2 and is strongly discouraged in earlier releases.		
1	The oldest previously active e-mail interaction, which started on October 12, 2010, was artificially terminated on June 4, 2012.	Agent activity, Outbound Contact facts: May 1, 2011, or earlier Voice and multimedia interactions: October 12, 2010, or earlier
2	The earliest active fact relates to an e-mail interaction that started on June 20, 2011.	All facts: May 1, 2011, or earlier

delayed-data-threshold

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

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

The context in which delayed data might be detected was extended in release 8.1.3. Starting with release 8.1.3, Genesys Info Mart applies the `delayed-data-threshold` to extraction delays, when there is no data available to be extracted from an active data source. In previous releases, Genesys Info Mart applied the `delayed-data-threshold` to transformation delays, when

Genesys Info Mart was unable to process a data chunk because dependent data was not available (in other words, had not been extracted yet).

Genesys recommends that you set an alarm on log message 55-20110, so that you can investigate the reasons for data delays in a timely manner and take appropriate action. For more information about appropriate actions, see "Delayed Data for Extraction" on page 376, as well as the section about recovering from data-source unavailability in the Troubleshooting chapter in the *Genesys Info Mart 8.1 Operations Guide*.

etl-start-date

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

Specifies the earliest date for which Genesys Info Mart considers IDB data for extraction in a new deployment or when the Info Mart database is re-initialized. IDB data that has timestamps earlier than the ETL start date is never extracted.

The option, which was introduced in release 8.1.103.07, is used only when `Job_InitializeGIM` initializes the database. If the `etl-start-date` option is not specified, the earliest starting point for Genesys Info Mart processing is IDB data that has timestamps 30 days prior to the Info Mart database initialization.

The main purpose of the option is to pre-empt performance and maintenance issues when Genesys Info Mart is introduced into a deployment with much older existing IDB data. Specifying the ETL start date enables users to:

- Shorten the backlog of IDB data if they do not need to include all of the old data in their reporting.
- In deployments with a partitioned Info Mart database, specify the starting point for creating partitions, which is important for proper partitioning.

Starting with release 8.1.4, Genesys Info Mart ignores this option value for Configuration details; after the database has been initialized or re-initialized, Genesys Info Mart extracts all `cfg` data going back as far as 2010, regardless of the value of `etl-start-date`.

extract-data-cfg-facts-chunk-size

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

Specifies the size of the time interval, in seconds, for which configuration relationship data is committed in one transaction. The data from this period is considered to be one data chunk, for the purpose of extract and transform. For example, if you set the value of this option to `90000`, Genesys Info Mart extracts 25 hours (90,000 seconds) of available configuration relationship fact

data. When any nonpositive value is set (for example, `0` or `-1`), data for all available time intervals is extracted in one chunk.

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

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

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

extract-data-chunk-size

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

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

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

extract-data-max-conn

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

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

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

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

extract-data-stuck-threshold

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

Specifies the time, in seconds, that Genesys Info Mart waits for stuck data to become available in IDB. For example, if new data is available until 3:00 PM and the value for `extract-data-stuck-threshold` is set to `28860` (8 hrs 1 min.), Genesys Info Mart extracts data that has a timestamp from 6:59 AM to 3:00 PM. IDB data from the time intervals before 6:59 AM is not, and will never be, extracted. For a more detailed example, see "Extraction Window Example" on page 374.

The value of this option sets the stuck data timeout for the extraction windows for the Voice, Multimedia, and Outbound Contact data domains. However, the actual time ranges of the extraction windows and HWMs for each data domain might be different (see "Separate Extraction Windows for Each Data Domain" on page 375).

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

extract-data-thread-pool-size

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

Specifies the maximum number of worker threads that are used to extract data concurrently. This option does not set a strong limit on the total number of threads that will be used in extraction processing, because certain extraction algorithms create additional helper threads. Instead, this option specifies the maximum number of logical partitions for concurrent extraction of subsets of data. Increasing the value of this option reduces the amount of time that is required for the extraction job, but it increases CPU and memory requirements,

especially if database links are not used. (At a rough estimate, each additional worker thread requires an additional 180 MB of RAM.) The optimal value of this option depends on the operating system, hardware (such as RAM and the number of CPUs), and the number of IDBs in the environment.

extract-last-second

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

Warning! This option *must* be set to `false` in a production environment. The option was discontinued in release 8.1.4.

In releases earlier than 8.1.4, specifies whether Genesys Info Mart extracts data from the last second of the specified time interval into the extracted data chunk. In a lab environment, setting this option to `true` speeds up data validation. If you are extracting the last second of data in a lab environment, make sure that all data from the executed test scenario is stored into IDB before Genesys Info Mart runs the extraction job.

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

The option was discontinued in release 8.1.4 because of issues with missing data. To speed up data validation in a lab environment, you can achieve a similar result by setting the ICON `NoData` timestamp (specified by the ICON `dss-no-data-tout` option—see page 167) to a small value, thus minimizing the processing delays that Genesys Info Mart provides to account for data-source inactivity.

max-call-duration

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

Specifies the maximum duration, in seconds, of voice calls in the deployed environment. This option controls the following timeouts:

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

For more information about the role of the stuck-link threshold in merge processing, see "Delayed or Missing Information" on page 387.

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

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

- The limit for waiting for After Call Work (ACW) before transforming interaction data.
- The limit for looking up party activity for multimedia interactions before transforming data. (Although a multimedia interaction can be very long-lasting, Genesys Info Mart considers that the duration of activity for any one party is comparable to the duration of a voice call.)
- In an HA deployment in which geolocation is a consideration, the amount of time that Genesys Info Mart will wait after a good ICON provider session has been re-established on the local IDB following a data disruption, before the ETL switches back to extract Voice details from the local IDB.

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

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

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

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

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

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

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

Specifies the amount of time, in hours, after which Genesys Info Mart ends campaign group states if the transformation process has not extracted a stopped row that brackets a previously extracted started row from the IDB `GO_CampaignHistory` table.

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

max-chain-processing-duration-in-hours

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

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

The option affects Genesys Info Mart behavior only when the Info Mart database is partitioned. On partitioned databases, Genesys Info Mart will not create `CONTACT_ATTEMPT_FACT` (CAF) records for activity that occurs after the maximum expected duration of chain processing. If the value of this option is `0` (zero), Genesys Info Mart will use the value of `max-camp-group-session-duration-in-hours` (see page 291).

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

max-chunks-per-job

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

Specifies the number of the extracted data chunks that the transform job processes in one ETL cycle. As long as it seems practical from the performance perspective, increase the option value to transform a larger amount of data in a single cycle.

max-parties-per-call

Default Value: `100`
Valid Values: `50–10000`
Dependencies: None
Changes Take Effect: At the next run of `Job_TransformGIM`

Note: This option applies only to releases 8.1.103.03 through 8.1.3.x.

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

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

Genesys recommends that you set an alarm on the log messages, to prompt you to locate and fix inappropriate routing strategies that result in huge numbers of records in IDB tables.

The option was discontinued in release 8.1.4 because improvements in incremental transformation removed the need to enforce a limit on the number of records to be read in each chunk. Starting with release 8.1.4, while Genesys Info Mart continues to abbreviate the representation of unsuccessful routing attempts in the Info Mart database, log message numbers 55-20120, 55-20121, and 55-20122 are no longer generated.

Note: For all Genesys Info Mart releases, if you cannot change the behavior of inappropriate strategies, Genesys recommends that you utilize Interaction Server and Interaction Routing Designer (IRD) functionality to hide strategy activity, as described on page 66. Hiding strategy activity is the preferred method of suppressing event data that is not important for historical reporting, and will result in better data quality in Genesys Info Mart.

max-session-duration-in-hours

Default Value: 24
Valid Values: 1–10000
Dependencies: None
Changes Take Effect: On the next ETL cycle

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

For related information, see the description of max-state-duration. See also the discussion about long-duration sessions or states in the chapter about populating Genesys Info Mart data in the *Genesys Info Mart 8.1 User's Guide*.

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

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

max-state-duration

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

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

The max-state-duration option enables you to:

- Improve performance of agent-activity transformation by reducing the amount of active data that Genesys Info Mart must maintain and by limiting the period of time for which lookups must be performed for agent-activity data (for example, when the transformation job populates the RES_PREVIOUS_SM_STATE_FACT_KEY in the IRF).
- Recognize in a timely manner when an agent session has gone inactive. In releases earlier than 8.1.3, when agents forgot to log out, the reported durations of sessions and states might be misleadingly high. While the max-session-duration-in-hours option puts a limit on the overall session duration, max-state-duration provides the ability to detect when a session has gone inactive. If a state times out and there are no other active states, Genesys Info Mart can conclude that there really is no agent activity within the session.

The default value of 14400 seconds (4 hours) for max-state-duration strikes a balance between considering long periods of idleness (no agent state

transitions) to be normal behavior and realizing the performance benefits mentioned above.

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

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

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

For related information, see the discussion about long-duration sessions or states in the chapter about populating Genesys Info Mart data in the *Genesys Info Mart 8.1 User's Guide*.

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

The option was introduced in release 8.1.3.

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

Default Value: `31` (`30` in release 8.1.1)
Valid Values: Any positive integer greater than `days-to-keep-active-facts`
Dependencies: `days-to-keep-active-facts`
Changes Take Effect: On the next ETL cycle

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

This option, which was introduced in release 8.1.1, affects how Genesys Info Mart identifies which interactions are associated with the same threads and, therefore, affects the thread-related data that is reported in IF and IRF records. For performance reasons, Genesys recommends that you set this option to the smallest value that is consistent with interaction patterns in your deployment, so that Genesys Info Mart does not have to maintain memory of inactive threads for unnecessarily long periods of time.

For more information about how interaction threads are identified, see the subsection about interaction threads, in the section about populating interaction data, in the *Genesys Info Mart 8.1 User's Guide*.

max-time-deviation

Default Value: `30`
Valid Values: `1-120`
Dependencies: None
Changes Take Effect: On the next ETL cycle

Specifies the maximum time deviation, in seconds, to take into account for time synchronization inaccuracy between host system clocks.

Genesys recommends the following relationship between the value of this option and the Advanced Disconnect Detection Protocol (ADDP) timeout, which is the `Local Timeout` parameter configured for ADDP connections to data sources on the `Connections` tab of the ICON `Application`:

[(*ADDP Local Timeout*) * 2] + (*actual maximum difference in time synchronization between hosts*) <= `max-time-deviation`

In HA deployments, `max-time-deviation` is used for reliable analysis of the "no data" information in IDB; (`NoData` – `max-time-deviation`) is considered to be reliable for all data sources for a particular ICON provider. For more information about the ICON `NoData` indicator, see the section about determining IDB availability in the *Interaction Concentrator 8.x User's Guide.*

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

memory-threshold

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

Specifies the percentage of available memory that must be exceeded before Genesys Info Mart logs a message (55-20101) that indicates that the memory threshold has been exceeded. If the value of this option is set to `0`, the feature will be disabled.

merge-chunk-size

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

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

option depends on the characteristics of your deployment. Consider the following guidelines:

- The option value should be large enough that it does not require transformation of an unnecessarily large number of data chunks and, therefore, does not interfere with normal processing of one chunk of extracted data. As a rule of thumb, allow 200000 rows for a call rate of 100 calls per second (cps).
- The option value should be small enough that it protects the transformation job from running out of memory. For example, for an environment with 8 GB of RAM dedicated to Genesys Info Mart, allow a number of rows that corresponds to 1 million root `G_IR` records. (Every root `G_IR` record corresponds to one interaction fact.)

merge-failed-is-link-timeout

Default Value: `0`
Valid Values: Any nonnegative integer
Dependencies: None
Changes Take Effect: On the next ETL cycle

Specifies the time interval, in seconds, for which the merge of failed Inter Server Call Control (ISCC) links (IS-Links) will be delayed to enable Genesys Info Mart to receive both sides of the links.

If the value of this option is greater than the value of `max-call-duration`, which specifies the timeout for stuck links (see page 290), then `max-call-duration` controls the timeout for the merge of failed IS-Links as well.

This option was introduced in release 8.1.1. The default value (`0`), which preserves legacy behavior, means that the merge operation will process unpaired failed links immediately, without waiting for the other side of the failed IS-Link. As a result of the partial merge, the transformation job might encounter dangling links.

partitioning-ahead-range

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

Specifies, in terms of number of days, how far ahead `Job_InitializeGIM` (in the first instance) and `Job_MaintainGIM` (on an ongoing basis) will create partitions for GIDB, Control, and Info Mart fact tables that are partitioned. Starting with release 8.1.2, these jobs add partitions for the number of days ahead of the time that the job is running. (`Job_InitializeGIM` also adds partitions from `etl-start-date` up to the time that the job is running.) In earlier releases, `Job_MaintainGIM` adds partitions for the number of days ahead of the extraction high-water mark (`extractHWM`). The number of partitions that `Job_MaintainGIM`

actually creates during each run depends on the partition sizes and the job frequency.

For example, if `partitioning-interval-size-gidb=86400` (1 day), `partitioning-interval-size-gim=604800` (1 week), and `partitioning-ahead-days=14`, `Job_MaintainGIM` will create as many additional partitions as necessary to provide partitions for GIDB, Control (in release 8.1.2 and later), and Info Mart fact tables up to 14 days ahead. If `Job_MaintainGIM` runs daily, this means that:

- For GIDB tables, each run of `Job_MaintainGIM` will create one new partition of size 1 day, for the fourteenth day. (Previous runs will have created partitions for the other days.)
- For Info Mart fact tables, as well as for Control tables, the maintenance job will create one new partition of size 7 days at the start of each week. (A previous run will have created a partition for the other week.)

  If the value of `partitioning-ahead-range` is not a multiple of `partitioning-interval-size-gim`, the maintenance job will create a new partition only when the last day of the partitioning ahead range falls in a week for which a partition has not yet been created. For example, if `partitioning-interval-size-gim=604800` but `partitioning-ahead-days=10`, two new partitions will be created on the very first run of `Job_MaintainGIM`, and the next partition will be created on the fifth run.

To guarantee that partitions are always available for use by the ETL, ensure that `run-maintain` (see page 318) is set to `true` (the default value).

This option applies only in deployments that use partitioning.

partitioning-interval-size-gidb

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

Specifies the size of partitions, in seconds, for GIDB tables that are partitioned. `Job_MaintainGIM` creates partitions of the specified size in the Info Mart database in preparation for future ETL cycles. The default size of GIDB table partitions is 24 hours (86400 seconds).

This option applies only in deployments that use partitioning.

partitioning-interval-size-gim

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

Specifies the size of partitions, in seconds, for Info Mart fact and Control tables that are partitioned. `Job_MaintainGIM` creates partitions of the specified

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

In PostgreSQL deployments, if you plan to retain dimensional-model data for more than a year (`days-to-keep-gim-facts` is greater than `365`), Genesys recommends that you increase the value of `partitioning-interval-size-gim` to `604800` or greater, so that the size of Info Mart fact and Control table partitions is at least 7 days.

This option applies only in deployments that use partitioning.

purge-thread-pool-size

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

Specifies the maximum number of concurrent purging transactions. The optimum value for this option depends on the characteristics and capacity of your deployment. Consider increasing the value of this option if you think that there is scope to improve performance of the purge operation.

purge-transaction-size

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

Specifies the number of deleted records per table that will be committed in a single transaction. For example:

- If there are 150,000 records in a particular table that are eligible for purging and `purge-transaction-size=100000, Job_MaintainGIM` will delete and commit 100,000 records in one transaction and 50,000 records in a separate transaction.
- If there are 90,000 records in one table, 10,000 records in another table, and `purge-transaction-size=100000, Job_MaintainGIM` will delete and commit 90,000 records from the first table in one transaction and 10,000 records from the second table in a separate transaction.

q-answer-threshold-voice

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

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

queues or ACD queues. Genesys Info Mart uses this value unless you configure an option that has the same name on an individual Virtual Queue or ACD Queue DN object in Configuration Manager.

To set an answer threshold on a specific Virtual Queue or ACD Queue DN object, see Procedure: Configuring a DN for ICON and Genesys Info Mart reporting, on page 261.

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

For similar options that control equivalent thresholds for multimedia interactions, see the "gim-etl-media-<media type> Section" on page 302.

q-short-abandoned-threshold-voice

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

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

You cannot set this value on an individual virtual queue or ACD Queue DN object.

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

For similar options that control equivalent thresholds for multimedia interactions, see the "gim-etl-media-<media type> Section" on page 302.

short-abandoned-threshold

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

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

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

For similar options that control equivalent thresholds for multimedia interactions, see the "gim-etl-media-<media type> Section" on page 302.

sm-resource-state-priority

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

Specifies a list of the state names—`BUSY, ACW, NOT_READY,` or `READY`—in order of decreasing priority. When an agent simultaneously has different states on different DNs for a given media type, Genesys Info Mart uses this list to determine which state has the highest priority when determining a summarized state to store in the `SM_RES_STATE_FACT` or `SM_RES_STATE_REASON_FACT` table.

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

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

Note: The list does not include the `LOGGED_IN` state, which always has the lowest priority.

user-event-data-timeout

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

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

The option value also has a role in enabling complete ACW information in IRF records in scenarios in which transformation of agent state and session data occurred in an earlier ETL cycle than transformation of interaction data (for example, when the first agent who handled a call transferred it almost immediately to another agent and completed ACW long before the call ended). Therefore, consider your call topologies carefully before reducing the value of this option.

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

gim-etl-media-<media type> Section

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

By default, the Genesys Info Mart `Application Template` includes the `[gim-etl-media-email]` and `[gim-etl-media-chat]` sections and options, for specifying thresholds for Genesys eServices/Multimedia e-mail and chat interactions.

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

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

q-answer-threshold

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

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

- On a `Media Type Business Attribute` for a particular tenant, for interactions of that media type for that tenant. (For more information, see Procedure: Setting Media Type Business Attribute object options for Genesys Info Mart reporting, on page 260.)
- On an individual Virtual Queue `DN` object, for multimedia interactions that come through that virtual queue. (For more information, see Procedure: Configuring a DN for ICON and Genesys Info Mart reporting, on page 261.)
- On an individual `Script` object that corresponds to a particular Multimedia `Interaction Queue` or Multimedia `Interaction Workbin`, for multimedia interactions that come through an interaction queue or workbin. (For more information, see Procedure: Setting Script object options for Genesys Info Mart reporting, on page 265.)

q-short-abandoned-threshold

Default Value: `10`

Valid Values: 1–1000
Dependencies: None
Changes Take Effect: On the next ETL cycle

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

Genesys Info Mart uses the option setting in the Genesys Info Mart Application object unless you override this value by configuring an option that has the same name in Configuration Manager:

- On a Media Type Business Attribute for a particular tenant, for interactions of that media type for that tenant. (For more information, see Procedure: Setting Media Type Business Attribute object options for Genesys Info Mart reporting, on page 260.)

You cannot set this value on an individual Virtual Queue, Interaction Queue, or Workbin Script object.

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

short-abandoned-threshold

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

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

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

Genesys Info Mart uses the option setting in the Genesys Info Mart Application object unless you override this value by configuring an option that has the same name in Configuration Manager:

- On a Media Type Business Attribute for a particular tenant, for interactions of that media type for that tenant. (For more information, see Procedure: Setting Media Type Business Attribute object options for Genesys Info Mart reporting, on page 260.)

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

gim-etl-populate Section

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

populate-mm-ixnqueue-facts

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

Enables or disables the population of eServices/Multimedia Interaction Queue activity to the MSF table. Genesys Info Mart uses this value for all configured eServices/Multimedia Interaction Queues, unless you configure an option that has the same name on an individual `Script` object for a specific Interaction Queue in Configuration Manager (see Procedure: Setting Script object options for Genesys Info Mart reporting, on page 265).

- `false`—The placement of an interaction in an Interaction Queue will not be represented in the MSF table.
- `true`—The placement of an interaction in an Interaction Queue will be represented in the MSF table. Strategy time is not included in the mediation duration, except for scenarios in which an interaction is bounced between an Interaction Queue and a strategy, as the strategy retries busy agents repeatedly. For more information, see the description of the `MEDIATION_DURATION` field in the section about populating MSFs in the *Genesys Info Mart 8.1 User's Guide*.

Note: Ideally, this option should be set to `true` only for Interaction Queues for which mediation reporting is desired, so that the MSF table does not become cluttered with unnecessary information. Some Interaction Queues are not useful for reporting, such as the Interaction Queues that are associated with Interaction Workbin objects only through a configured Interaction Queue View.

populate-mm-workbin-facts

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

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

Genesys Info Mart considers workbin time as hold, see the description of the `populate-workbin-as-hold` option on page 307.

Genesys Info Mart uses this value for all configured eServices/Multimedia Interaction Workbins, unless you configure an option that has the same name on an individual `Script` object for a specific Interaction Workbin in Configuration Manager (see Procedure: Setting Script object options for Genesys Info Mart reporting, on page 265).

- `false`—The placement of an interaction in an Interaction Workbin will not be represented in the MSF table.
- `true`—Provided that Genesys Info Mart does not consider the workbin time as hold, the placement of an interaction in an Interaction Workbin will be represented in the MSF table. The MSF record references a `WORKBIN` dimension that identifies the type of resource that is associated with the workbin and the specific resource that is associated with the mediation. Strategy time is not included in the mediation duration, except for scenarios in which an interaction is bounced between an Interaction Workbin and a strategy, as the strategy retries busy agents repeatedly. For more information, see the description of the `MEDIATION_DURATION` field in the section about populating MSFs in the *Genesys Info Mart 8.1 User's Guide*.

populate-sm-chat-resource-activity

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

Enables or disables the population of Genesys eServices/Multimedia chat resource activity in the `SM_RES_SESSION_FACT`, `SM_RES_STATE_FACT`, and `SM_RES_STATE_REASON_FACT` tables. When the value of this option is `false`, Genesys Info Mart does not populate these tables with activity data for a Genesys eServices/Multimedia chat resource. When the value of this option is `true`, these tables are populated with the activity data for a Genesys eServices/Multimedia chat resource.

populate-sm-email-resource-activity

Default Value: `true`
Valid Values: `true, false`
Dependencies: None for release 8.1.1 and later; in earlier releases, has an effect only on the tables that are populated when the values of the following options are set to `true`: `populate-sm-resource-session-facts`,

populate-sm-resource-state-facts, and
populate-sm-resource-state-reason-facts.
Changes Take Effect: On the next ETL cycle

Enables or disables the population of Genesys eServices/Multimedia e-mail resource activity in the SM_RES_SESSION_FACT, SM_RES_STATE_FACT, and SM_RES_STATE_REASON_FACT tables. When the value of this option is false, Genesys Info Mart does not populate these tables with activity data for a Genesys eServices/Multimedia e-mail resource. When the value of this option is true, these tables are populated with the activity data for a Genesys eServices/Multimedia e-mail resource.

populate-sm-resource-session-facts

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

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

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

populate-sm-resource-state-facts

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

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

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

populate-sm-resource-state-reason-facts

Default Value: true
Valid Values: true, false

Dependencies: `populate-sm-resource-session-facts`, `populate-sm-resource-state-facts`
Changes Take Effect: On the next ETL cycle

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

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

populate-sm-voice-resource-activity

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

Enables or disables the population of voice resource activity in the `SM_RES_SESSION_FACT, SM_RES_STATE_FACT,` and `SM_RES_STATE_REASON_FACT` tables. When the value of this option is `false,` Genesys Info Mart does not populate these tables with activity data for a voice resource. When the value of this option is `true,` these tables are populated with the activity data for a voice resource.

populate-workbin-as-hold

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

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

- `false`—Workbin time is considered to be mediation. Whether or not the workbin activity will be represented by an MSF record depends on the value of the `populate-mm-workbin-facts` option (see page 304).
- `true`—Workbin time is considered to be hold if the handling resource places the interaction into its own personal workbin. The hold ends when any resource takes the interaction out of the workbin. Various hold metrics

are reported in the IRF record that is associated with that handling resource. (For the meaning of the term *personal workbin,* as used in Genesys Info Mart, see "Workbin Instance and Personal Workbin" on page 40.)

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

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

Table 23 on page 308 summarizes the results of different permutations of the `populate-workbin-as-hold` and `populate-mm-workbin-facts` options, in combination with the specifics of the interaction flow. If the values of the `populate-workbin-as-hold` and `populate-mm-workbin-facts` options are both set to `false`, the workbin activity is not represented in the dimensional model.

Table 23: Workbin Reporting Results Matrix

<table>
<tr><th rowspan="2">Scenario</th><th colspan="3">Reporting Result for Workbin Activity</th></tr>
<tr><th>populate-workbin-as-hold=false
populate-mm-workbin-facts=true</th><th>populate-workbin-as-hold=true
populate-mm-workbin-facts=true</th><th>populate-workbin-as-hold=true
populate-mm-workbin-facts=false</th></tr>
<tr><td>An Agent/Place resource places the interaction into its personal workbin.</td><td rowspan="3">MSF record.*
This is the default behavior.</td><td colspan="2">Hold metrics in the IRF record for the association of that Agent/Place with the interaction.</td></tr>
<tr><td>Another resource places the interaction into the personal workbin of an Agent/Place.</td><td rowspan="2">MSF record.*</td><td rowspan="2">Not represented in the dimensional model.</td></tr>
<tr><td>Any resource places the interaction into an AgentGroup or PlaceGroup workbin.</td></tr>
<tr><td colspan="4">*Starting with release 8.1.1, the MSF record includes a WORKBIN_KEY value that identifies the workbin instance that is associated with the mediation.</td></tr>
</table>

gim-transformation Section

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

adjust-vq-time-by-strategy-time

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

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

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

  In multimedia scenarios in which an interaction is bounced between a mediation resource (for example, an Interaction Queue or a Workbin) and a strategy, as the strategy retries busy agents repeatedly, this option comes into effect only for the virtual queue that is associated with the last strategy party, before the interaction is routed successfully or else terminated. Genesys Info Mart does not report on virtual-queue activity that overlaps the repeated interim mediations.

canceled-queues

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

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

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

The option was introduced in release 8.1.3. In earlier 8.x releases, Genesys Info Mart handled these situations as transfers to a queue. For related information, see "Archive Queues" on page 100.

completed-queues

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

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

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

The option was introduced in release 8.1.3. In earlier 8.x releases, Genesys Info Mart handled these situations as transfers to a queue. For related information, see "Archive Queues" on page 100.

default-ivr-to-self-service

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

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

- `false`—The IVR is treated as a nonself-service IVR (in other words, as a mediation device).
- `true`—The IVR is treated as a self-service IVR (in other words, as a handling resource).

ignored-reason-codes

Default Value: `INTERACTION_WORKSPACE`
Valid Values: A comma-separated list of reason codes
Dependencies: None

Changes Take Effect: At the next run of Job_TransformGIM

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

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

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

The option was introduced in release 8.1.4.

irf-io-parallelism

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

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

msf-target-route-thru-queue

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

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

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

This option affects the technical results and targets that are reported in the MSF records for virtual queues and ACD queues, as well as the mediation segments and resources that are referenced in the IRF record. For more

information, see the section about populating MSFs in the *Genesys Info Mart 8.1 User's Guide*.

pipeline-timeout-in-hours

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

Specifies the maximum expected duration, in hours, of the execution of a single transformation pipeline. (*Single transformation pipeline* refers to the separate pipelines for in-memory transformation of separate data streams, such as for voice, multimedia, Outbound Contact, and agent data.) If the timeout is exceeded, Genesys Info Mart tries to abort the pipeline, and the transformation job fails.

The option was introduced in release 8.1.3, with the default value set to preserve legacy behavior. For performance reasons, Genesys recommends that you retain the default value unless Genesys Customer Care advises you to change it.

ud-io-parallelism

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

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

log Section

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

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

standard

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

Specifies the location of the log output.

verbose

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

Specifies the minimum level of logging.

log4j Section

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

console-pattern-layout

Default Value: `%d{ISO8601} %-5p %-12t %m%n`
Dependencies: None
Changes Take Effect: Immediately

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

file-pattern-layout

Default Value: `%d{ISO8601} %-5p %-12t %m%n`
Dependencies: None
Changes Take Effect: Immediately

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

log4j.appender.ConsoleLogger.Threshold

Default Value: `info`
Dependencies: None
Changes Take Effect: Immediately

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

log-file-name

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

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

logging-level

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

Determines whether local logging is enabled, and specifies the minimum level of events to log. The debug, info, and warn values correspond to the Genesys Management Layer DEBUG, TRACE, and STANDARD logging values, respectively.

In releases earlier than 8.1.3, the valid values were debug, info, warn, and none.

max-backup-index

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

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

max-log-file-size

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

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

schedule Section

This configuration section specifies the schedule that Genesys Info Mart Server uses to launch the ETL jobs. The Genesys Info Mart Server enables options to be modified while it is running. For those options that specify a time, the time format is HH:mm, where HH represent the number of hours (00–24), and mm represents the number of minutes (00–59).

aggregate-duration

Default Value: 05:00
Valid Values: 00:00–24:00
Dependencies: run-aggregates, aggregate-schedule
Changes Take Effect: Immediately

Specifies the amount of time, in 24-hour format, that Job_AggregateGIM will run after it is launched. When the run-aggregates option (see page 317) is set

to true, the scheduler will stop the aggregation job when this interval expires. The aggregation job is launched in accordance with a schedule defined by the aggregate-schedule option. After the aggregation job is launched, it runs continuously until the aggregation-duration interval expires.

aggregate-schedule

Default Value: 0 1
Valid Values: A valid CRON expression
Dependencies: run-aggregates
Changes Take Effect: Immediately

Specifies the daily schedule for Job_AggregateGIM to start. The job will start in accordance with this schedule when aggregation is being controlled by the scheduler (in other words, the run-aggregates option is set to true). Between them, the aggregate-schedule and aggregate-duration options define daily time intervals within which Job_AggregateGIM will run continuously.

The schedule is defined in the format of a CRON expression that represents a set. The expression comprises two fields, which are separated by whitespace:

- The first field specifies minutes. Valid values are 0–59 and optional special characters (see below).
- The second field specifies hours. Valid values are 0–23 and allowed special characters.

Special Characters

The following special characters are allowed in the CRON expression:

- , (comma)—Separates items in a list. For example, specifying the first field (minutes) as 0,30,45 means the 0th, 30th, and 45th minutes of the hour.
- - (hyphen)—Defines a range. For example, specifying the first field (minutes) as 30-35 means every minute between the 30th and 35th minute of the hour, inclusive; this is the same as specifying 30,31,32,33,34,35.
- * (asterisk)—Indicates that the CRON expression will match for all values of the field. For example, specifying the second field (hours) as * means every hour in the day.
- / (forward slash)—Describes increments. For example, specifying the first field (minutes) as 0/10 means the 0th minute of the hour and every 10 minutes thereafter.

Examples

The following values for aggregate-schedule illustrate sample schedules:

- 0 1 means that the aggregation job will be launched once a day at 01:00.
- 30 0,3/2 means that the aggregation job will be launched every day at 00:30, 03:30, and every 2 hours after that for the rest of the day.

  This schedule assumes that the value of aggregate-duration is 02:00 or less. The scheduler will not launch a new instance of Job_AggregateGIM while an existing instance is running. For aggregation to run on the specified schedule, the value of aggregate-duration must not exceed the intervals between scheduled start times.

- `30 *` means that the aggregation job will be launched every hour during the day on the half-hour (00:30, 01:30, 02:30, and so on), assuming that the value of `aggregate-duration` is `01:00` or less.

Genesys recommends against configuring a schedule that has the aggregation job running in a series of short bursts—for example, `aggregate-schedule=30 *` and `aggregate-duration=00:15`. When the time specified by `aggregate-duration` expires, the scheduler immediately stops the aggregation job, even if it is in the middle of processing a batch of data.

If you want `Job_AggregateGIM` to run continuously for 24 hours a day, without any breaks for maintenance activities (which is not recommended), set `aggregate-schedule=0 0` and `aggregate-duration=24:00`.

etl-end-time

Default Value: `22:00`
Valid Values: `00:00`–`23:59`
Dependencies: `run-scheduler, etl-start-time`
Changes Take Effect: Immediately

Specifies the time of day, in 24-hour format, when the last ETL cycle can start running. If the value that you specify is before the ETL start time, the end time is for the next day (past midnight).

Ensure that you configure `etl-start-time` and `etl-end-time` so that there is sufficient time for the last ETL cycle of the day to complete and for `Job_MaintainGIM` to run (see the `maintain-start-time` option on page 317), before the start of the first ETL cycle of the next day. In particular, do not configure `etl-start-time=etl-end-time`; if you do, no ETL jobs will run.

etl-frequency

Default Value: `1`
Valid Values: `0`–`1440`
Dependencies: None
Changes Take Effect: On the next ETL cycle

Specifies the number of minutes that pass between the start times of each ETL cycle. If the amount of time that it takes to complete a cycle is shorter than the specified value, the next cycle is delayed until the time elapses. If the amount of time that it takes to complete a cycle is longer than the specified value, the next cycle is started immediately.

The ETL frequency must not be greater than the chunk size for data extraction, as specified by the `extract-data-chunk-size` option (see page 288). Otherwise, Genesys Info Mart will not be able to keep pace with ICON. When it checks the deployment (see "Configuration Check" on page 402), Genesys Info Mart verifies the internal consistency between the ETL frequency and extraction chunk size.

By default, the value of `etl-frequency` is much smaller than the value of `extract-data-chunk-size`. Genesys recommends that you retain this

relationship, to minimize data latency. For example, say that `extract-data-chunk-size=900` (15 minutes), `etl-frequency=1`, and all data from the last chunk has been processed; when the next ETL cycle starts 1 minute later, there is only 1 minute's worth of new data, and this can be processed very quickly. Alternatively, if there is a backlog of data, and it takes less than 15 minutes to process a 15-minute chunk, the next ETL cycle starts almost immediately, to continue catching up.

etl-start-time

Default Value: `06:00`
Valid Values: `00:00–23:59`
Dependencies: `run-scheduler`
Changes Take Effect: Immediately

Specifies the time of day, in 24-hour format, when the first ETL cycle starts running.

maintain-start-time

Default Value: `03:00`
Valid Values: `00:00–23:59`
Dependencies: `run-maintain`
Changes Take Effect: Immediately

Note: If the time of day that is represented by the new value has already passed, the new value is applied to the following day.

Specifies the time of day, in 24-hour format, when `Job_MaintainGIM` is started. This job is scheduled to start at this time when the `run-maintain` option is set to `true`. The value that you specify must be outside the range that is specified by `etl-start-time` and `etl-end-time`.

run-aggregates

Default Value: `false`
Valid Values: `true, false`
Dependencies: None
Changes Take Effect: Immediately

Specifies whether the scheduler will manage the aggregation job, to run the aggregation engine inside the Genesys Info Mart process.

When the value of this option is set to `true`, the scheduler will start `Job_AggregateGIM` at the scheduled time, as specified by the `aggregate-schedule` option (see page 315); `Job_AggregateGIM` will then run continuously until the scheduler stops the job after the scheduled interval, as specified by the `aggregate-duration` option (see page 314). The scheduler will not allow a second aggregation process to be launched while the job is running. The scheduler will also not allow any other aggregation process to be launched

outside the intervals that are defined by the aggregate-schedule and aggregate-duration options.

For example, if run-aggregates=true, aggregate-schedule=0 1, and aggregate-duration=05:00, the aggregation job will run continuously between 01:00 AM and 06:00 AM daily. The scheduler will not allow you to launch a second instance of Job_AggregateGIM manually from the management GUI (Genesys Info Mart Manager or the Genesys Info Mart Administration Console) within that time period. Furthermore, if you try to launch an instance of Job_AggregateGIM manually from the management GUI outside that time period (for example, at 08:00 AM), the scheduler will identify that the job is not supposed to be running at that time and will stop it. If you want to run Job_AggregateGIM manually from the management GUI outside the scheduled times, you must first set run-aggregates to false.

For more information about starting Job_AggregateGIM and managing the aggregation job from the management GUI, see the *Genesys Info Mart 8.1 Operations Guide*.

run-maintain

Default Value: true
Valid Values: true, false
Dependencies: None
Changes Take Effect: Immediately

Specifies whether to run Job_MaintainGIM at the scheduled time, as specified by the maintain-start-time option (see page 317).

run-scheduler

Default Value: false
Valid Values: true, false
Dependencies: None
Changes Take Effect: Immediately

Specifies whether to stop or start the scheduler. If the value of this option was set to true, so that the scheduler is currently scheduling jobs, and you change the value of this option to false, the scheduler pauses, with no effect on any jobs that might already be running. If you then reset the value to true, the scheduler resumes at the point at which it stopped.

run-update-stats

Default Value: false
Valid Values: true, false
Dependencies: None
Changes Take Effect: Immediately

Specifies whether Job_UpdateStats will run in PostgreSQL deployments.

When the value of this option is set to true, the scheduler will start the job at the time specified by the update-stats-schedule option (see page 319) and

will then periodically run the job at the intervals specified by the `update-stats-schedule` option.

timezone

Default Value: `GMT`
Valid Values: Any valid Java time zone
Dependencies: None
Changes Take Effect: Immediately

Specifies the time zone in which the schedule is defined. Internally, Genesys Info Mart maintains the schedule in UTC time. For convenience, you can use this option to specify a local time zone that makes it easier for you to plan and manage the schedule. You can use any valid time zone that is supported by the version of the JRE that runs the Genesys Info Mart Server.

For more information about supported time zones, see the documentation about calendar time zones on the Java developer website or other public resources. For sample reference sites, see the description of the `date-time-tz` option on page 273.

update-stats-schedule

Default Value: `0/10 *`
Valid Values: A valid CRON expression
Dependencies: `run-update-stats`
Changes Take Effect: Immediately

Defines the time intervals at which `Job_UpdateStats` will run. The job will start and then run periodically in accordance with this schedule. By default, the job runs every 10 minutes throughout the day.

`Job_UpdateStats` can run in conjunction with the ETL jobs, but not in conjunction with `Job_MaintainGIM`. However, the schedule that you configure for `Job_UpdateStats` does not need to specifically allow for a maintenance window: A running instance of `Job_UpdateStats` does not prevent `Job_MaintainGIM` from starting, and once `Job_MaintainGIM` has started as part of the schedule, the scheduler suspends the schedule for `Job_UpdateStats` until `Job_MaintainGIM` finishes.

The schedule is defined in the format of a CRON expression that represents a set. The expression comprises two fields, which are separated by whitespace:

- The first field specifies minutes. Valid values are `0`–`59` and optional special characters (see below).
- The second field specifies hours. Valid values are `0`–`23` and allowed special characters.

Special Characters

The following special characters are allowed in the CRON expression:

- , (comma)—Separates items in a list. For example, specifying the first field (minutes) as `0,30,45` means the 0th, 30th, and 45th minutes of the hour.

- - (hyphen)—Defines a range. For example, specifying the first field (minutes) as `30-35` means every minute between the 30th and 35th minute of the hour, inclusive; this is the same as specifying `30,31,32,33,34,35`.
- * (asterisk)—Indicates that the CRON expression will match for all values of the field. For example, specifying the second field (hours) as `*` means every hour in the day.
- / (forward slash)—Describes increments. For example, specifying the first field (minutes) as `0/10` means the 0th minute of the hour and every 10 minutes thereafter.

For values that illustrate sample schedules, see the examples for the `aggregate-schedule` option on page 315.

Chapter

15 Installing Genesys Info Mart Components

This chapter includes all the information that you need in order to prepare the Genesys Info Mart Server host for installation. It also describes how to install the Genesys Info Mart application and the user interfaces (*management GUIs*) for administering Genesys Info Mart—the Genesys Info Mart Manager or the Genesys Info Mart Administration Console.

This chapter contains the following sections:

Overview

Before You Proceed

Before you can install Genesys Info Mart components, make sure that you have:

- Configured the `Application` object for Genesys Info Mart Server.
- Set any Genesys Info Mart–related options in the `Tenant`, `Switch`, and `DN` configuration objects.

If you have not already done so, follow the instructions in Chapter 12 on page 231 before you continue.

Task Flow for Installing Genesys Info Mart Components

Task Summary: Installing Genesys Info Mart Components summarizes the task flow to install Genesys Info Mart Server and the Genesys Info Mart management GUI.

Task Summary: Installing Genesys Info Mart Components

Objective	Related Procedures and Actions
1. Prepare the Genesys Info Mart Server host.	Verify and, if necessary, install or modify system information about the following supporting software components: • Genesys Local Control Agent (LCA) • Java Development Kit (JDK) or Server Java Runtime Environment (Server JRE), including the `PATH` and `JAVA_HOME` environment variables • The Java Database Connectivity (JDBC) driver, including the `CLASSPATH` environment variable For more information, see "Preparing the Genesys Info Mart Server Host" on page 323.
2. (Optional) Enable Transport Layer Security (TLS) protocol on the connections from Genesys Info Mart Server to Configuration Server and Message Server.	See Task Summary: Enabling Secure Connections, on page 155.
3. Install the Genesys Info Mart Server application.	Install the Genesys Info Mart Server application, as follows: 1. Review and follow recommendations in "Preparing the Genesys Info Mart Server Host" on page 323. 2. Install Genesys Info Mart Server, using instructions that are appropriate to your environment: ◆ Procedure: Installing the Genesys Info Mart application (Windows), on page 327 ◆ Procedure: Installing the Genesys Info Mart application (UNIX), on page 328 3. If necessary, follow the recommendations in "Installing the Genesys Info Mart Administration Console" on page 332.
4. Verify the Genesys Info Mart Server host setup.	Verify that the operational parameters of the Genesys Info Mart Server host conform with Genesys Info Mart requirements. For more information, see "Verifying Host Requirements" on page 329.

Task Summary: Installing Genesys Info Mart Components (Continued)

Objective	Related Procedures and Actions
5. Install a graphical user interface (GUI) application (a *management GUI*) for monitoring and controlling Genesys Info Mart jobs.	Install and then open one or both of the following: • (Starting with release 8.1.4) Genesys Info Mart Manager, as instructed in: a. Procedure: Installing the Genesys Info Mart Manager plug-in, on page 331 b. Procedure: Accessing Genesys Info Mart Manager, on page 340 • The Genesys Info Mart Administration Console, as instructed in: a. Procedure: Installing the Genesys Info Mart Administration Console (Windows), on page 333 b. Procedure: Accessing the Genesys Info Mart Administration Console, on page 342
6. Ensure time synchronization in the deployment.	Verify that the system times are synchronized on all hosts on which Genesys applications are running. Otherwise, Genesys Info Mart might report inaccurate data.

Preparing the Genesys Info Mart Server Host

To prepare the Genesys Info Mart Server host for the Genesys Info Mart installation, perform the following tasks, which are described in detail in this chapter:

- If you want to use the Genesys Solution Control Interface (SCI) to control the operation of the Genesys Info Mart Server, install, or verify the installation of, Genesys Local Control Agent (LCA).
- Install the Java Development Kit (JDK) or Server Java Runtime Environment (JRE) and modify, or verify the content of, the `PATH` and the `JAVA_HOME` environment variables. For the Java versions that Genesys Info Mart supports, see the *Genesys Supported Operating Environment Reference Guide.*
- Install, or verify the installation of, a Java Database Connectivity (JDBC) driver for each RDBMS that Genesys Info Mart Server will access (Oracle and Microsoft SQL Server). Modify, or verify the content of, the `CLASSPATH` environment variable.

Installing LCA

If you plan to monitor or control Genesys Info Mart through the Management Layer, you must also configure and install Management Layer components—in particular, LCA.

To monitor the status of Genesys Info Mart components through the Management Layer, you must load an LCA instance on every host that is running Info Mart and DB Server instances. Without LCA, the Management Layer cannot monitor the status of these components.

If you do not use the Management Layer, you do not need LCA.

You will need a Genesys Management Framework product CD in order to install the components of the Management Layer. For more information about these Framework components, including deployment instructions, see the *Framework Deployment Guide* and the *Framework Management Layer User's Guide*.

Installing Java

Install the JDK or Server JRE on the Genesys Info Mart Server host, according to instructions in the Java documentation.

Windows If you plan to run Genesys Info Mart Server on Windows:

- If you will use Genesys Solution Control (SCI) to start and stop Genesys Info Mart Server, you must modify the `PATH` and the `JAVA_HOME` environment variables for either the system account or the user account under which you plan to start the LCA service.
- If you will run the Genesys Info Server as a Windows service, you must modify the `PATH` and the `JAVA_HOME` environment variables for either the system account or the user account under which you plan to start the Genesys Info Mart Server Windows service.

To modify the `PATH` environment variable, add `<java-install-dir>\bin` to the path (where `<java-install-dir>` is the path where you installed the JDK or Server JRE), so that it appears earlier in the path than any other reference to Java.

For the `JAVA_HOME` environment variable, specify `<java-install-dir>`.

UNIX If you plan to run Genesys Info Mart Server on UNIX:

- If you will use LCA to start and stop Genesys Info Mart Server, you must modify the `PATH` and the `JAVA_HOME` environment variables for the user account under which you plan to start the LCA service.
- If you will not use LCA, you must modify the `PATH` and the `JAVA_HOME` environment variables for the user account under which you plan to start Genesys Info Mart Server.

To modify the `PATH` environment variable, add `<java-install-dir>/bin` to the path (where `<java-install-dir>` is the path where you installed the JDK or

Server JRE), so that it appears earlier in the path than any other reference to Java.

For the JAVA_HOME environment variable, specify <java-install-dir>.

Installing JDBC Drivers

JDBC Driver for Oracle

You must use Oracle Thin Driver version 11.2.0.1.0 or later, regardless of the Oracle RDBMS version that you are using. The required Type 4 JDBC thin client driver ships with the Oracle client, or you can download it from the Oracle website. The class name of the driver that Genesys Info Mart uses for Oracle is oracle.jdbc.driver.OracleDriver. This class is from the ojdbc6.jar file.

You must also modify your CLASSPATH environment variable, so that Genesys Info Mart can locate the JDBC driver. The specific CLASSPATH environment variable that you modify depends on the operating system and user account under which the Genesys Info Mart Server runs.

Windows If you plan to run Genesys Info Mart Server on Windows:

- If you will use Genesys Solution Control to start and stop Genesys Info Mart Server, you must modify the CLASSPATH environment variable for either the system account or the user account under which you plan to start the LCA service.
- If you will run Genesys Info Mart Server as a Windows service, you must modify the CLASSPATH environment variable for either the system account or the user account under which you plan to start the Genesys Info Mart Server Windows service.

To modify the CLASSPATH environment variable, add <Oracle-jdbc-driver-dir>\ojdbc6.jar (where <Oracle-jdbc-driver-dir> is the path where you installed the Oracle JDBC driver), so that it appears earlier in the class path than any other reference to an Oracle JDBC driver.

UNIX If you plan to run Genesys Info Mart Server on UNIX:

- If you will use Genesys Solution Control to start and stop Genesys Info Mart Server, you must modify the CLASSPATH environment variable for the user account under which you plan to start LCA.
- If you will not use Genesys Solution Control, you must modify the CLASSPATH environment variable for the user account under which you plan to start Genesys Info Mart Server.

To modify the CLASSPATH environment variable, add <Oracle-jdbc-driver-dir>/ojdbc6.jar (where <Oracle-jdbc-driver-dir> is the path where you installed the Oracle JDBC driver), so that it appears earlier in the class path than any other reference to an Oracle JDBC driver.

JDBC Driver for Microsoft SQL Server

The Microsoft SQL Server Client is not shipped with a JDBC driver. Therefore, you must download and install the driver separately. The required driver is available from the Microsoft Download Center website by searching for "JDBC Driver." Genesys Info Mart requires JDBC Driver 1.2 or later.

Install the driver by following the installation instructions that are supplied with it.

The name of the `.jar` file that contains the driver is `sqljdbc4.jar`. After the installation, you must update the `CLASSPATH` environment variable to include the path to the folder that contains this `.jar` file.

To modify the `CLASSPATH` environment variable, add `<MSSQL-jdbc-driver-dir>\sqljdbc4.jar` (where `<MSSQL-jdbc-driver-dir>` is the path where you installed the Microsoft SQL Server JDBC driver), so that it appears earlier in the class path than any other reference to Microsoft SQL Server JDBC driver.

JDBC Driver for PostgreSQL

You must download and install the JDBC driver. The required driver is available from the PostgreSQL website by searching for "JDBC Driver." Genesys Info Mart requires JDBC4 Postgresql Driver Version 9.2-1003 or later.

Install the driver by following the installation instructions that are supplied with it.

The name of the `.jar` file that contains the driver varies by version. For example, the minimum supported version is named `postgresql-9.2-1003.jdbc4.jar`. After the installation, you must update the `CLASSPATH` environment variable to include this `.jar` file.

To modify the `CLASSPATH` environment variable, add `<PostgreSQL-jdbc-driver-dir>\<postgresql.jar>` (where `<PostgreSQL-jdbc-driver-dir>` is the path where you installed the PostgreSQL JDBC driver and `<postgresql.jar>` is the name of the `.jar` file for the version), so that it appears earlier in the class path than any other reference to PostgreSQL JDBC driver.

Installing the Genesys Info Mart Application

You can install the Genesys Info Mart application on either a Windows or a UNIX operating system.

Windows Installation

Procedure: Installing the Genesys Info Mart application (Windows)

Purpose: To install the Genesys Info Mart 8.1 application on a host that is running a Windows operating system.

When you install Genesys Info Mart on a Windows operating system, Genesys Info Mart is also installed as a Windows service with a startup type of `Automatic`. This means that if the host machine is restarted, the Windows service will start Genesys Info Mart automatically.

Note: You can install more than one Genesys Info Mart application in Configuration Manager.

Prerequisites

- The host has been prepared, as described in "Preparing the Genesys Info Mart Server Host" on page 323.

Start of procedure

1. Insert the Genesys Info Mart CD into the CD-ROM drive of the machine on which you want to install Genesys Info Mart.
2. Navigate to, and open, the `genesys-info-mart\windows` directory.
3. Double-click the `setup.exe` file, and then follow the directions in the installation wizard.

Warning! If you are installing on a 64-bit operating system, Genesys recommends that you specify an installation folder under `/gcti/gim_etl/`. There must not be any parentheses in the path name. For example, if you specify an installation folder under `/Program Files (x86)/`, the installation will fail.

If you want to enable client-side port definition for the initial connection to Configuration Server, specify the required parameters on the `Client Side Port Configuration` page in the wizard. For more information, see the procedure about configuring a connection to Configuration Server on Windows in the chapter about client-side port definition in the *Genesys 8.x Security Deployment Guide*.

Note: The following error message might appear. If it does, you should ignore it:

```
Unable to find configuration information. Either you have not
used configuration wizards and the GCTISetup.ini file was not
created or the file is corrupted.
```

End of procedure

Next Steps

- Verify the Genesys Info Mart host setup. See "Verifying Host Requirements" on page 329.
- (Optional) Modify the default arguments for the Genesys Info Mart application. For more information, see "Modifying the Genesys Info Mart Default Arguments" on page 359.
- Install the management GUI. For more information, see Procedure: Installing the Genesys Info Mart Manager plug-in, on page 331 or Procedure: Installing the Genesys Info Mart Administration Console (Windows), on page 333.

UNIX Installation

Procedure: Installing the Genesys Info Mart application (UNIX)

Purpose: To install the Genesys Info Mart 8.1 application on a host that is running a UNIX operating system.

Prerequisites

- The host has been prepared, as described in "Preparing the Genesys Info Mart Server Host" on page 323.

Start of procedure

1. Insert the Genesys Info Mart CD into the CD-ROM drive of the machine on which you want to install Genesys Info Mart.
2. Locate the correct installation directory for your platform—for example, `genesys_info_mart/gim_etl/solaris`.
3. Save the contents of this directory to a local folder.

4. Locate and run the `install.sh` shell script. Enter requested information when you are prompted to do so.

   If you want to enable client-side port definition for the initial connection to Configuration Server, specify the required parameters when you get the `Client Side Port Configuration` prompt in the wizard. For more information, see the procedure about configuring a connection to Configuration Server on UNIX in the chapter about client-side port definition in the *Genesys 8.x Security Deployment Guide*.

End of procedure

Next Steps

- Verify the Genesys Info Mart host setup. See "Verifying Host Requirements" on page 329.
- (Optional) Modify the Genesys Info Mart startup arguments. For more information, see "Modifying the Genesys Info Mart Default Arguments" on page 359.
- Install the Genesys Info Mart management GUI. For more information, see Procedure: Installing the Genesys Info Mart Manager plug-in, on page 331 or Procedure: Installing the Genesys Info Mart Administration Console (Windows), on page 333.

Verifying Host Requirements

As described in "Preparing the Genesys Info Mart Server Host" on page 323, Genesys Info Mart has specific requirements for supporting system software components and operational parameters, such as JDBC drivers and environment variables. This section describes important system information that you cannot verify or modify before installation.

gim_etl_paths.bat File

After the installation, verify that the path in the `gim_etl_paths.bat` file correctly points to the Java executable. For example:

```
set JAVACMD=java.exe
```

Update the path in the `gim_etl_paths.bat` file, if necessary.

For related information about modifying the `PATH` and `JAVA_HOME` environment variables, see "Installing Java" on page 324.

Modifying the Configuration Checkup Properties File

On startup, or whenever the configuration of the Genesys Info Mart application changes, the Genesys Info Mart Server automatically performs a configuration

check, which includes verification that the correct JDBC driver version is available. If the available version is not as specified in a properties file (`gim_cfg_checkup.properties`), the configuration check logs an error, and the Genesys Info Mart Server will not run any jobs until the error has been fixed.

By default, the properties file specifies the required minimum version of Oracle Thin Driver (`oracle.jdbc.driver.OracleDriver=11.2.0.1.0`). For later Oracle versions, for Microsoft SQL Server deployments, or for PostgreSQL, you must modify the properties file to specify the installed version of the driver for Genesys Info Mart to use. After installation, the `gim_cfg_checkup.properties` file is available in the `resources` subfolder in the Genesys Info Mart installation folder.

Installing Genesys Info Mart Manager

Genesys Info Mart Manager is a web-based GUI that was introduced in Genesys Info Mart release 8.1.4, to enable the monitoring and real-time administration of Genesys Info Mart jobs. Genesys Info Mart Manager is included as a separate installation package (IP) on the Genesys Info Mart CD. Genesys Info Mart Manager is a Genesys Administrator Extension (GAX) plug-in, which you install on the same host as your GAX application.

Genesys Info Mart Manager interfaces directly with Genesys Info Mart Server to monitor the status of Genesys Info Mart jobs and to start and stop jobs on an ad-hoc basis.

Genesys Info Mart Manager provides essentially the same functionality as the Genesys Info Mart Administration Console (see page 36). You can use both GUIs in your deployment at the same time, if you want to enable additional administration resources but do not want to decommission an existing Administration Console.

Before You Proceed

Before you install Genesys Info Mart Manager:

1. Ensure that GAX release 8.5 or later has been installed, and verify that the Genesys Info Mart `Application` object is in the contact center configuration represented in GAX.

   For instructions about installing GAX, see the *Genesys Administrator Extension Deployment Guide,* 8.5.x release.

2. Verify that the users who will use Genesys Info Mart Manager have been assigned to access groups that have the necessary permissions and role privileges to access GAX plug-ins and the Genesys Info Mart Server.

(*Permissions* control access to objects configured in the Configuration Layer; *privileges* control access to functionality in GUI-based applications, such as GAX.)

If necessary, work with your GAX administrator to create and assign Users, Access Groups, or Roles as required to provide the following minimum access permissions and role privileges:

- User must have `Read` permission for the user's own `Persons` object and `Read` and `Execute` permissions for the GAX `Application` object. These permissions enable the user to log in to GAX.
- User must be a member of a Role that has the `Read Plugins` privilege. This enables the user to access Genesys Info Mart Manager after logging in to GAX.
- User must have `Read` permissions for the Genesys Info Mart `Application` object. This enables the user to monitor Genesys Info Mart Server. To enable the user to start and stop jobs as well, also provide `Execute` permissions.

For information about configuring permissions and privileges, see the *Genesys Administrator Extension 8.5 Deployment Guide.*

The page about managing plug-ins in the *Genesys Administrator Extension 8.5 Deployment Guide* describes how to use the Software Installation Wizard, together with the Genesys Deployment Agent (GDA), to install a plug-in. However, for Genesys Info Mart Manager specifically, Genesys recommends the following procedure, which is similar to the method for installing language packs. This method does not require you to have any particular privileges to administer GAX.

Procedure: Installing the Genesys Info Mart Manager plug-in

Purpose: To install Genesys Info Mart Manager on the host on which GAX is installed.

Prerequisites

- You have completed the tasks described in "Before You Proceed" on page 330.

Start of procedure

1. Stop GAX if it is running.
2. Insert the Genesys Info Mart CD into the CD-ROM drive of the machine on which GAX is installed.
3. Navigate to the `InfomartMgr\windows` or `InfomartMgr/<UNIX OS>` directory, as applicable, and open it.

4. To install on:
   - Windows, double-click the `setup.exe` file.
   - UNIX, run the `install.sh` script.

   Follow the directions in the installation wizard.

   The installation wizard copies three files into the GAX installation directory, including the `gim-manager.jar` file, which goes in the `<GAX installation directory>\webapp\WEB-INF\lib` folder.

End of procedure

Next Steps

- If desired, install language packs to localize Genesys Info Mart Manager.
  - For the available language packs, see the list of international release notes in docs.genesys.com.
  - For information about installing the language packs, see the page about managing plug-ins in the *Genesys Administrator Extension 8.5 Deployment Guide.*
- Perform post-installation activities. For more information, see Chapter 16 on page 337.

Installing the Genesys Info Mart Administration Console

The Genesys Info Mart Administration Console is a graphical user interface that enables the monitoring and real-time administration of some aspects of the Genesys Info Mart jobs. It is included as a separate installation package on the Genesys Info Mart CD. You install it on the same host as your Genesys Configuration Manager.

The Genesys Info Mart Administration Console is implemented by using the existing Genesys Framework. It interfaces with Genesys Info Mart Server in order to start, schedule, and stop Genesys Info Mart jobs on an ad-hoc basis.

Note: The Genesys Info Mart Administration Console is not supported on UNIX operating systems.

Before You Proceed

Before you install the Genesys Info Mart Administration Console, you must complete the following tasks:

1. Install and configure a DB Server, to enable the Genesys Info Mart Administration Console to access the Info Mart database.

   For complete instructions, see Appendix A, "Standard Configuration Procedure," in the *Framework 8.x Deployment Guide.*
2. Create and configure a Genesys Info Mart Administration Console DAP. For complete instructions, see Procedure: Configuring a dedicated Administration Console DAP, on page 224.
3. Add the Genesys Info Mart Administration Console DAP to the `Connections` tab of the Genesys Info Mart `Application` object. See Step 7 on page 237.

Procedure: Installing the Genesys Info Mart Administration Console (Windows)

Purpose: To install the Genesys Info Mart Administration Console on a host that is running a Windows operating system.

Prerequisites

- You have completed the tasks that are described in "Before You Proceed" on page 321.
- Genesys recommends that you close Configuration Manager before you install the Genesys Info Mart Administration Console.

Start of procedure

1. Insert the Genesys Info Mart CD into the CD-ROM drive of the machine on which you want to install the Genesys Info Mart Administration Console.
2. Navigate to the `genesys-info-mart\Admin-Console\windows` directory, and open it.
3. Double-click the `setup.exe` file, and then follow the directions in the installation wizard.

End of procedure

Next Steps

- Perform post-installation activities. For more information, see Chapter 16 on page 337.

Part

3 Post-Deployment Activities

Part 3 of this document provides information about required and optional activities after you have installed the Genesys Info Mart application.

This information appears in the following chapters:

- Chapter 16, “Post-Installation Activities,” on page 337
- Chapter 17, “Starting and Stopping Genesys Info Mart Server,” on page 357

Chapter

16 Post-Installation Activities

This chapter describes the tasks that you must complete after you finish configuring and installing Genesys Info Mart and the Genesys Info Mart management GUI. It contains the following sections:

Task Flow for Post-Installation Activities

Task Summary: Post-Installation Activities summarizes the task flow for activities that are required to complete the Genesys Info Mart deployment.

Task Summary: Post-Installation Activities

Objective	Related Procedures and Actions
1. Initialize Genesys Info Mart.	Start the Genesys Info Mart Server application. On startup, Genesys Info Mart automatically performs a configuration check and runs `Job_InitializeGIM`, which initializes the Info Mart database. Starting with release 8.1.2, `Job_InitializeGIM` also automatically executes the scripts to modify the Interaction Database (IDB) schemas for use with Genesys Info Mart (see "Preparing IDBs" on page 178). For more information about how to start and stop Genesys Info Mart, see Chapter 17 on page 357.

Task Summary: Post-Installation Activities (Continued)

Objective	Related Procedures and Actions
2. Access the management GUI.	See one of the following, as applicable for your deployment: • Procedure: Accessing Genesys Info Mart Manager, on page 340 • Procedure: Accessing the Genesys Info Mart Administration Console, on page 342.
3. Verify the deployment.	1. Review the `gim_etl` log file to confirm the results of the configuration check. If necessary, modify the configuration option settings or connections in the Genesys Info Mart, Interaction Concentrator (ICON), or data-source `Application` objects in your deployment, to ensure that they have been configured for correct Genesys Info Mart functioning. For more information about the configuration check, see "Deployment Verification" on page 401. For more information about how to configure the required applications for Genesys Info Mart, see Part 2 on page 133. **Note:** After Genesys Info Mart checks the configuration, the console and the `gim_etl` log file might report results for configuration options that you cannot modify. Any options that are not documented in Chapter 14 on page 271 are not configurable. 2. Review the status of `Job_InitializeGIM` in the management GUI, to verify successful completion. Completion of the job indicates successful initialization of the Info Mart database and, starting with release 8.1.2, successful update of the IDBs. **Note:** If `Job_InitializeGIM` fails, the `Status` column in the management GUI will display `FAILED` to the right of the job. To determine the nature of the error, refer to the `gim_etl` log file.
4. (Optional) Define online media types for 3rd Party Media interactions.	Execute an `INSERT` command to add a record for each online media type to the `MEDIA_TYPE` dimension table, with the `IS_ONLINE` flag set to `1` (online). Otherwise, when Genesys Info Mart encounters the unknown media type during transformation, it will add the media type dynamically to the `MEDIA_TYPE` table, but with the `IS_ONLINE` flag set to `0` (offline). Therefore, you do not have to add offline media types in advance. For more information about how to add the media type to the `MEDIA_TYPE` table, see Procedure: Setting up media types for online interactions, on page 344. For more information about online and offline media types in Genesys Info Mart, see "Online and Offline Interactions" on page 98.

Task Summary: Post-Installation Activities (Continued)

Objective	Related Procedures and Actions
5. Run the first ETL cycle.	1. If you did not set scheduling options when you configured the Genesys Info Mart `Application` object, review the configuration options in the `[schedule]` section and, if necessary, modify them to suit your deployment. For more information about the scheduling options, see "schedule Section" on page 314. 2. Wait for the first scheduled ETL cycle to run, or run the ETL jobs manually for the first time. 3. Review the jobs status in the management GUI and the logs to verify that the jobs completed successfully. For more information about scheduling considerations and about how to run ETL jobs, see the chapter about working with ETL jobs in the *Genesys Info Mart 8.1 Operations Guide*.
6. Complete deployment of the Info Mart database, to prepare it for reporting queries.	1. (Optional) Create custom calendars. For more information, see "Creating Custom Calendars" on page 347. 2. Using your prepared RDBMS user accounts, which have been configured with the proper privileges, run the scripts to create Genesys Info Mart read-only views. This procedure is strictly required only if you have a multi-tenant deployment, but Genesys recommends this for all deployments. See Procedure: Creating Genesys Info Mart Read-Only Tenant Views, on page 348. **Note:** You cannot create the read-only views until the first ETL cycle has completed and the `TENANT` dimension has been populated.
7. (Optional) Disable processing of unwanted interaction subtypes.	Execute an `ALTER TABLE` command to set the `IGNORE` flag to `1` (ignore) in the `INTERACTION_TYPE` record for each unwanted interaction subtype. You can change the `IGNORE` flag for a particular interaction subtype during runtime.

Accessing the Management GUIs

To monitor and manage the day-to-day operations of Genesys Info Mart, you can use Genesys Info Mart Manager or the Genesys Info Mart Administration Console. Genesys Info Mart Manager is available starting with Genesys Info Mart release 8.1.4.

Accessing Genesys Info Mart Manager

Genesys Info Mart Manager is a web-based GUI that is available through Genesys Administrator Extension (GAX).

Multiple instances of the Genesys Info Mart Manager can be running at the same time, each associated with the same instance of the Genesys Info Mart Server.

Procedure: Accessing Genesys Info Mart Manager

Purpose: To access Genesys Info Mart Manager to manage Genesys Info Mart jobs.

Prerequisites

- The Genesys Info Mart Manager plug-in has been installed, as described in "Installing Genesys Info Mart Manager" on page 330.
- A Genesys user account with the appropriate permissions and role privileges to access plug-ins and the Genesys Info Mart application has been provisioned. For more information, see "Before You Proceed" on page 330.
- The Genesys Info Mart Server application is started.

Start of procedure

1. Use a web browser to log in to GAX.
2. Select `Administration > GIM Manager`.

   The GIM Manager screen appears, displaying the status of the Genesys Info Mart jobs (see Figure 16 on page 341).
3. If there is more than one Genesys Info Mart Server in your deployment, use the drop-down list box to choose a server to manage.

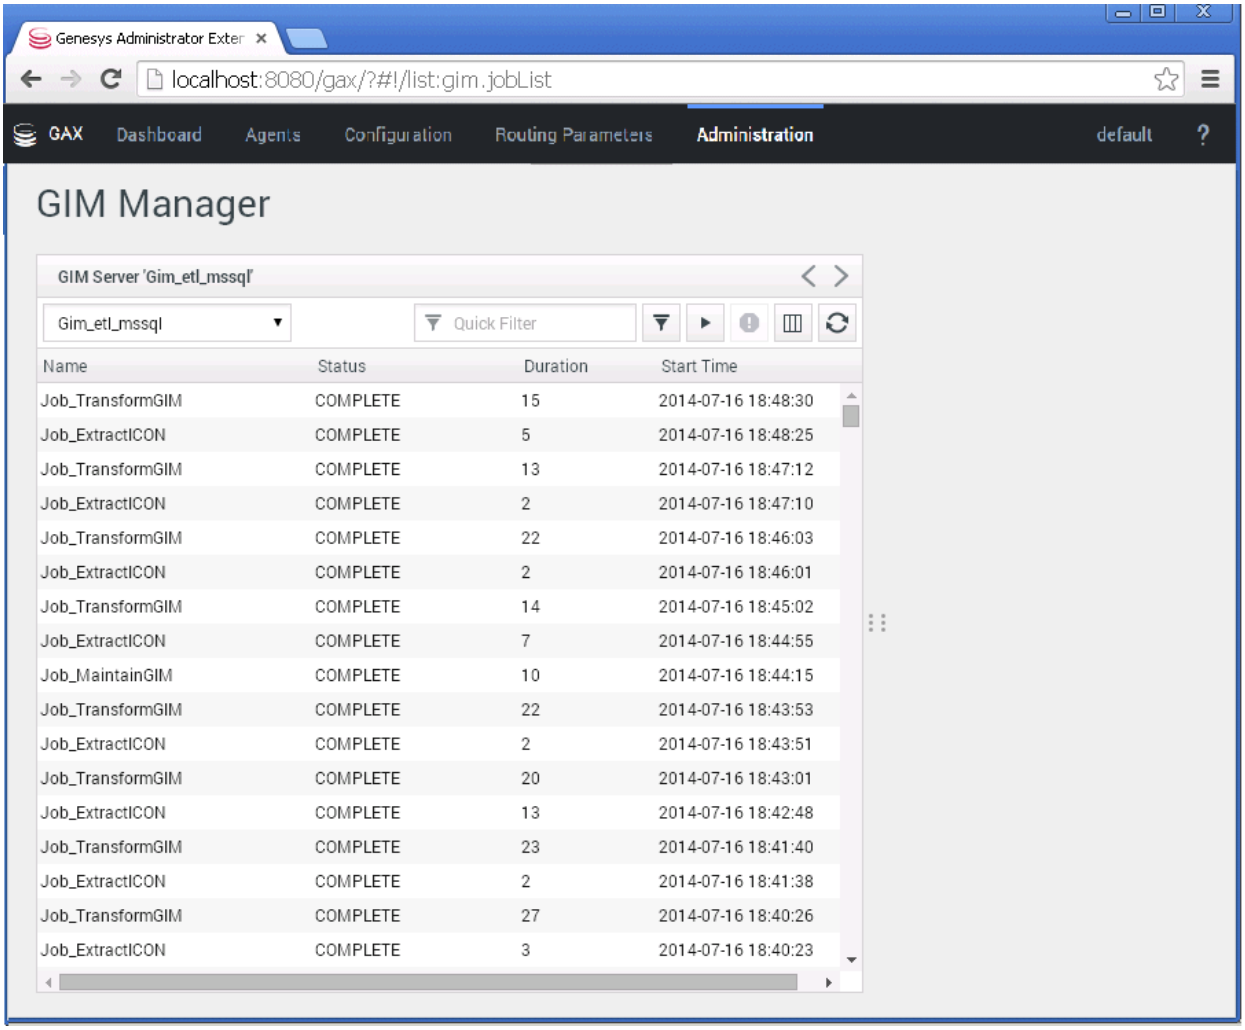

Figure 16: Genesys Info Mart Manager

End of procedure

Next Steps

- For more information about how to use Genesys Info Mart Manager to start and stop jobs, and to view job status, see the *Genesys Info Mart 8.1 Operations Guide.*
- Perform additional post-installation activities, as required. For an overview of post-installation tasks, see “Task Flow for Post-Installation Activities” on page 337.

Accessing the Genesys Info Mart Administration Console

In Genesys Info Mart 8.x, the Genesys Info Mart Administration Console is implemented as an extension to Genesys Configuration Manager by using the existing Wizard Framework.

Multiple instances of the Genesys Info Mart Administration Console can be running at the same time, each associated with the same instance of the Genesys Info Mart Server.

Procedure: Accessing the Genesys Info Mart Administration Console

Purpose: To access the Genesys Info Mart Administration Console to manage Genesys Info Mart jobs.

Prerequisites

- The Genesys Info Mart Administration Console has been installed, as described in "Installing the Genesys Info Mart Administration Console" on page 332.
- You must start the Genesys Info Mart Server application before you start the Genesys Info Mart Administration Console.

Start of procedure

1. Open Configuration Manager.
2. Select the `Application` object for the Genesys Info Mart ETL that you want to manage.
3. Right-click the Genesys Info Mart ETL `Application`, and then select `Wizard > Configure` (see Figure 17 on page 343).

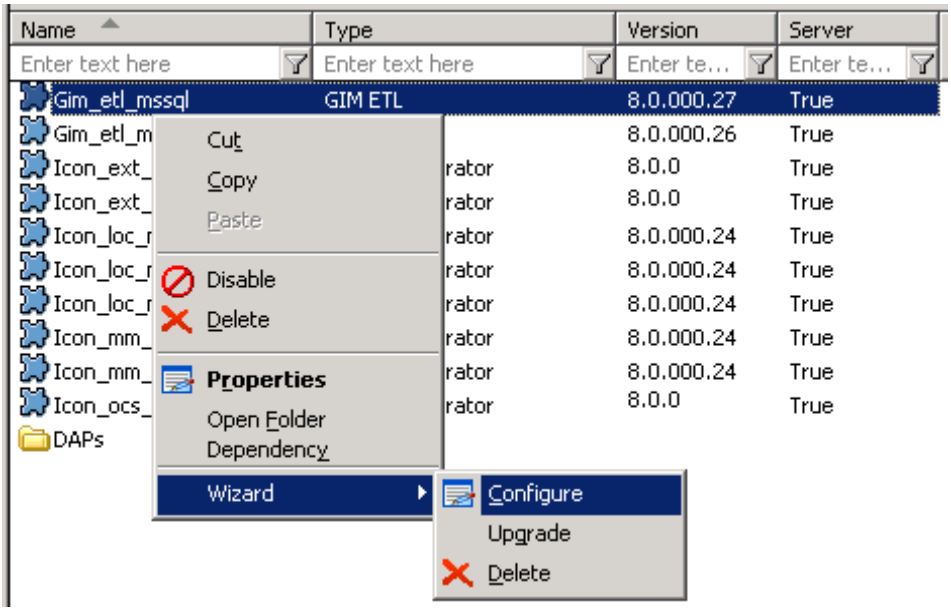

Figure 17: Accessing the Genesys Info Mart Administration Console from the Configuration Manager

The `GIM Admin Console` dialog box appears, displaying the status of the Genesys Info Mart jobs (see Figure 18).

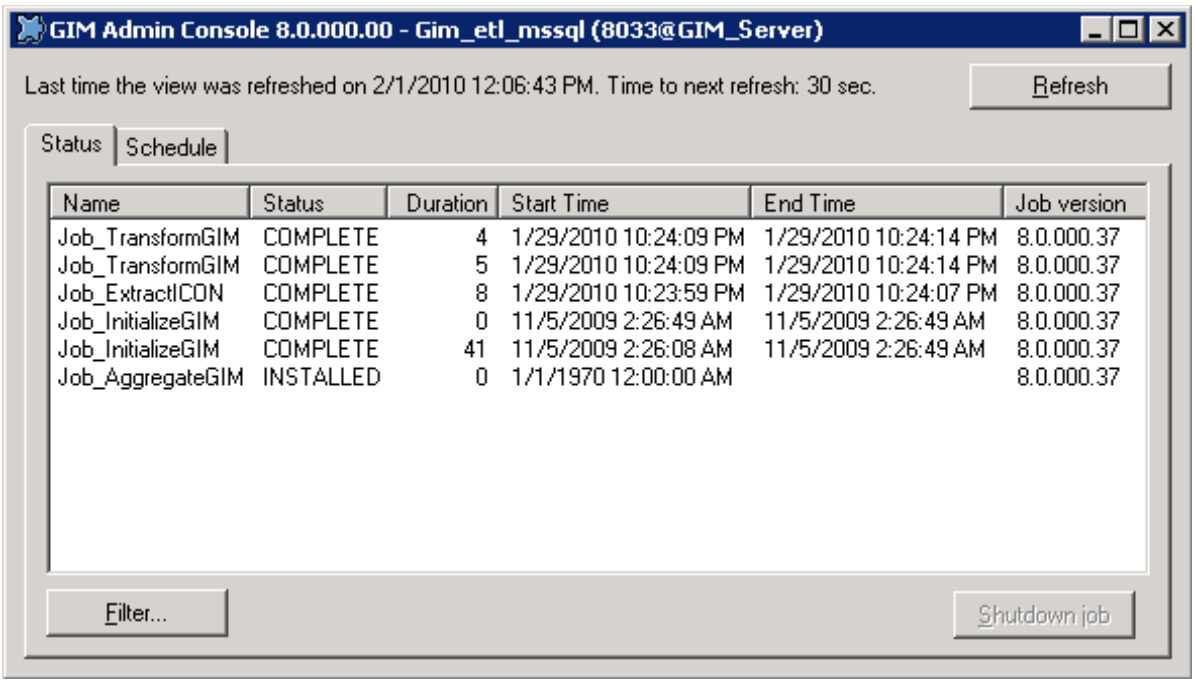

Figure 18: GIM Admin Console Dialog Box

End of procedure

Next Steps

- For more information about how to use the Genesys Info Mart Administration Console to start and stop jobs, and to view job status, see the *Genesys Info Mart 8.1 Operations Guide.*
- Perform additional post-installation activities, as required. For an overview of post-installation tasks, see "Task Flow for Post-Installation Activities" on page 337.

Preparing the Info Mart Database for 3rd Party Media

The Genesys Info Mart transformation job processes online and offline interactions differently. (For definitions of *online* and *offline* as used in Genesys Info Mart, see page 98.) The `IS_ONLINE` field in the `MEDIA_TYPE` dimension table identifies whether a particular media type will be associated with online interactions or with offline interactions.

Genesys Info Mart will dynamically add to the `MEDIA_TYPE` table any unknown media types that it encounters during transformation, but it identifies them as media types for offline interactions (`IS_ONLINE=0`, by default). Genesys recommends specifying new online media types in advance, to ensure that they are represented correctly in interaction data from the start.

Note: It is only new *online* media types that Genesys recommends that you add manually in advance, because *offline* media types that are added dynamically by Genesys Info Mart will be identified correctly as offline from the start.

The following procedure describes how to set up the Info Mart database to handle online 3rd Party Media interactions.

Procedure: Setting up media types for online interactions

Purpose: To prepopulate the `MEDIA_TYPE` dimension table with media types that Genesys Info Mart might encounter in 3rd Party Media interactions, so that they will be processed as online interactions.

You can perform this procedure at any time after the Info Mart database schema has been initialized.

Prerequisites

- The Info Mart database has been initialized (see Step 1 on page 337).
- You have the required permissions to alter the Info Mart database.

Start of procedure

1. Log in to the Info Mart database as any user with `INSERT` and `ALTER` permissions.
2. Check the current contents of the `MEDIA_TYPE` table to identify the last `MEDIA_TYPE_KEY` value. Media types for 3rd Party Media interactions have values of 1001 or higher.
3. Execute an SQL `INSERT` command to add the desired media types to the `MEDIA_TYPE` table. For each new media type, specify values for the following fields:
   - `MEDIA_TYPE_KEY`—The unique primary key. Using the information from Step 2, specify the next available value greater than `1000` (`1001` or higher).
   - `MEDIA_NAME`—Any name, up to 64 characters, that you want to use to identify the media type.
   - `MEDIA_NAME_CODE`—The name of the media type that Genesys Info Mart uses internally. The code name can be up to 32 characters, must be unique, and must exactly match the name of the corresponding `Media Type Business Attribute`, including case.
   - `IS_ONLINE`—The online/offline flag. For online media types, specify a value of `1`.
   - `CREATE_AUDIT_KEY`—The lineage for data creation. Specify a value of `-1`, which indicates that the row was not inserted by Genesys Info Mart.
   - `UPDATE_AUDIT_KEY`—The lineage for data update. Specify a value of `-1`, which indicates that the row was not updated by Genesys Info Mart.

   Figure 19 on page 346 shows the contents of a sample `MEDIA_TYPE` table in which:
   - `None`, `Voice`, `Email`, and `Chat` are the default media type dimensions.
   - `SampleOnlineMedia1` is an online 3rd Party Media media type that was added manually in advance.
   - `SampleOfflineMedia` is an offline 3rd Party Media media type that was added dynamically by Genesys Info Mart during runtime.
   - `SampleOnlineMedia2` is an online 3rd Party Media media type that was not added in advance (for example, because the user forgot or because it is new to the deployment). The media type was added dynamically by Genesys Info Mart (as an offline media type) during runtime, and then the `IS_ONLINE` field in the record was manually altered to indicate that it is an online media type.

MEDIA_TYPE_KEY	MEDIA_NAME	MEDIA_NAME_CODE	IS_ONLINE	CREATE_AUDIT_KEY	UPDATE_AUDIT_KEY
0	None	NONE	1	2	0
1	Voice	VOICE	1	2	0
2	Email	EMAIL	0	2	0
3	Chat	CHAT	1	2	0
1001	SampleOnlineMedia1	SampleOnlineMedia1	1	-1	-1
1002	SampleOfflineMedia	SampleOfflineMedia	0	5	0
→ 1003	SampleOnlineMedia2	SampleOnlineMedia2	1	5	-1

Figure 19: Sample MEDIA_TYPE Table

For more information about the required data types and the meaning of the fields, see the *Genesys Info Mart 8.1 Reference Manual* for your RDBMS.

If you are executing this procedure during runtime and you encounter a constraint violation, go to Step 4. Otherwise, go to Step 5 on page 347.

4. Constraint violations, if they occur, are most likely to occur if you are manually adding media types while the transformation job is running.
   - You might encounter a constraint violation when you execute the INSERT command if:
     - The media type has already been added. Check the contents of the MEDIA_TYPE table to identify if a record with that MEDIA_NAME_CODE already exists, and verify the value of the IS_ONLINE flag. If the record exists because Genesys Info Mart added it dynamically during transformation, execute an ALTER TABLE command to change the value of the IS_ONLINE flag to 1.
     - Genesys Info Mart dynamically added a media type to the MEDIA_TYPE table and used a MEDIA_TYPE_KEY value that you were planning to use. Repeat Steps 2 and 3 on page 345 to add the desired media type with a different MEDIA_TYPE_KEY value.
   - Genesys Info Mart might encounter a constraint violation if the transformation job was coincidentally trying to add an unknown media type at the exact time that you were manually adding it or using the same MEDIA_NAME_KEY to add another media type, and your insert took effect first. In this case, no action is required.
     - The transformation job will complete unsuccessfully.
     - In the next ETL cycle, the transformation job will take the appropriate action—either recognizing the media type that you inserted, or else using a different MEDIA_TYPE_KEY to insert the new media type.

5. Commit the change to the database.

End of procedure

Creating Custom Calendars

To enable you to customize the calendars that you use in your reporting, Genesys Info Mart supports multiple custom calendar dimensions. The following procedure describes the steps to create a custom calendar.

Procedure: Configuring custom calendars

Purpose: To modify the Genesys Info Mart `Application` object and Info Mart database schema to provide additional calendar dimensions for use in reports.

Prerequisites

- You have the required permissions to access and modify the Genesys Info Mart `Application` object, if required (see Step 1).
- The database access account that you use to create the Info Mart database schema is available and has the required privileges (see Table 3 on page 95).

  Refer to the "Installation Worksheets" beginning on page 421 to determine the ID to use.

Start of procedure

1. Configure the Genesys Info Mart application to support the required custom calendars, if you did not configure the required options before installing the application.
   a. On the `Options` tab of the Genesys Info Mart `Application` object, create a new section that has a section name that starts with the prefix *date-time-* (for example, `[date-time-emea]`).
   b. Create new options that have the same names as the options in the default `[date-time]` section, and configure the options as required for your reporting purposes.

      For more information about the `[date-time]` options, see "date-time Section" on page 271.

2. Prepare the script to create the custom dimension tables in the Info Mart database schema:
   a. Locate the `make_gim.sql` or `make_gim_partitioned.sql` script in the `sql-scripts` folder in your Genesys Info Mart 8.1 installation package. (This is the script that you ran to create the Info Mart database schema.)
   b. Copy the SQL statements that create the default `DATE_TIME` table, and use them as the basis for a custom script to create additional calendar table(s), with the same structure and column names as the default `DATE_TIME` table.

      The table name must match the value of the `date-time-table-name` option that you specified in the custom `[date-time-*]` configuration section (see Step 1 on page 347).
3. Log in to the Info Mart database by using the Info Mart owner ID.
4. Execute the script that you created in Step 2.

End of procedure

Next Steps

- Verify that the custom dimension tables have been populated.

  If you configure the custom calendars before you start Genesys Info Mart for the first time, `Job_InitializeGIM` will populate the custom calendar dimension tables. If you configure the custom calendars after `Job_InitializeGIM` has run, the custom calendar dimension tables will be populated at the next run of `Job_MaintainGIM`.
- If you have not already done so, create read-only views for your reporting application to query. See Procedure: Creating Genesys Info Mart tenant views—Oracle.

Creating Genesys Info Mart Read-Only Tenant Views

Read-only views allow you to control end-user access to Info Mart data.

Creating read-only, tenant-specific views is strictly required only for multi-tenant deployments. However, even in single-tenant deployments, Genesys recommends that you consider setting up a tenant-specific, read-only view for security reasons: a read-only view helps protect the data, because your queries do not work directly on the database tables.

The Genesys Info Mart 8.1 installation package includes a script, `make_gim_view_for_tenant.sql`, that you run in order to create tenant-specific, read-only views on the Genesys Info Mart tables.

In 8.1.0 and 8.1.1 releases, Genesys Info Mart supports creating table views for only one tenant per Tenant User schema. Since each Tenant User schema has a view on a single DATE_TIME table, each schema supports a single time zone. Starting with release 8.1.2, Genesys Info Mart supports creating table views for more than one tenant in the same Tenant User schema. This functionality simplifies deployment of the reporting solution.

Reports in Multiple Time Zones

For example, to provide reports in multiple time zones, the downstream report developer must use a separate Tenant User schema for each time zone. In releases prior to 8.1.2, the Tenant Admin must create a separate Tenant User schema for each combination of time zone and tenant. Starting with release 8.1.2, while separate schemas must still be created for each time zone, the Tenant Admin can include all tenants, or a group of tenants, in a single schema per time zone.

Figure 20 on page 350 illustrates how the tenant-specific, read-only Genesys Info Mart Views function in a Genesys Info Mart deployment. The diagram shows a Tenant User schema that contains table views for only one tenant.

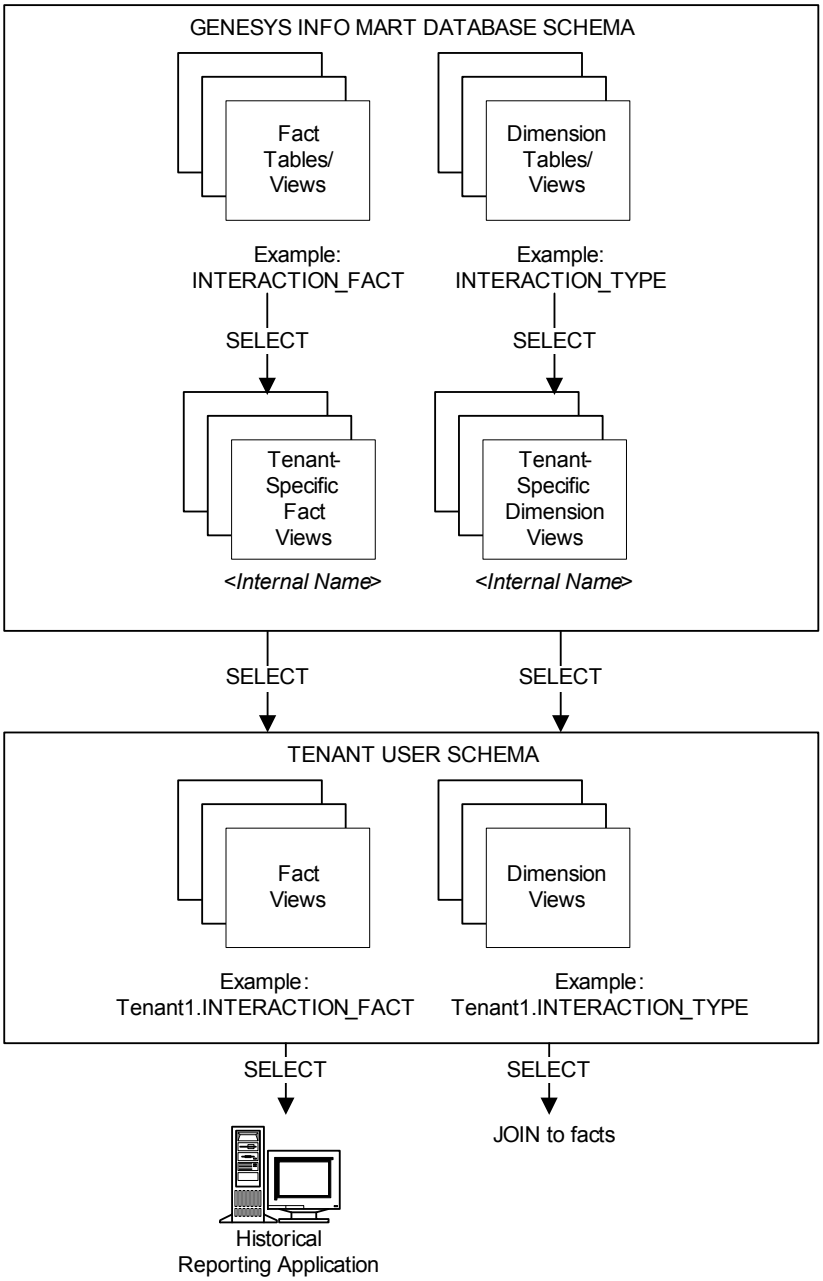

Figure 20: Multi-Tenant Filtered Views or Single-Tenant Read-Only View

The following procedures describe the modifications you must make to the `make_gim_view_for_tenant.sql` script for your RDBMS, in order to create the tenant views in your deployment.

- Procedure: Creating Genesys Info Mart tenant views—Oracle, on page 351
- Procedure: Creating Genesys Info Mart tenant views—Microsoft SQL Server, on page 352
- Procedure: Creating Genesys Info Mart tenant views— PostgreSQL, on page 354

The script that creates each tenant view first drops the view, so the script can be rerun if necessary (for example, if you add new custom user-data tables to the Info Mart database schema).

Procedure: Creating Genesys Info Mart tenant views—Oracle

Purpose: In a tenant user schema, to create tenant-specific, read-only views on database objects (database tables and views) that are in the Info Mart schema in Oracle.

Repeat this procedure for each tenant in your Genesys Info Mart configuration.

Prerequisites

- Use the "Worksheet for Oracle Databases" on page 422 to identify the database connection parameters for each database.
- You have obtained the `TENANT_KEY`, which Genesys Info Mart uses to identify the tenant, in one of the following ways:
  - You have obtained the database identifier that Configuration Server assigned to the tenant (Tenant DBID). The `TENANT_KEY` matches the Tenant DBID.
  - You have successfully run the first ETL cycle, so that the `TENANT` dimension has been populated, and then you queried the `TENANT` dimension table in the Info Mart view that is named `TENANT`, to obtain the `TENANT_KEY`. (The `TENANT` view was created by the `make_gim.sql` or `make_gim_partitioned.sql` script when you created the Info Mart database schema.)
- The Tenant User user accounts are available. Refer to the "Worksheet for Oracle Databases" on page 422 to determine the IDs to use.

Start of procedure

1. Log in to the Info Mart database as the Tenant Admin database user.
2. In a text editor, modify the `make_gim_view_for_tenant.sql` script to provide the required parameters. The modifications for release 8.1.2 and later deployments are different from the modifications you must make for release 8.1.0 and 8.1.1 deployments.

Release 8.1.2 and Later

For release 8.1.2 and later, make the following modifications:

- `&&1`—Replace all instances of *&&1* with the name of the Tenant User schema.
- `&&2`—Replace all instances of *&&2* with the name of the Info Mart schema.

- `&&3`—Replace all instances of *&&3* with one of the following:
  - — `all_tenants`—To create views for all the tenants that have been configured for the deployment
  - — A list of Tenant keys (for example: `101, 102, 103`)—To create views for a group of tenants
  - — The Tenant key—To create views for a single tenant

Release 8.1.0 or 8.1.1

For release 8.1.0 or 8.1.1, make the following modifications:

- `&&1`—Replace all instances of *&&1* with the name of the tenant user who will access the view (Tenant User).
- `&&2`—Replace all instances of *&&2* with the name of the Info Mart user.
- `&&3`—Replace all instances of *&&3* with the number of the Genesys Info Mart Tenant key.

Note: Ensure that you use upper case for the user names. For example, if the name of the Tenant User is `Tenant1`, replace all instances of `&&1` with `TENANT1`.

3. Run the modified `make_gim_view_for_tenant.sql` script.

End of procedure

Next Steps

- If you are enabling optional functionality, see the task summaries that are provided in "Enabling Specific Functionality" on page 140. Otherwise, your Genesys Info Mart deployment process is complete.

Procedure: Creating Genesys Info Mart tenant views—Microsoft SQL Server

Purpose: In a tenant user schema, to create tenant-specific, read-only views on database objects (database tables and views) that are in the Info Mart schema in Microsoft SQL Server.

Repeat this procedure for each tenant in your Genesys Info Mart configuration.

Prerequisites

- Use the "Worksheet for Microsoft SQL Databases" on page 424 to identify the database connection parameters for each database.
- You have obtained the `TENANT_KEY`, which Genesys Info Mart uses to identify the tenant, in one of the following ways:

- You have obtained the database identifier that Configuration Server assigned to the tenant (Tenant DBID). The `TENANT_KEY` matches the Tenant DBID.
- You have successfully run the first ETL cycle, so that the `TENANT` dimension has been populated, and then you queried the `TENANT` dimension table in the Info Mart view that is named `TENANT`, to obtain the `TENANT_KEY`. (The `TENANT` view was created by the `make_gim.sql` or `make_gim_partitioned.sql` script when you created the Info Mart database schema.)

- The Tenant User user accounts are available. Refer to the "Worksheet for Microsoft SQL Databases" on page 424 to determine the IDs to use.

Start of procedure

1. Log in to the Info Mart database as the Tenant Admin database user.
2. In a text editor, modify the `make_gim_view_for_tenant.sql` script to provide the required parameters. The modifications for release 8.1.2 and later deployments are different from the modifications you must make for release 8.1.0 and 8.1.1 deployments.

Release 8.1.2 and Later

For release 8.1.2 and later, make the following modifications:

- `&&1`—Replace all instances of *&&1* with the name of the Tenant User schema.
- `&&2`—Replace all instances of *&&2* with the name of the Info Mart schema.
- `&&3`—Replace all instances of *&&3* with one of the following:
  - `all_tenants`—To create views for all the tenants that have been configured for the deployment
  - A list of Tenant keys (for example: `101, 102, 103`)—To create views for a group of tenants
  - The Tenant key—To create views for a single tenant
- `&&4`—Replace all instances of *&&4* with the name of the tenant user who will access the view.

Release 8.1.0 or 8.1.1

For release 8.1.0 or 8.1.1, make the following modifications:

- `&&1`—Replace all instances of *&&1* with the name of the Tenant User schema.
- `&&2`—Replace all instances of *&&2* with the name of the Info Mart schema.
- `&&3`—Replace all instances of *&&3* with the number of the Genesys Info Mart Tenant key.
- `&&4`—Replace all instances of *&&4* with the name of the tenant user who will access the view (Tenant User).

3. Run the modified `make_gim_view_for_tenant.sql` script.

End of procedure

Next Steps

- If you are enabling optional functionality, see the task summaries that are provided in "Enabling Specific Functionality" on page 140. Otherwise, your Genesys Info Mart deployment process is complete.

Procedure: Creating Genesys Info Mart tenant views—PostgreSQL

Purpose: In a tenant user schema, to create tenant-specific, read-only views on database objects (database tables and views) that are in the Info Mart schema in PostgreSQL.

Repeat this procedure for each tenant in your Genesys Info Mart configuration.

Prerequisites

- Use the "Worksheet for PostgreSQL Databases" on page 426 to identify the database connection parameters for each database.
- You have obtained the `TENANT_KEY`, which Genesys Info Mart uses to identify the tenant, in one of the following ways:
  - You have obtained the database identifier that Configuration Server assigned to the tenant (Tenant DBID). The `TENANT_KEY` matches the Tenant DBID.
  - You have successfully run the first ETL cycle, so that the `TENANT` dimension has been populated, and then you queried the `TENANT` dimension table in the Info Mart view that is named `TENANT`, to obtain the `TENANT_KEY`. (The `TENANT` view was created by the `make_gim.sql` or `make_gim_partitioned.sql` script when you created the Info Mart database schema.)
- The Tenant User user accounts are available. Refer to the "Worksheet for PostgreSQL Databases" on page 426 to determine the IDs to use. Starting with release 8.1.4, the name of the Tenant User schema must be the same as the name of the corresponding tenant user.

Start of procedure

1. Log in to the Info Mart database as the Tenant Admin database user.
2. In a text editor, modify the `make_gim_view_for_tenant.sql` script to provide the required parameters. The modifications for release 8.1.4 and later deployments are different from the modifications you must make for release 8.1.3, when support for PostgreSQL was introduced.

Release 8.1.4 and Later

For release 8.1.4 and later, make the following modifications:

- `&&1`—Replace all instances of *&&1* with the name of the Tenant User schema.
- `&&2`—Replace all instances of *&&2* with the name of the Info Mart schema.
- `&&3`—Replace all instances of *&&3* with one of the following:
  - `all_tenants`—To create views for all the tenants that have been configured for the deployment
  - A list of Tenant keys (for example: `101, 102, 103`)—To create views for a group of tenants
  - The Tenant key—To create views for a single tenant

Release 8.1.3

For release 8.1.3, make the following modifications:

- `&&1`—Replace all instances of *&&1* with the name of the Tenant User schema.
- `&&2`—Replace all instances of *&&2* with the name of the Info Mart schema.
- `&&3`—Replace all instances of *&&3* with one of the following:
  - `all_tenants`—To create views for all the tenants that have been configured for the deployment
  - A list of Tenant keys (for example: `101, 102, 103`)—To create views for a group of tenants
  - The Tenant key—To create views for a single tenant
- `&&4`—Replace all instances of *&&4* with the name of the tenant user who will access the view.

3. Run the modified `make_gim_view_for_tenant.sql` script.

End of procedure

Next Steps

- If you are enabling optional functionality, see the task summaries that are provided in "Enabling Specific Functionality" on page 140. Otherwise, your Genesys Info Mart deployment process is complete.

Chapter

17 Starting and Stopping Genesys Info Mart Server

This chapter describes the prerequisites for Genesys Info Mart Server startup and provides instructions for starting and stopping Genesys Info Mart Server. It contains the following sections:

Overview

You can start and shut down Genesys Info Mart Server by using the Genesys Management Layer, a startup file, a manual procedure, or Services Manager.

All of these methods usually require command-line parameters for a server application, as well as an executable file name. This chapter includes information about the command-line parameters that are common to most Genesys server applications. It also describes various methods of starting and stopping the Genesys Info Mart Server application.

Note: For information about how to use the Management Layer, startup files, and Services Manager for startup, see the *Framework 8.x Deployment Guide.*

Before You Begin

The following issues are important for you to consider before you attempt to start Genesys Info Mart Server.

Verifying Genesys Info Mart Connections and Configuration

Before you attempt to start Genesys Info Mart Server, confirm that the connections and configuration options that have been configured for your Genesys Info Mart `Application` object are correct for your deployment.

Connections

In general, do not change any connections on the `Connections` tab of the Genesys Info Mart `Application` object during startup. For more information about how to configure connections, see the step about configuring the "Connections Tab" on page 237, in the procedure to configure the Genesys Info Mart `Application` object.

Configuration Options

Do not make changes to Genesys Info Mart configuration options during startup. You can make changes to Genesys Info Mart configuration options during runtime; in almost all cases, you do not need to restart Genesys Info Mart for the changes to take effect. For more information, see "Configuring Options for Genesys Info Mart" on page 238.

Other Applications

Genesys recommends that the following applications be running before you start Genesys Info Mart Server:

- Configuration Server
- Message Server
- The relational database management system (RDBMS)
- The data sources and Interaction Concentrator (ICON) applications from which Genesys Info Mart obtains data

If your deployment includes attached data, ensure that there is a proper attached data specification file in the ICON working directory, suitably modified for use with Genesys Info Mart. (By default, ICON uses the `ccon_adata_spec.xml` file.) For more information about modifying the attached data specification file, see "Customizing Your ICON Attached Data Specification File" on page 176.

For detailed instructions about starting the Genesys components on which Genesys Info Mart depends, see the:

- *Framework 8.x Deployment Guide*
- *Framework 8.x T-Server Deployment Guide* for your particular T-Server type
- *Framework 8.x DB Server User's Guide*

- *Interaction Concentrator 8.x Deployment Guide*
- *Outbound Contact 8.x Deployment Guide*
- *eServices (Multimedia) 8.x Deployment Guide*

Command-Line Parameters

Genesys Info Mart supports the following startup command-line parameters:

`-host`	The name of the host on which Configuration Server is running
`-port`	The communication port that client applications must use to connect to Configuration Server
`-app`	The exact name of an application as configured in the Configuration Database
`-service`	(Windows only) The name of the Windows service

Genesys Info Mart also supports the following optional command-line parameters, which configure client-side port definition for a secure connection to Configuration Server:

`-transport-port`	The port number that the client application will use for the TCP/IP connection
`-transport-address`	The IP address that the client application will use for the TCP/IP connection

In addition, the `gim_etl_server` file, which is included in the Genesys Info Mart installation package (IP), specifies a number of Java Virtual Machine (JVM) startup parameters. For information about JVM parameters that you might want to modify, see "Modifying JVM Startup Parameters" on page 360.

Modifying the Genesys Info Mart Default Arguments

You might want to change the Genesys Info Mart application's default arguments; for example, change the host on which Configuration Server is running.

From Configuration Manager

If you plan to use Genesys Solution Control to start and stop Genesys Info Mart Server, you can edit the default command-line arguments directly on the `Start Info` tab of the Genesys Info Mart `Application` object in the Configuration Manager. These changes will take effect when you use Genesys Solution Control to start the application.

The following example (provided for a Windows operating system) shows the command-line arguments that are created under the `Start Info` tab of the Genesys Info Mart `Application` object in Configuration Manager after the installation is complete:

```
-host<lhostname> -port <port_number> -app InfoMart -service InfoMart
```

Note: The Windows service name of the Genesys Info Mart application is identified by the `-service` parameter. If Genesys Info Mart is installed multiple times, there will be a Windows service for each installation, in which the value of the parameter `-service` will be the service name of the corresponding Genesys Info Mart application in the Windows service. Genesys recommends that you *not* change the value of this parameter.

Modifying JVM Startup Parameters

The `gim_etl_server` file specifies a number of JVM parameters that set the overall context for Genesys Info Mart functioning—for example, Java memory settings. You might need to modify the file for performance reasons or to enable certain Genesys Info Mart or RDBMS functionality.

In particular, consider modifying the `gim_etl_server` file to set the following startup parameters:

- `-DfilterUserData=false`—Starting with release 8.1.101.07, enables the extraction job to extract user data that was sourced from the `UserData`, `Reasons`, or `Extensions` attributes in TEvents or Interaction Server events. Otherwise, the extraction job extracts KVPs from the `UserData` attribute only. For more information, see "Source Attributes in Events" on page 107.
- `-DqueryParallelism=<degree of parallelism>`—(For Oracle only) Starting with release 8.1.103.03, enables you to control the degree of query parallelism that Oracle will use, to improve performance. Genesys Info Mart uses the value that is specified by the startup parameter to provide a parallelism hint in certain transformation queries.

  By default, Genesys Info Mart now uses query parallelism at a level that is expected to be optimal for most deployments (`queryParallelism=4`).

  For more information about how parallel execution works, see the Oracle documentation.
- `-Dcfg.protocol.timeout=<timeout value in seconds>`—Starting with release 8.1.202.01, enables you to change the timeout value for a response from Configuration Server. The default timeout of 60 seconds (30 seconds in Genesys Info Mart releases earlier than 8.1.202.01) is appropriate in most installations. If you encounter error messages that the Genesys Info

Mart Server timed out while reading configuration data, increase the value of the `cfg.protocol.timeout` parameter. Genesys recommends values between `60` and `1800` seconds.

Warning! Do not change the `-Duser` JVM parameters specified in the `gim_etl_server` file. For example, do not change the time zone specified in `-Duser.timezone=GMT`. For example, do not change the time zone specified in `-Duser.timezone=GMT`; changing this parameter will result in incorrect population of the `DATE_TIME` calendar.

The following procedure provides the syntax for modifying the `gim_etl_server` file.

Procedure: Modifying the gim_etl_server file

Prerequisites

- Genesys Info Mart has been installed.

Start of procedure

1. Locate the `gim_etl_server` file for your operating system. After Genesys Info Mart has been installed, the `gim_etl_server` file is available in the installation directory.
2. Edit the file to add or change command-line parameters as required. JVM-related parameters use the prefix "-X" (for example, `-Xmx900m`); system-related parameters use the prefix "-D" (for example, `-DfilterUserData=false`).
   - On Windows, use the following syntax:
     ```
     set JVM_PARAMETERS=%JVM_PARAMETERS% <parameter=value>
     ```
     Use a separate line for each new system-related parameter.
   - On UNIX, use the following syntax:
     ```
     ${JAVACMD} <other command-line paramaters> <parameter=value>
     com.genesyslab.gim.etl.server.GIMServer "$@"
     ```
3. Save the file to the Genesys Info Mart installation directory.
4. Start or restart the Genesys Info Mart Server.

End of procedure

Starting Genesys Info Mart Server

You can start Genesys Info Mart Server in any of the following ways:

- From SCI (see page 362)
- Manually on UNIX (see page 362)
- Manually on Windows (see page 363)

Starting Genesys Info Mart Server with Solution Control Interface

Complete the following procedure to start Genesys Info Mart Server by using Solution Control Interface (SCI).

Procedure: Starting Genesys Info Mart Server by using SCI

Start of procedure

1. On the list pane in the SCI `Applications` view, select your Genesys Info Mart Server.
2. Do one of the following:
   - On the toolbar, click the `Start` button.
   - From the `Action` menu, select `Start`.
   - Right-click the `Application` object to access the shortcut menu, and then select `Start`.
3. In the confirmation box that appears, click `Yes`.

   SCI starts your Genesys Info Mart Server.

End of procedure

Starting Genesys Info Mart Server on UNIX

The following procedure describes how to start Genesys Info Mart Server manually on a UNIX system.

Procedure:
Starting Genesys Info Mart Server manually (UNIX)

Start of procedure

1. Open a console window.
2. Go to the directory in which you have installed Genesys Info Mart.
3. Enter the name of the Genesys Info Mart executable, followed by the appropriate command-line parameters. Use the following syntax:

   `./gim_etl_server -host <hostname> -port <portno> -app <application>`

   where:
   - *hostname* is the name of the host on which Configuration Server is running.
   - *portno* is the communication port that client applications must use to connect to Configuration Server.
   - *application* is the name of the Genesys Info Mart `Application` object, as defined for Configuration Server.

Note: If the host name or application name contains spaces or hyphens (-), enclose the name in double quotation marks.

For example, to start Genesys Info Mart Server by using command-line parameters that specify the host as `cs-host`, the port as `2020`, and the name as `Genesys_Info_Mart_81`, enter the following:

```
./gim_etl_server -host "cs-host" -port 2020 -app
"Genesys_Info_Mart_81"
```

End of procedure

Starting Genesys Info Mart Server on Windows

Use one of the following procedures to start Genesys Info Mart Server manually on Windows:

- Procedure: Starting Genesys Info Mart Server from the Start menu (Windows)
- Procedure: Starting Genesys Info Mart Server from a batch file (Windows), on page 364
- Procedure: Starting Genesys Info Mart Server as a Windows service

Note: On Microsoft Windows platforms, by default, the installation process installs Genesys Info Mart Server as a Windows Service.

Procedure:
Starting Genesys Info Mart Server from the Start menu (Windows)

Purpose: To start Genesys Info Mart Server manually from the `Start` menu.

Start of procedure

1. Go to `Start > Programs > Genesys Solutions > Genesys Info Mart (<Application Name>) > ETL Server`.
2. Select `ETL Server`.

   The Genesys Info Mart Server application starts.

End of procedure

Procedure:
Starting Genesys Info Mart Server from a batch file (Windows)

Purpose: To start Genesys Info Mart Server manually by using the batch file that is provided in the installation package (IP).

Start of procedure

1. In Windows Explorer, go to the directory in which you installed Genesys Info Mart.
2. Double-click the `startServer.bat` file.

End of procedure

Procedure:
Starting Genesys Info Mart Server as a Windows service

On Microsoft Windows platforms, by default, the installation process installs Genesys Info Mart Server as a Windows Service. If you stopped Genesys Info Mart Server from running as a Windows Service and need to start it again as a Windows Service, complete this procedure.

Start of procedure

1. Open the Windows `Control Panel`, and then double-click the `Services` icon. The `Services` dialog box opens.
2. In the `Services` list box, select your Genesys Info Mart Server service, and then click `Start`. (If you disabled Genesys Info Mart Server from operating as a Windows Service, the `Start` option for this application will not be available.)

Note: You can install the Local Control Agent (LCA) as a Windows Service with the user interface disabled. In this case, all servers that are started through SCI are started without a console, unless you specifically select the `Allow Service to Interact with Desktop` check box for both LCA and Genesys Info Mart Server.

End of procedure

Stopping Genesys Info Mart Server

You can stop Genesys Info Mart Server in any of the following ways:

- From SCI (see page 365). (This is the recommended method.)
- Manually on UNIX (see page 366).
- Manually on Windows (see page 367).

Note: To prevent Genesys Info Mart Server from self-starting, make sure that you clear the `autorestart` property in the Genesys Info Mart `Application` object in Configuration Manager.

Stopping Genesys Info Mart Server with Solution Control Interface

If you are using LCA and Solution Control Server (SCS), complete the following procedure to stop Genesys Info Mart by using SCI.

Procedure:
Stopping Genesys Info Mart Server by using SCI

Start of procedure

1. On the list pane in the SCI `Applications` view, select your Genesys Info Mart Server.
2. Do one of the following:
   - On the toolbar, click `Stop`.
   - From the `Action` menu, select `Stop`.
   - Right-click the `Application` object to access the shortcut menu, and then select `Stop`.
3. In the confirmation box that appears, click `Yes`.

   SCI stops your Genesys Info Mart Server.

End of procedure

Stopping Genesys Info Mart Server on UNIX

The following procedure describes two methods for stopping Genesys Info Mart Server manually on UNIX: from the command line and from a console window.

Note: If you are using LCA and SCS, you can also use SCI to stop Genesys Info Mart Server (see "Stopping Genesys Info Mart Server with Solution Control Interface").

Procedure:
Stopping Genesys Info Mart Server manually (UNIX)

Purpose: To stop Genesys Info Mart Server manually on a UNIX system from the command line or a console window.

Start of procedure

- Do one of the following:
  - On the command line, enter the following command:

    `kill -SIGTERM <processid>`

    where `<processid>` is the UNIX process ID of the application.
  - In the active console window, press `CTRL+C`.

End of procedure

Stopping Genesys Info Mart Server on Windows

Use one of the following procedures to stop Genesys Info Mart manually on Windows:

- If Genesys Info Mart Server is running as an application (not as a Windows Service), Procedure: Stopping Genesys Info Mart Server from the console window (Windows).
- If Genesys Info Mart Server is running as a Windows Service, Procedure: Stopping Genesys Info Mart Server running as a Windows service, on page 367.

Procedure: Stopping Genesys Info Mart Server from the console window (Windows)

Start of procedure

- In the application's console window, press `CTRL+C`.

End of procedure

Procedure: Stopping Genesys Info Mart Server running as a Windows service

If you are running Genesys Info Mart Server as a Windows Service, you should stop it only from the Services Control Manager, as described in this procedure.

Start of procedure

1. Open the `Control Panel`, and then double-click the `Services` icon. The `Services` dialog box opens.
2. In the `Services` list box, select your Genesys Info Mart Server service, and then click `Stop`.

End of procedure

Part

4 How Genesys Info Mart Works

Part 4 of this document provides more information for users who want an in-depth understanding about Genesys Info Mart functioning, particularly for users who want to modify Genesys Info Mart `Application` settings. This information appears in the following chapters:

- Chapter 18, “ETL Processing,” on page 371
- Chapter 19, “Maintenance and Other Activities,” on page 401
- Chapter 20, “High Availability,” on page 411

Chapter

18 ETL Processing

Both data quality in the Genesys Info Mart database and the performance of the Genesys Info Mart jobs that extract, transform, and load (ETL) data depend, to a large degree, on the amount and timeliness of new and late-arriving data that a given ETL cycle has to process. This chapter provides some details on extraction and transformation processing to explain ETL functionality.

This chapter contains the following sections:

Extraction Processing

If you plan to modify the default Genesys Info Mart application settings, review the information in this section to help answer the following questions during the planning process, to enable you to configure ETL jobs appropriately for your deployment:

- How much new data should be extracted and transformed in each cycle?
- What is the optimal transaction size?
- How long should ETL processing wait for missing data, so that reports reflect the most accurate data, but the jobs do not get stuck?

This section provides information about the following topics:

Extraction Roles

The algorithm that `Job_ExtractICON` uses to extract data depends on the extraction role that you configured in the database access point (DAP) that enables Genesys Info Mart to access IDB. The extraction roles are the following:

- `ICON_CFG`—To extract Configuration details from an IDB that contains data from Configuration Server
- `ICON_CORE`—To extract Voice details from an IDB that contains data from T-Server
- `ICON_OCS`—To extract Outbound Contact details from an IDB that contains data from Outbound Contact Server (OCS)
- `ICON_MM`—To extract Multimedia details from an IDB that contains data from Interaction Server

Notes:

- Extracting Voice details from IDB is optional, provided that you are extracting Multimedia (`ICON_MM`) details from an IDB. Similarly, 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 plan to extract both Voice and Multimedia details, you must configure different ICON applications to store ICON Voice and ICON Multimedia details into separate IDBs. This means that you cannot configure both the `ICON_CORE` and `ICON_MM` roles for the same DAP. You can configure multiple DAPs for the `ICON_CORE` or `ICON_MM` roles, to extract Voice or Multimedia details, respectively, from more than one IDB.
- Extracting Outbound Contact details from IDB is optional. If you do configure a DAP for `ICON_OCS`, the Genesys Info Mart Server will run `Job_ExtractICON` for the DAP with this role.
- If you plan to extract both Voice and Outbound Contact details, you must configure different ICON applications to store ICON Voice and ICON Outbound Contact details into separate IDBs. This means that you cannot configure both the `ICON_CORE` and `ICON_OCS` roles for the same DAP. You can configure multiple DAPs for the `ICON_OCS` role, to extract Outbound Contact details from more than one IDB.

- If you do not wish to extract a particular type of ICON details, do not configure a DAP for the applicable role.

For more information about DAP requirements and the topologies that Genesys Info Mart supports, see "Interaction Concentrator Topologies" on page 70.

For more information about how to configure the DAPs for data extraction, see "Configuring Required DAPs" on page 211.

For more information about how the extraction job extracts data for the various extraction roles, see "Extracting ICON Details" on page 377.

High-Water Marks

Genesys Info Mart 8.x extracts data by using the latest timestamp values as a *high-water mark,* which indicates the starting point for the next extraction. For example, if data was already extracted until 3:00 PM, the extraction job checks that the high-water mark (HWM) value from the previous cycle is 3:00 PM and extracts the next group of data that has timestamps that start from 3:00 PM. The timestamps are provided by the server that produces the data.

Transaction Size for Extract

To better manage database performance, data is extracted in chunks, which are defined by a configurable time interval. A chunk of data for each table (or group of tables) is committed in a separate transaction.

Extraction Windows

To align the data in time, the extraction job extracts—from all tables—all data that fits inside the same extraction window. The extraction window advances forward if there is available data and waits a configurable, limited time for stuck data. If data from active data sources is missing forever (for example, because a given ICON instance was stopped and then was uninstalled), the extraction job moves the HWM forward after the stuck threshold expires, and transformation proceeds with the available data.

To minimize the possibility of data loss, the extraction window does not advance if the previous run of the extraction job failed.

Figure 21 on page 374 illustrates the extraction window.

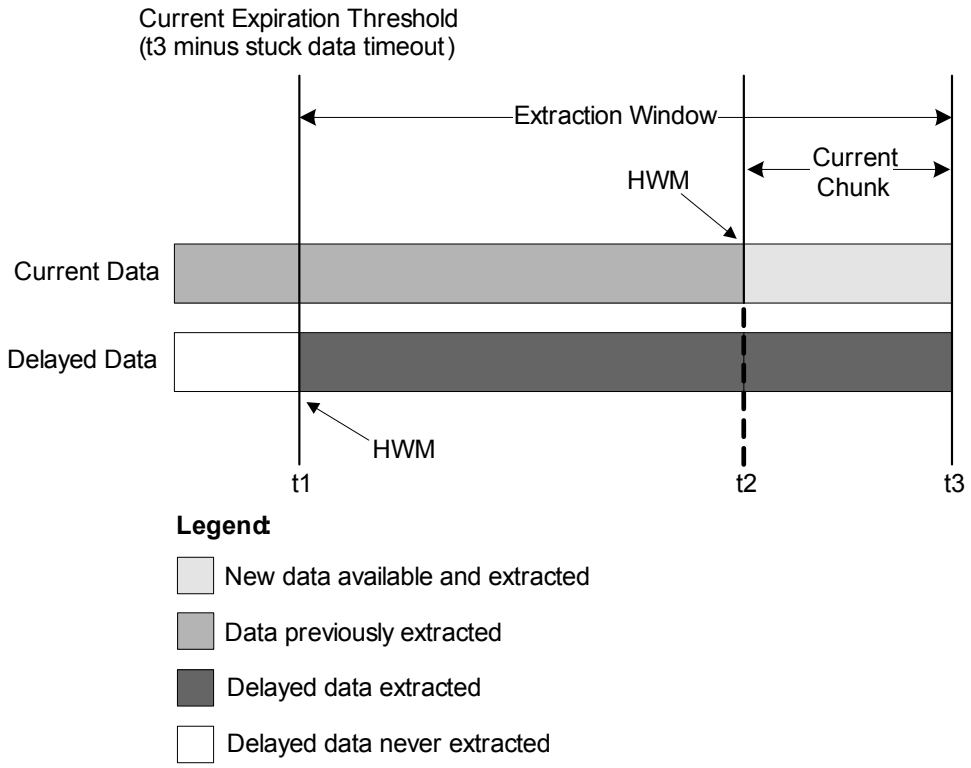

Figure 21: Extraction Window for a Data Domain

Extraction Window Example

For example, say that the configured chunk size is 15 minutes and the stuck data threshold is set to the default value of 28860 (8 hrs 1 min.). If the previous extraction cycle extracted data until 2:45 PM (`t2` in Figure 21), and if data is available beyond 3:00 PM (`t3`), the extraction window is from 6:59 AM (`t1`) to 3:00 PM in the current ETL cycle. If there is unextracted Interaction Database (IDB) data that has timestamps earlier than 6:59 AM, that data is not, and will never be, extracted. In this scenario, the Genesys Info Mart Server identifies that all data has been processed by the extraction job up to 6:59 AM; in other words, Genesys Info Mart adjusts all delayed extract HWMs (before `t1`) to 6:59 AM, and the transformation job proceeds to transform available data (see "Data Transformation" on page 388).

To look at the example in another way, if Genesys Info Mart extracted interactions that have timestamps up to 6:59 AM (`t1`) but some data is available up to 6:44 AM (`t0`) only, the ETL will wait for the delayed data until it has extracted data up to 3:00 PM (`t3`). If the delayed data (in other words, that has timestamps up to 6:59 AM) arrives in IDB after the leading edge of the extraction window has moved past 3:00 PM, the delayed data will never be extracted, because the extract HWM will have moved past the 6:59 AM timestamp.

When the extraction window advances, the extract HWM is set in the control tables. Subsequently changing the stuck data threshold does not undo the

HWM. Therefore, if data has been delayed past the stuck data threshold in an ETL cycle, it is lost to Genesys Info Mart.

Separate Extraction Windows for Each Data Domain

The extraction job manages extraction windows for different data sources (server types) separately. Thus, up to four extraction windows could coexist concurrently—one for data from each of the following server types: Configuration Server, T-Server, Interaction Server, and Outbound Contact Server (OCS).

However, the actual time ranges of the extraction windows and HWMs for each data domain might be different. To extend the preceding example, if new Voice details data is available but there is no new Outbound Contact details data, the extraction window and HWM for Voice details will advance as described in the example, but the extraction window and HWM for Outbound Contact details will not. So the concurrent extraction windows could be, say, from 6:59 AM to 3:00 PM for the Voice details data domain (T-Server data source) and from 6:44 AM to 2:45 PM for the Outbound Contact details data domain (OCS data source).

Controlling Data Extraction

The extraction window defines how much data to extract in one extraction cycle. Except for configuration *object* data, you can control the amount of data that is extracted in each ETL cycle. All available configuration object data is always extracted in one extraction cycle, as soon as it becomes available in IDB, so that it is available for the transformation of the other data. For all other types of data, configuration settings for the chunk size and stuck data waiting period affect how much data will be extracted during each ETL cycle.

A separate configuration option, `extract-data-cfg-facts-chunk-size` (see page 287) enables you to configure a larger extraction window for configuration *relationship* data than for other types of data, so that all available configuration relationship data can be extracted and transformed in one extraction cycle.

For Voice details, Multimedia details, and Outbound Contact details, the `extract-data-chunk-size` (see page 288) and `extract-data-stuck-threshold` (see page 289) configuration options determine the size of the extraction window for all three data domains. However, the actual time ranges for which data is being extracted and the HWMs for the different data domains might not be synchronized because of differences in data availability, as described above.

For more detailed information about how Genesys Info Mart extracts data for the various data domains, see “Extracting ICON Details” on page 377.

Data Quality vs. Data Availability

Genesys Info Mart extraction processing balances requirements for data quality with requirements for data availability. You can adjust the balance if you are willing to trade off the risks of loss of data against the risks of delayed reporting data or interrupted processing. For a summary of important configuration options that enable you to adjust ETL processing, see "Configuration Recommendations" on page 391.

Delayed Data for Extraction

Delayed data from an active data source does not delay extraction of data from other data sources in that domain. However, as described in "Delayed Data for Transformation" on page 389, missing data from one data source can delay transformation. Furthermore, particularly in a non-HA deployment, delayed data can result in data loss, as described in "Limitations" on page 377. It is important, therefore, to identify and fix extraction delays as soon as possible.

Starting with release 8.1.3, the `delayed-data-threshold` configuration option (see page 286) enables you to specify a timeout after which Genesys Info Mart generates a log event if expected data does not arrive. The log message (55-20110) includes information about the data sources and IDB tables from which data was expected. (In earlier releases, log message 55-20110 was generated only if related data was delaying a transformation chunk.)

Genesys recommends that you set an alarm on this message, so that you can investigate the reasons for data delays in a timely manner and take appropriate action. Depending on the reasons for the delay and your reporting requirements, appropriate action might include:

- Temporarily excluding an unavailable data source from extraction, so that transformation of data from other data sources in the same data domain will not be delayed. For more information about the steps to temporarily exclude a data source, see the section about recovering from data-source unavailability in the Troubleshooting chapter of the *Genesys Info Mart 8.1 Operations Guide*.
- Interrupting the job or temporarily stopping the ETL schedule, so that the next run of the extraction job does not advance the high-water mark (HWM) to a point that would result in permanent loss of the delayed data.

Concurrency

Data is extracted concurrently by using a configurable number of threads. In general, data from each table in a source database that is recorded from each server (T-Server, Interaction Server, and OCS) can be extracted concurrently. When the configured number of threads is `1` (one), all extraction tasks are executed sequentially. Setting the number of threads to a higher number can

help to improve performance of the extraction job; however, it requires more hardware resources: operational memory (RAM) and database transaction log.

Limitations

Genesys Info Mart does not support extraction of data that has an event timestamp that is earlier than the previous HWM. Therefore, the extraction job might skip data if:

- The data source HWM was not advanced by ICON. For example, this might occur if data was stuck in ICON's persistent queue.
- One data source was excluded from extraction, as described in the section about recovering from data-source unavailability in the Troubleshooting chapter of the *Genesys Info Mart 8.1 Operations Guide*. In this scenario, Genesys Info Mart continues to extract from active data sources for that data domain, and the extraction HWM continues to advance. By the time that data from the data source becomes available again, the expiration threshold for the extraction window might have advanced past some of the data, so that the data will never be extracted.

Extracting ICON Details

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

This subsection contains information about the following topics:

- Extracting Configuration Details, page 377
- Extracting Voice Interaction Data, page 378
- Extracting Multimedia Interaction Data, page 379
- Extracting Agent State, Agent State Reason, and DND Usage Data, page 380
- Extracting Agent Login Session Data, page 380
- Extracting Outbound Contact Details, page 380

Extracting Configuration Details

ICON Configuration details include data about:

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

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

Extracting Configuration Object Data

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

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

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

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

Extracting Configuration Object Relationship Data

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

The ETL extracts object relationship records from an available IDB based on comparing the extract HWM against the IDB timestamp of when the record was written to IDB. New relationship fact records are inserted into GIDB, and updates are simply merged into existing historical records. The ETL might create a particular GIDB table record based on data extracted from one IDB, and then might subsequently update the record based on data extracted from another IDB.

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

Extracting Voice Interaction Data

ICON Voice details include data about:

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

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

State Reason, and DND Usage Data" on page 380 and "Extracting Agent Login Session Data" on page 380.

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

The extract job initially extracts eligible call-related data to Merge tables, and it does not move the data to GIDB until it has run the merge procedure to establish the associations between related voice interactions. For more information about merge processing, see "Merge" on page 386.

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

In an HA deployment, when there are no disruptions, the ETL continues to extract each table independently from a reliable IDB that it used in the previous extract cycle. When disruptions occur, the ETL will switch to extracting from another IDB, if possible. For more information, see "Extracting Data in an HA Deployment" on page 381.

Extracting Multimedia Interaction Data

ICON Multimedia details include data about:

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

Agent-related data for multimedia activity is extracted in the same way as for voice activity. For more information, see "Extracting Agent State, Agent State Reason, and DND Usage Data" on page 380 and "Extracting Agent Login Session Data" on page 380.

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

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

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

Note: For a detailed discussion of Agent States, Agent State Reasons, and DND, see the section about populating agent activity in the chapter about populating Genesys Info Mart data in the *Genesys Info Mart 8.1 User's Guide*.

Agent state–related records are inserted once and never updated. These records do not contain artifacts that are created by ICON. Therefore, in an HA deployment, they do not have to be synchronized between HA ICON instances.

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

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

Extracting Agent Login Session Data

The configuration objects that are the subjects of agent login session records are uniquely identified by DBIDs. Agent login session data is processed similarly to configuration object relationship data (see page 378).

Extracting Outbound Contact Details

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

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

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

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

Extracting Data in an HA Deployment

Genesys Info Mart uses time-based switchover to extract all types of reporting data in an HA deployment. However, there are some differences in the extraction-processing details for different types of data and data flows. In particular, the fact that Genesys Info Mart processes voice interaction data only for completed calls, while ICON does not record data for voice calls that were already in progress when ICON established the connection with T-Server, has additional implications for HA data.

This subsection contains information about the following topics:

For more information about how the ETL extracts all types of data, see "Extracting ICON Details" on page 377.

General Data-Extraction Approach

For simplicity, the general description of HA data extraction in this subsection does not include consideration of geolocation. For more information about the effect of the additional consideration of geolocation on Genesys Info Mart extraction behavior, see "HA Data Extraction with Geolocation" on page 385.

To extract data from redundant IDBs when one of the IDBs experiences a disruption, Genesys Info Mart:

1. Determines the timestamps of data flow interruption and the timeframes of reliable data.
2. Extracts data from an IDB that it identifies as available and reliable, until the timestamp of the disruption.
3. Switches to another IDB that it identifies as available and reliable, and continues to extract all data from the other IDB.

HA Data-Extraction Example

The following example illustrates the overall approach to data extraction in an HA environment.

- An HA pair of ICONs (ICON-1 and ICON-2) populates an HA pair of IDBs (IDB-1 and IDB-2).

- Connection events occur at the following times:

Time	Event
t0	ICON-1 starts.
t1	ICON-2 starts.
t2	ICON-1 disconnects/terminates unexpectedly.
t3	ICON-1 reconnects/restarts.
t4	Both IDBs become unavailable to either ICON or Genesys Info Mart. They remain unavailable for a period that exceeds the configured stuck data timeout for the extraction window.
t5	Expiration threshold for the next extraction window (t6 – the stuck data timeout).
t6	IDB-1 reconnects/restarts.

  The intervals [t0–t1], [t2–t3], and from t6 onward represent a partial failure—data is not available from one of the IDBs but is available from the other.

  The interval [t4–t6] represents a total failure—data is not available from any IDB. In this scenario example, in which the failure occurs on the IDB side, data for the interval [t4–t6] might continue to accumulate in the ICON persistent queue during the outage and, subsequently, be stored in IDB-1.

Genesys Info Mart extracts data from the highly available ICONs and IDBs as follows:

- [t0–t2]: IDB-1
- [t2–t4]: IDB-2
- [t5–next disruption]: IDB-1

Data for the period [t4–t5] is never extracted.

Figure 22 on page 383 illustrates HA data extraction in this scenario.

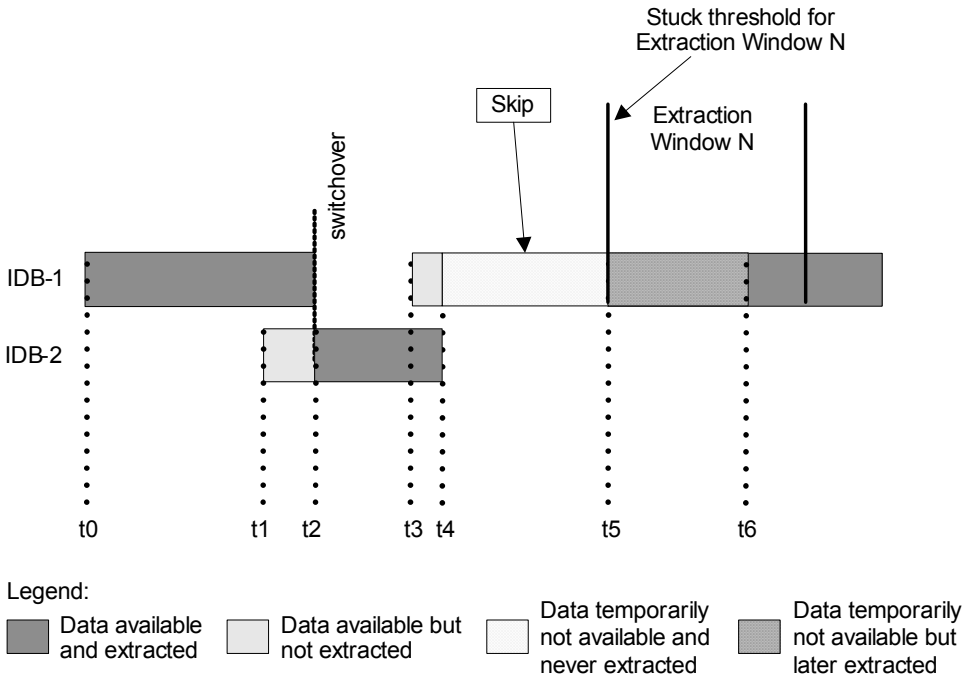

Figure 22: HA Data Extraction

Extracting Voice Interaction Data

The ETL processes only completed voice interactions. The extraction job initially extracts eligible interaction-related data to Merge tables, and it does not move the data to GIDB until related voice interactions have been merged.

The ETL extracts interaction data from a reliable IDB in the usual way, based on a comparison of the extraction HWM against the termination timestamp in the record. For reliable call data, the call must be visible to ICON for the entire call duration, from the time of call creation until call completion.

- The ETL extracts each table independently from a reliable IDB. Data from the `G_IR`, `G_CALL`, and `G_IS_LINK` tables is always extracted from the same IDB.

  When there are no disruptions:

  - If geolocation is a consideration, the ETL uses the local IDB. If Genesys Info Mart had switched to use a remote IDB because of a previous disruption, the ETL switches back to use the local IDB, as described in "HA Data Extraction with Geolocation" on page 385.
  - If geolocation is not a consideration, the ETL continues to use the reliable IDB that it used in the previous extraction cycle.

- When disruptions occur, the data that Genesys Info Mart extracts from the redundant IDB starts with the first call that has a termination timestamp later than the last completed call from the original IDB. Data about calls that started, but did not complete, before the disruption will be lost if there are no redundant IDBs that cover those calls entirely.

Figure 23 illustrates a partial failure scenario for HA extraction of voice interaction data. There is no data loss if the redundant ICON and IDB were operating before the disruption.

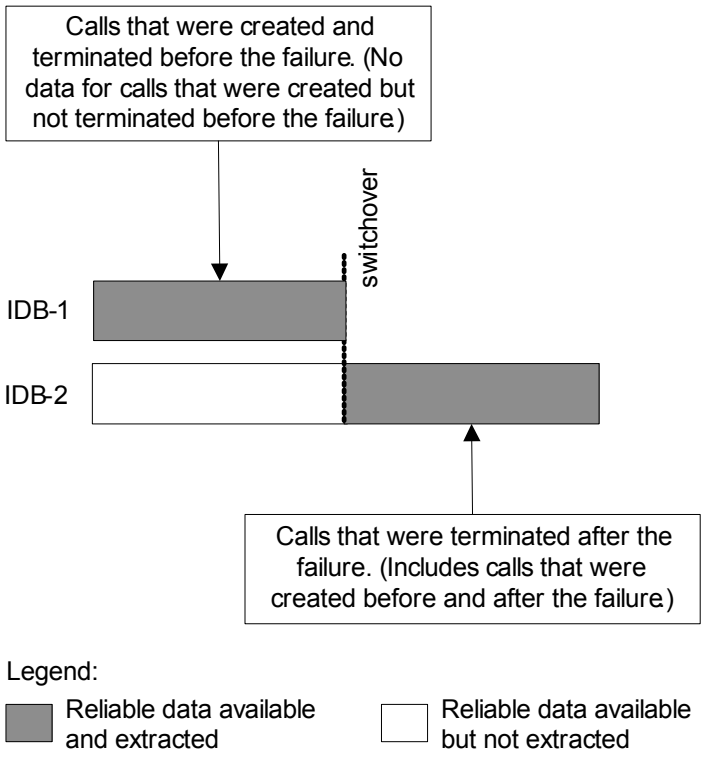

Figure 23: HA Data Extraction for Voice Calls—Partial Failure

Figure 24 on page 385 illustrates the total failure scenario for HA extraction of voice interaction data. Data will be lost for calls that end or start during the period when neither ICON is recording data.

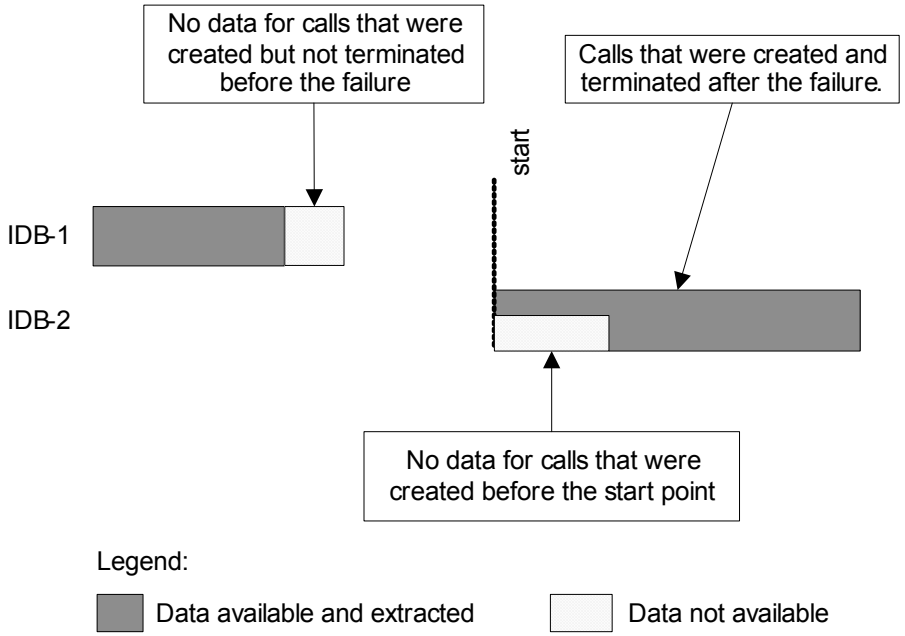

Figure 24: HA Data Extraction for Voice Calls—Total Failure

HA Data Extraction with Geolocation

In deployments in which the `geo-location` option has been configured in the DAPs that Genesys Info Mart uses to access the IDBs (extraction DAPs) and the Info Mart database (Info Mart DAP), Genesys Info Mart uses geolocation as a tie-breaker to select the best IDB from among equally reliable IDBs in an HA set. The `geo-location` option enables you to control whether the IDB is considered to be local to the Info Mart database or remote, regardless of where the hosts are physically located. For more information about configuring the `geo-location` option in order to identify local and remote IDBs, see pages 215 and 221.

The extraction job gives preference to the IDB that is considered to be local. Identifying preferred IDBs enables you to improve extraction performance by factoring in network management considerations (for example, minimize network traffic, adjust for or minimize network latency, and so on).

If the local IDB experiences a disruption, the ETL will switch over to use a remote IDB that has better data quality, if one is available, as described in "General Data-Extraction Approach" on page 381 and "Extracting Voice Interaction Data" on page 383. For more information about how Genesys Info Mart determines data quality, see "Criteria for Best IDB" on page 413.

Switching Back to the Local IDB

Following a disruption, when Genesys Info Mart determines that the local IDB is again providing reliable data, the ETL switches back to use the local IDB. For Voice details, Genesys Info Mart does not consider a new or restored

ICON session to be reliable until the `max-call-duration` interval has elapsed—in other words, until all voice calls that started after the disruption are monitored from start to end.

Figure 25 extends the previous example of the partial failure scenario for Voice details to illustrate how Genesys Info Mart switches back to use the local IDB after a disruption.

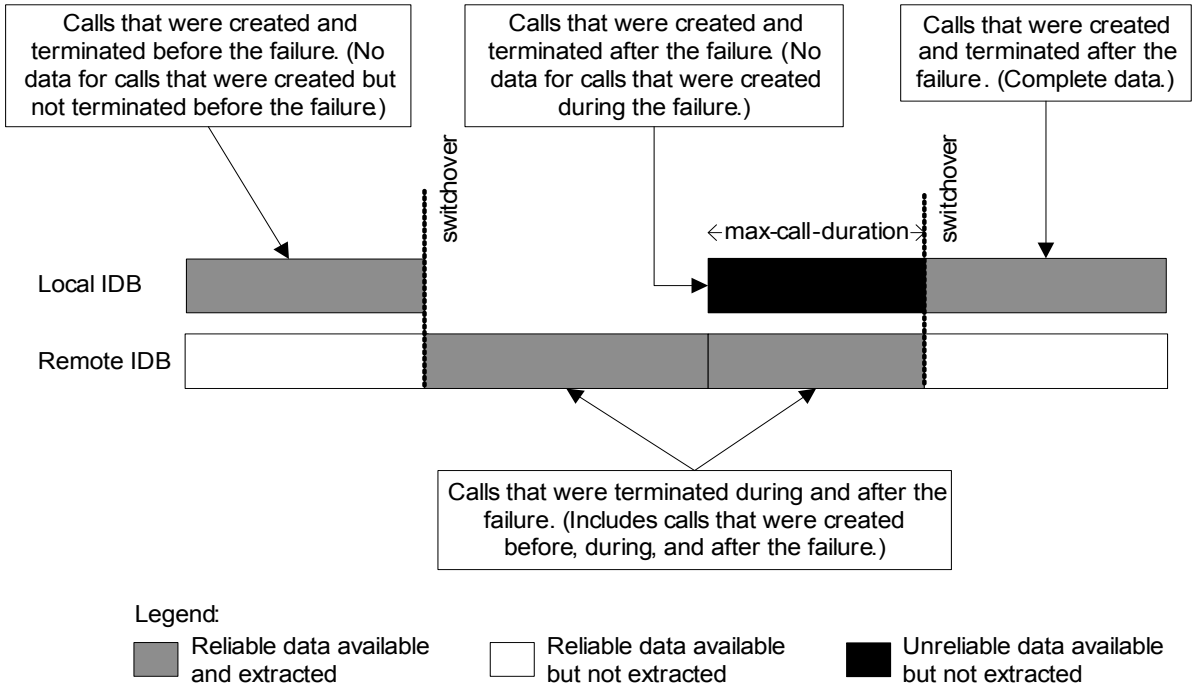

Figure 25: HA Data Extraction with Geolocation—Voice Details

Merge

The merge process is the process of merging related voice calls after extracting their data from one or more IDBs.

- In a single-site deployment, the merge establishes associations between voice calls that appear to be independent but are, in fact, part of the same voice interaction.
- In a multi-site deployment, the merge also processes interswitch voice calls—that is, the calls that involve parties from different monitored switches.

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

end-to-end interaction, and then Job_ExtractICON moves the successfully merged voice interactions to the GIDB tables.

Partially Merged Calls

If Genesys Info Mart does not have information about both sides of an IS-Link, the merge operation is able to perform only a partial merge. When the extraction job moves partially merged interactions into GIDB, the interaction records have unresolved linkages, which are referred to as *dangling links*.

The usual reasons that link information is missing are:

Unmonitored Sites

- One of the sites in a multi-site call is not monitored by ICON—In a multi-site environment in which ICON does not monitor one or more of the sites (T-Servers), data from any unmonitored site(s) will never arrive. To streamline processing, identify unmonitored sites in the GSYS_DNPREMOTELOCATION table in the Merge area, so that the merge operation does not wait for the stuck-link timeout to expire before the procedure finalizes the partial merge of interactions with links to the unmonitored site. For more information about how to configure unmonitored (remote) locations, see Procedure: Configuring Info Mart database for merge, on page 193.

Delayed or Missing Information

- All sites are monitored, but link information is delayed or missing—Link information from a monitored site might be delayed (for example, because of intermittent connectivity issues) or permanently missing (for example, because of extended outage of an ICON application during an extraction cycle). In these situations, the merge operation will wait for the missing data until the stuck-link timeout expires.
  - If the data arrives before the timeout expires (delayed data), the interaction is processed normally.
  - If the data does not arrive before the timeout expires (missing data), the ETL finalizes the partial merge, and the extraction job moves the partially merged interaction data to GIDB.

For information about how the transformation job handles dangling links, see “Transforming Partially Merged Calls” on page 390.

Merge Operation Stuck Threshold

The merge operation uses the max-call-duration configuration option (see page 290) to determine the stuck-link timeout. The timeout is determined in relation to the earliest HWM for all extracted data; the calculation is: min(HWM for all extracted data) – max-call-duration. If an unpaired G_IS_LINK record has a link-initiation timestamp that is earlier than the stuck-link threshold, the merge operation determines that the link is stuck and moves the record into GIDB.

The merge-failed-is-link-timeout configuration option (see page 297) enables you to configure a separate timeout for failed links, to provide a

sufficient buffer for Genesys Info Mart to receive both sides of a failed link, without introducing unnecessary data-quality issues that are related to dangling links.

Configuring for Merge

If you collect reporting data for voice interactions, merge-related application and database settings can affect Genesys Info Mart performance significantly. To streamline the functioning of the merge operation, you must identify unmonitored switches and ensure that merge-related configuration option settings are appropriate for your deployment. For more information, see Procedure: Configuring Info Mart database for merge, on page 193.

Data Transformation

`Job_TransformGIM` transforms GIDB data and then loads it into the Info Mart database as the last step of the transformation process.

This section describes important concepts that are related to `Job_TransformGIM` processing that you need to consider when you are configuring Genesys Info Mart.

Transaction Size in Transform

To make transaction size manageable, Genesys Info Mart 8.x transforms data in chunks. After a chunk of data is transformed, a new transaction is committed.

Transformation Chunk—Voice

A chunk (or transaction size) for voice interactions is defined as all data in a given primary (main) table that is marked by the same audit key, plus related data in a secondary (details) table that is marked with its own audit key. For example, `GIDB_G_IR_V` is the primary `GIDB` table for data that is transformed into `INTERACTION_RESOURCE_FACT` (IRF) or `MEDIATION_SEGMENT_FACT` (MSF) tables, while the secondary GIDB table in this case is `GIDB_G_PARTY_HISTORY_V`.

Transformation Chunk—Multimedia

For multimedia interactions, a chunk is based on minimum and maximum timestamps (`MinTS` and `MaxTS`)—`MaxTS` is taken from `CTL_AUDIT_LOG.MAX_CHUNK_TS` values associated with the audit keys for chunks that have been extracted into `GIDB_G_PARTY_HISTORY_MM`, and `MinTS` is the `MaxTS` from the previous chunk `+ 1` second. To determine which interactions will be transformed in the current chunk, the transformation job queries the `GIDB_G_PARTY_HISTORY_MM`, `GIDB_G_CALL_MM`, `GIDB_G_VIRTUAL_QUEUE_MM`, `GIDB_G_USERDATA_HISTORY_MM`, and `GIDB_G_SECURE_UD_HISTORY_MM` tables. If any record for an interaction was added to one or more of these tables in the time between `MinTs` and `MaxTs`, then that interaction is taken into the current chunk.

Starting with release 8.1.4, the minimum and maximum timestamps of the records that are taken into the current chunk are then used to query all of the

GIDB_*_MM tables, and previously transformed GIDB data for existing active interactions is taken from STG tables. In 8.1 releases earlier than 8.1.4, the transformation job queried the GIDB_*_MM from the beginning of the oldest interaction that was taken into the current chunk and retransformed all the GIDB data for existing active interactions.

Transaction Size for Transformation and Extraction

The transformation job uses the audit key that the extraction job assigned to GIDB data during the extraction. This means that the transaction size for the transformation job depends on the transaction size of the preceding extraction job within the same cycle. Using the same configuration parameters, you control the transaction size for both jobs. Therefore, when you set the configuration options to control the volume of data that is extracted (see "Controlling Data Extraction" on page 375), consider also the transaction size for transformation transactions.

Horizontal Transform

Horizontal transformation is tied closely to chunking. The data from the primary table (such as GIDB_G_IR_* in the previous example) is transformed only when the data in the secondary table(s) is available (GIDB_G_PARTY_HISTORY_* in the previous example). In other words, a dependency between the primary and secondary tables is built into the transformation logic. Before it transforms a chunk of data, Job_TransformGIM checks that all related data has been extracted at least up to the maximum timestamp in the chunk. If data arrival in a secondary IDB table is delayed, a current data chunk in the primary table is not processed either.

Under normal conditions, the transformation job does not get stuck, because the data that is currently missing would be extracted during a subsequent ETL cycle. It is important that you configure a wait time for stuck data (the extract-data-stuck-threshold configuration option) that allows enough time for expected or acceptable delays before data appears in the source database. The default stuck threshold is 8 hrs 1 min.

Each time that the transformation job detects references to unavailable data, it generates a log message.

Delayed Data for Transformation

If data from one data source is delayed, transformation of data from all data sources for that data domain is delayed while the transformation job waits for the missing data.

A configuration option in the [gim-etl] section, delayed-data-threshold (see page 286), enables you to specify a timeout after which Genesys Info Mart generates log message 55-20110. To prevent excessive or unnecessary delays if expected data does not arrive, you can set an alarm on this message, to alert you to take action in a timely fashion.

Late- and Early-Arriving Virtual Queue Data

Because the T-Server or Interaction Server that handles virtual-queue data might not be the same T-Server or Interaction Server that handles the associated interaction, the virtual queue and interaction data streams are not necessarily synchronized. Virtual-queue data might be available in GIDB at the same time as the interaction data, before the interaction data arrives, or after some delay.

Late- and early-arriving data requires additional processing and causes additional updates to fact tables in the dimensional model. The reasonably intensive processing has a performance impact.

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

Transforming Partially Merged Calls

As described in "Partially Merged Calls" on page 387, there are two situations that result in partially merged voice interactions:

- One of the sites is registered as unmonitored—The transformation job ignores dangling links that point to unmonitored sites and transforms the interaction as if it were an inbound or outbound call to a remote site. For example, if the dangling link provides information about the target for an internal transfer, the interaction will be transformed as an inbound call, instead of as a transfer.
- Information is missing—The transformation job treats dangling IS-Links as a data inconsistency, which it handles as described in "Data Inconsistencies During Transformation" on page 397. The interaction-level error-policy option that controls behavior when the transformation job encounters this data inconsistency is `error-policy-islink-dangling` (see page 278).

Partially merged calls have implications for data quality. For more information, see the section about data issues in a partially monitored environment in the *Genesys Info Mart 8.1 User's Guide*.

Transforming Calls in Parallel Queues

If a call is handled by a mediation resource while the call is in a queue, `Job_TransformGIM` identifies the fact that the same interaction is being processed by two resources at once and combines the records for that interaction into a single IRF.

For example, a call arrives at a Routing Point, which diverts the call to an ACD queue. While the call is queued at the ACD, the ACD interacts with an IVR, which collects an extension number from the customer and returns the collected digits to the ACD, after which the ACD diverts the call to the extension number that was entered by the caller.

For a graphic that shows a call flow of this scenario, see the chapter about validated voice interaction flows in the *Genesys Info Mart 8.1 User's Guide*.

No Retransform

If the transformation job completes successfully, Genesys Info Mart advances the transformation HWM to the next chunk of data. Genesys Info Mart does not support retransformation of any previous chunk of data for voice interactions or, starting with release 8.1.4, multimedia interactions. (In releases earlier than 8.1.4, Genesys Info Mart did retransform data for old multimedia interactions in GIDB.)

Multimedia Data Transformation

Genesys Info Mart provides integrated processing of voice and multimedia interactions. However, because of differences in the characteristics of voice and multimedia interactions, there are additional special requirements for the internal processing of multimedia interactions. For more information about considerations for multimedia interaction processing, including their implications for configuration and management of your Genesys Info Mart deployment, see Chapter 5 on page 97.

Configuration Recommendations

The following configuration options control important aspects of ETL operation:

- `extract-data-chunk-size` controls the size, in seconds, of the time interval for which the data is committed in one transaction. This option affects `Job_ExtractICON` and `Job_TransformGIM`. For more information about this option, see page 288.
- `extract-data-cfg-facts-chunk-size` controls the size, in seconds, of the time interval for which configuration relationship data is committed in one transaction. This option affects `Job_ExtractICON` and `Job_TransformGIM`. For more information about this option, see page 287.
- `extract-data-stuck-threshold` specifies the time, in seconds, that Genesys Info Mart waits for delayed data from active data sources. This option affects `Job_ExtractICON` and `Job_TransformGIM`. For more information about this option, see page 289. For the meaning of *active data source,* see "Active Data Sources" on page 40.

- `extract-data-thread-pool-size` sets the maximum number of threads that will be used to extract data concurrently. For more information about this option, see page 289.
- `extract-data-max-conn` sets the maximum number of connections that can be used to extract data concurrently from each database access point (DAP) through which Genesys Info Mart accesses IDB. For more information about this option, see page 288.
- `max-call-duration` specifies the time, in seconds, that:
  - The merge operation waits for delayed IS-Link information during extraction processing.
  - The transformation job waits for After Call Work (ACW) information.
  - The transformation job performs party activity lookup for multimedia interactions.

  For more information about this option, see page 290.

  See also related information about the `merge-failed-is-link-timeout` option (see page 297), which controls the merge-processing timeout for failed IS-Links.
- `merge-chunk-size` controls the size, in rows, of a chunk of merged data for transformation. For more information about this option, see page 296.
- `delayed-data-threshold` controls the amount of time, in seconds, after which Genesys Info Mart generates a log event to notify you that expected data is delayed. You can set an alarm on this event. For more information about this option, see page 286.

When you select values for these options, keep in mind the following recommendations:

- Chunk size should be small enough to process database transactions quickly, yet large enough to avoid the overhead that is caused by almost empty chunks. Note that the number and frequency of transactions affect database performance.

  For example, if you have one month's worth of data in IDBs in your lab environment, but data density is low (under 10,000 calls per day), it makes sense to increase `extract-data-chunk-size` temporarily to, say, two hours (7200 seconds). By contrast with the default chunk size of 900 seconds (15 minutes), this new setting would enable Genesys Info Mart to speed up the processing of the backlog of IDB data.

  On the other hand, setting the chunk size excessively high (for example, 90000 seconds [25 hours]) in an attempt to speed up backlog processing causes transformation problems. Genesys does not recommend setting the chunk size to a value greater than 7200 seconds.
- Transformation chunk size is defined by the extraction chunk size. Therefore, consider the size of the transform transaction when you select the value for `extract-data-chunk-size`.

For Voice details, the supplementary `merge-chunk-size` configuration option enables you to limit the chunk size for merged data, to assist you in managing situations in which the Merge tables grow excessively large (for example, if disruptions during extraction processing result in the Merge tables containing more than one extraction chunk's worth of unmerged data).

- The waiting period for stuck data should be large enough to ensure adequate data quality, yet small enough so as not to cause the jobs to wait for missing data longer than is acceptable for adequate performance.
- The recommended setting range for `extract-data-thread-pool-size` is between 16 and 32 threads.

You must also specify in the configuration how often the ETL cycle (extraction and transformation) runs. Make sure that the transaction size for the extraction job (`extract-data-chunk-size`) is greater than or equal to the interval that is configured between extraction cycles.

Error Handling

The main categories of error that Genesys Info Mart might encounter are the following:

- Invalid Configuration, page 393
- Unavailable Data, page 393
- Missing Configuration Data, page 396
- Data Inconsistencies During Transformation, page 397

Invalid Configuration

Genesys Info Mart checks the configuration of Genesys Info Mart, Interaction Concentrator, and the database environment. For more information about Genesys Info Mart behavior with regard to required configuration settings, see "Deployment Verification" on page 401.

Unavailable Data

In the path from the data source through ICON to Genesys Info Mart, there are a number of points at which data processing can be interrupted.

At the start of each extraction job (as described in "Data-Source Availability" on page 403), Genesys Info Mart compares a list of active data sources against a list of available data sources. Any discrepancy between the lists enables Genesys Info Mart to infer that a failure has occurred.

During extraction processing, Genesys Info Mart uses information in the `G_DSS_*_PROVIDER` tables in IDB to monitor activity on each data-source session. In particular, Genesys Info Mart compares last event and `NoData`

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

During transformation processing, if expected data does not arrive within a configurable time interval, Genesys Info Mart generates a log message to alert you about the delay from a particular data source. For more information, see "Delayed Data for Transformation" on page 389.

The following subsection provides more information about the scenarios in which delays might occur and the implications of delays for data quality.

Failure Scenarios and Recommendations

Figure 26 on page 394 shows three points of failure that can arise from either component or network outages. Each point is discussed in the correspondingly numbered paragraph that follows Figure 26.

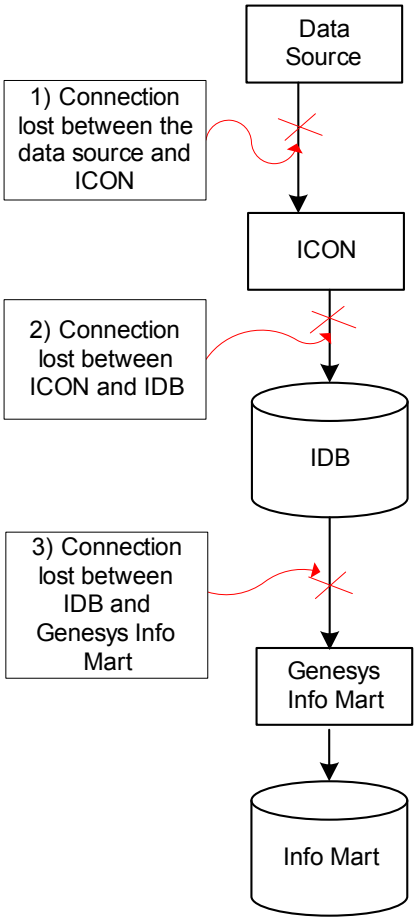

Figure 26: Possible Failure Points

1) Connection failure between ICON and the data source

If this connection fails, data transmission is interrupted. When the connection is restored, ICON writes whatever data the data source can provide to IDB, from which Genesys Info Mart then extracts it.

Recommendation: To prevent the loss of data from a connection failure between the data source and ICON, set up your environment to use HA at the data-source level.

In addition, Genesys recommends that you configure an alert for log message 55-20110, so that you can investigate delays promptly and take appropriate action, as described in "Delayed Data for Transformation" on page 389.

2) Communication failure between ICON and IDB

When the interruption to communication occurs between an ICON and the IDB that it populates, Genesys Info Mart continues extracting data from any other available IDBs for the same data domain. The result depends on two factors:

- Whether the data interruption lasts long enough that the extraction window passes entirely beyond the point at which the data interruption occurred (see Figure 27)
- Whether your Genesys Info Mart settings can accommodate backlog extraction

Figure 27 illustrates the scenario in which the extraction window passes the point at which the interruption occurred.

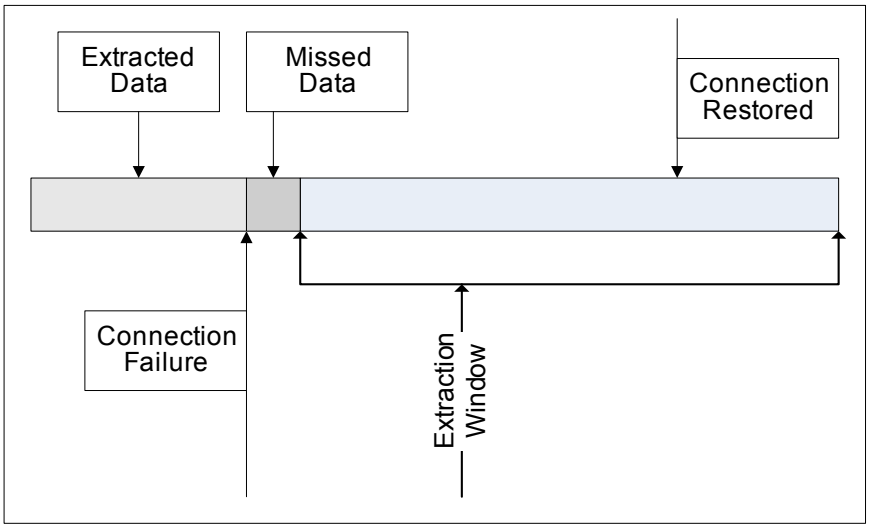

Figure 27: Data Extraction Window

By adjusting the extraction chunk size or even halting purging, you can ensure that backlogged data is extracted and transformed correctly. The outcome is delay, but no data loss.

If the configuration options that control extraction chunk size and ETL scheduling do not accommodate efficient processing of a backlog of ICON data, you might fail to extract some ICON data if the communication failure lasts long enough for the extraction window to pass entirely beyond the time at which the communication failure initially occurred. Alternatively, if the configuration options that control data retention in GIDB do not allow for a backlog of unprocessed data, you might fail to transform some extracted data before it is purged.

Recommendation: To prevent the loss of data from a connection failure between ICON and IDB, set up your environment to use HA at the ICON level.

3) Communication failure between Genesys Info Mart and IDB

There are a number of outcomes in this scenario, depending on the nature of the connection interruption.

- Genesys Info Mart might not be able to access the data in IDB and, therefore, cannot read the session information from IDB. In this case, the extraction job fails.
- Genesys Info Mart tries to match the list of configured data sources against the list of data sources for the available IDB(s). If it finds no mismatch, Genesys Info Mart:
  - Logs an event that indicates the DAP connection that is missing.
  - Proceeds with extraction.

Note: If Genesys Info Mart has no connection to the Info Mart database, it will not start the extraction job at all.

Missing Configuration Data

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

- For Configuration details themselves (configuration object records or configuration relationship records), missing configuration objects are an error that indicates a problem in the source data. For example, the IDB schema has no constraint against foreign-key violations; therefore, a configuration relationship record might reference a configuration object for which there is no record. (For example, if an agent was created before ICON started, and then the agent was added to a group and subsequently deleted, there will be a record in the `GCX_GROUP_AGENT` table but no corresponding record in the `GC_AGENT` table.)

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

- When Genesys Info Mart transforms other kinds of fact data, the transformation job treats missing configuration objects as delayed dimension data. In other words, the transformation job does not yet have sufficient information to identify the configuration object, but it expects to receive that information later. In the meantime, the transformation job has sufficient information (object type and unique configuration ID) to enable it to create an incomplete record as a placeholder for the missing object, so that transformation can continue. If data about the missing configuration objects subsequently arrives from Configuration Server, the incomplete records are populated with full information.

Missing Virtual Queue Resource Example

For example, starting with Universal Routing Server (URS) release 8.0, URS provides the database identifier that Configuration Server assigns to the virtual queue (VQ DBID) in `EventRouteUsed` for the `Virtual Queue`, and ICON stores the DBID. If an interaction is routed through a virtual queue for which Genesys Info Mart has not yet created a Resource record, Genesys Info Mart uses the available data about VQ DBID to create a placeholder `VirtualQueue` record in the `RESOURCE_` table, and the MSF record references that Resource. When full information arrives from Configuration Server, Genesys Info Mart updates the `VirtualQueue` record in the `RESOURCE_` table.

Data Inconsistencies During Transformation

Genesys Info Mart does not fully validate data consistency during processing. However, the transformation job will report data inconsistencies if the transformation algorithm encounters them.

Genesys Info Mart has a two-level approach to handling errors during transformation:

- Interaction level—Error-policy options enable you to configure the behavior of the transformation job when it encounters various types of data inconsistency during processing (for example, missing IS-Link information, missing party creation data, or duplicated party creation data). The options enable you to specify different treatment for different kinds of errors.

  If the option that specifies an error-handling policy at the interaction level is set to:

  - Interrupt processing with an exception, for all interaction-level error policies except the policy controlled by `error-policy-campaign-group-missing`, the exception is handled as specified by the error-policy option for the job. For `error-policy-campaign-group-missing`, an exception causes `Job_TransformGIM` to fail.
  - Resume processing without generating an exception, `Job_TransformGIM` processes the data as described for the applicable interaction-level option in “error-policy Section” on page 276. The type of error is recorded in the `STATUS` field of the `INTERACTION_FACT` record.

- Job level—An error-policy option, `error-policy-irf-exception` (see page 277), controls the behavior of the transformation job when it encounters an exception. A configurable filter (see `error-policy-irf-exception-resumable` on page 277) enables you to fine-tune this behavior.

The default setting of almost all the interaction-level error-policy options is `resume`, to favor uninterrupted operation and data inclusiveness over data consistency.

Table 24 summarizes the results of various combinations of the interaction-level and job-level error policies. For a summary list of the available error-policy options, see the "Transformation error handling"section in Table 15 on page 248. For full details about the available error-policy options, see "error-policy Section" on page 276.

Table 24: Error Policy Reporting Results Matrix

Interaction-Level Error Policy	Job-Level Error Policy	Reporting Result
`resume` (default for most options)	• `error-policy-irf-exception`= Any valid value • `error-policy-irf-exception-resumable`= Any valid Java expression	Interaction is transformed with a possible impact on data quality. Information about the error is reported in the `STATUS` field of the `INTERACTION_FACT` record.
For all except `error-policy-campaign-group-missing`: `exception`	• `error-policy-irf-exception=log_db_resume (default)` • `error-policy-irf-exception-resumable`= A value that matches the actual exception class	Interaction is discarded, and corresponding information is written to the `STG_TRANSFORM_DISCARDS` table.
	• `error-policy-irf-exception=resume` • `error-policy-irf-exception-resumable`= A value that matches the actual exception class	Interaction is discarded, and no information is written to the database.
	• `error-policy-irf-exception=log_db_resume (default)` • `error-policy-irf-exception-resumable`= A value that does not match the actual exception class	The transformation job fails.
	• `error-policy-irf-exception=resume` • `error-policy-irf-exception-resumable`= A value that does not match the actual exception class	The transformation job fails.
	`exception`	The transformation job fails.

Table 24: Error Policy Reporting Results Matrix (Continued)

Interaction-Level Error Policy	Job-Level Error Policy	Reporting Result
For `error-policy-campaign-group-missing: exception`	Any value	The transformation job fails.

Chapter

19 Maintenance and Other Activities

This chapter provides information about aspects of Genesys Info Mart functioning other than direct extract, transform, and load (ETL) processing. It contains the following sections:

Deployment Verification

The Genesys Info Mart Server checks the integrity of the Genesys Info Mart deployment. Genesys Info Mart automatically performs this check on startup; whenever the configuration of the Genesys Info Mart `Application` changes (option settings or connections); and whenever the configuration of any ICON `Applications` that are in the Genesys Info Mart `Application` connections changes. Genesys Info Mart also performs a data availability check at the start of each extraction cycle.

Genesys Info Mart considers only those Interaction Concentrator applications (ICONs) and database access points (DAPs) that are in the Genesys Info Mart `Application` connections to be part of the deployment. In addition, starting with release 8.1.2, the ICONs and the extraction DAPs (the DAPs that Genesys Info Mart uses to access the associated Interaction Databases [IDBs]) must be enabled in order for Genesys Info Mart to consider them part of the deployment.

Configuration Check

The Genesys Info Mart Server verifies the following items in the deployment. If it finds any errors, the Genesys Info Mart Server logs an error and does not start any new jobs until the error is corrected.

Genesys Info Mart Configuration

- The version of the Genesys Info Mart database schema is correct.

  For information about how to create the Genesys Info Mart database schema, see "Preparing the Info Mart Database" on page 191.

- Values for Genesys Info Mart `Application` options are valid.

  Starting with release 8.1.0, Genesys Info Mart verifies that the ETL frequency, as specified by the `etl-frequency` option in the `[schedule]` section (see page 316), is not greater than the chunk size for each extraction cycle, as specified by the `extract-data-chunk-size` option in the `[gim-etl]` section (see page 288). This check ensures that Genesys Info Mart is set up to keep pace with ICON.

  Starting with release 8.1.2, Genesys Info Mart verifies that the `days-to-keep-gim-facts` option (see page 284) is greater than `days-to-keep-gidb-facts` option and `days-to-keep-active-facts`.

  Genesys Info Mart does not perform any further checks to validate internal consistency.

  For information about the Genesys Info Mart configuration options, see Chapter 14 on page 271.

ICON Configuration

- Mandatory ICON `Application` options are present and their values are correct. In the `callconcentrator` section:
  - `use-dss-monitor=true`
  - `calls-in-the-past=true`
  - `om-force-adata=true`
  - `gls-active-reason-codes=true`
  - `partition-type=2`
  - The `gcc`, `gud`, and `gls` values for the role option always occur together.

  Any changes in the ICON `Application` configuration options or connections trigger the configuration check, but Genesys Info Mart checks only that these mandatory options have been set as required.

  For more information about ICON configuration options that affect Genesys Info Mart functioning, see "Preparing ICON" on page 162.

Genesys Info Mart Configuration on Related Objects

- If Genesys Info Mart–related configuration options are present on `DN`, `Switch`, or `Media Type Business Attribute` configuration objects, the options have valid values.

  For more information about Genesys Info Mart–related configuration options that can be set on related configuration objects, see Chapter 13 on page 253.

General Deployment

- The correct version of the JDBC driver is available.

  Genesys Info Mart verifies the version of the JDBC driver against the version that is specified in the `gim_cfg_checkup.properties` file. By default, the specified driver is `oracle.jdbc.driver.OracleDriver`. For information about how to specify an alternative JDBC driver, see "Modifying the Configuration Checkup Properties File" on page 329.
- There is at least one ICON in the Genesys Info Mart `Application` connections.
- There is one ICON (or one HA set) that provides Configuration details in the Genesys Info Mart `Application` connections.
- Voice details and Multimedia details are in separate IDBs. Prior to release 8.1.301.07, Genesys Info Mart also required Outbound Contact details to be in a separate IDB.
- Based on the roles that are configured in the DAPs for extracting data, Genesys Info Mart is set up to extract Voice details, Multimedia details, and Outbound Contact details from separate IDBs.
- The configuration includes an Info Mart DAP.

IDB Schema Compatibility

Starting with release 8.1.2, at the start of each extraction job, Genesys Info Mart checks each IDB for which Genesys Info Mart has an extraction DAP, to identify whether the IDB schema needs to be modified for Genesys Info Mart. If any IDB does not have the correct version, Genesys Info Mart enters the migration state and will not perform any other activity until you run the migration job, which automatically executes the required scripts to update IDB. You must manually run `Job_MigrateGIM` from the management GUI.

For more information about the update scripts, see "Preparing IDBs" on page 178.

Data-Source Availability

At the start of each extraction job, Genesys Info Mart verifies the availability of configured data sources.

Genesys Info Mart uses session information from the `G_DSS_*_PROVIDER` tables in the IDBs to identify all available data sources. If the same data source is present in session information in more than one IDB, high availability (HA) mode is in effect automatically.

Genesys Info Mart also uses information about the data sources in ICON connections to verify the availability of *active data sources* (as defined on page 40).

If required, Genesys Info Mart performs maintenance on high-water marks (HWMs). For example, adjustment of HWMs might be necessary if a data

source that was previously in the deployment is removed from ICON connections.

Success Criteria for the Extraction Job

The extraction job proceeds if data from all active data sources is available in at least one IDB. Starting with release 8.1.3, the extraction job logs an error message if there is no data from an active data source for more than the time specified by the `delayed-data-threshold` option (see page 286).

The extraction job succeeds if Genesys Info Mart is able to extract all data that was planned for a particular extraction cycle.

Preventing Extraction Delays

If you do not want an unavailable data source to delay extraction processing, consider temporarily excluding the data source from extraction. For more information, see "Delayed Data for Extraction" on page 376.

Maintenance

Genesys Info Mart provides a maintenance job, `Job_MaintainGIM`, to maintain the Info Mart database and calendar data. You can schedule the maintenance job to run on a regular basis, or you can run the job manually from the Genesys Info Mart management GUI, as required.

The main functions of the maintenance job are:

- Maintaining database partitions in partitioned databases
- Maintaining calendars (see page 405)
- Purging the Info Mart database (see page 405)

Additional Maintenance for PostgreSQL

For PostgreSQL deployments, Genesys Info Mart provides a supplementary maintenance job, `Job_UpdateStats`, which updates out-of-date statistics and runs a vacuum process to reclaim storage space from updated or deleted rows. You must schedule `Job_UpdateStats` to run regularly, outside the window for `Job_MaintainGIM`. You cannot run the job manually from the Genesys Info Mart management GUI. For information about scheduling `Job_UpdateStats`, see the option descriptions for `run-update-stats` (page 318) and `update-stats-schedule` (page 319),

Maintaining Database Partitions

In deployments that use a partitioned Info Mart database, `Job_MaintainGIM` creates partitions ahead of the ETL cycle, to provide storage for GIDB, Control, and dimensional model fact data in upcoming ETL cycles.

The size of the partitions, as well as the number of days ahead for which the maintenance job creates partitions, are configurable. For more information, see the descriptions of the `partitioning-interval-size-*` and `partitioning-ahead-range` options, starting on page 298.

If your database is partitioned, Genesys strongly recommends that you configure a schedule for the maintenance job. For more information, see "Scheduling the Maintenance Job" on page 409.

Maintaining Calendars

`Job_MaintainGIM` performs calendar maintenance by populating the calendar table(s) for future reports. The *calendar tables* are the default `DATE_TIME` dimension table and any custom calendar tables that you create to support your reporting.

`Job_InitializeGIM` initially populates the calendar table(s) as far ahead as you specify (see the `date-time-max-days-ahead` option, on page 272), so that calendar dimensions are available for your reports. `Job_MaintainGIM` continues to populate the calendar tables when the next batch of calendars is required (as specified by the `date-time-min-days-ahead` option, on page 272).

For more information about how to create custom calendars, see Procedure: Configuring custom calendars, on page 347. For information about how to modify existing calendars, see the procedure about changing calendar dimensions in the *Genesys Info Mart 8.1 Operations Guide*.

Purging Info Mart Data

Note: Purge functionality with respect to active interactions changed significantly between Genesys Info Mart release 8.1.0 and release 8.1.1. Information in the following section applies to behavior starting with release 8.1.1. For the benefit of Genesys Info Mart 8.1.0 customers, Appendix E on page 451 provides information about purge functionality in previous 8.x releases.

`Job_MaintainGIM` purges:

- Completed and artificially terminated fact data from GIDB, in accordance with data-retention policies that are specified by the `days-to-keep-gidb-facts` and `days-to-keep-active-facts` options (see pages 283 and 281).
- Completed and artificially terminated fact data from the dimensional model, in accordance with data-retention policies that are specified by the `days-to-keep-gim-facts` and `days-to-keep-active-facts` options (see pages 284 and 281).
- Starting with release 8.1.4, deleted records from the `GROUP_ANNEX` and `RESOURCE_ANNEX` dimension tables, in accordance with a data-retention policy specified by the `days-to-keep-deleted-annex` option (see page 282).

- Discarded operational data from discard tables, as well as data lineage and ETL job history data from various Control tables, in accordance with a data-retention policy that is specified by the `days-to-keep-discards-and-job-history` option (see page 283).

Notes:

- `Job_MaintainGIM` does *not* purge old aggregate data, dimension data, or configuration data (except for the `*_ANNEX` tables in release 8.1.4 and later).
- Genesys Info Mart artificially terminates any active interactions before it purges them.

Retention Periods

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

For the configuration options that determine the retention periods, the values that you choose must allow the Info Mart database to retain the data long enough for you to complete deployment-specific tasks (such as 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 retention period for Info Mart fact data to 380 days. This allows for the extra day in a leap year, plus two extra weeks for deployment-specific tasks.

Retention Periods for Facts

The subjects of fact records are entities such as interactions, agent login sessions, or Outbound Contact campaigns.

`Job_MaintainGIM` purges data based on the start time of facts in the applicable tables. Depending on configuration, the `days-to-keep-gim-facts` and `days-to-keep-gidb-facts` options either directly specify the retention period or else work in conjunction with the `days-to-keep-active-facts` option to determine the actual retention period for a particular situation.

Figure 28 illustrates the retention periods for GIDB and the dimensional model, respectively, for the situation in which `days-to-keep-active-facts` is smaller than `days-to-keep-gidb-facts` or `days-to-keep-gim-facts`.

The starting point from which Genesys Info Mart calculates the retention periods (in other words, the time from which Genesys Info Mart starts counting backward) changed between releases 8.1.1 and 8.1.2. In release 8.1.1, retention periods were calculated from the extraction HWM. Starting with

release 8.1.2, the retention period is usually calculated from the transformation HWM, with one exception: In the case of inactive data domains (for example, if your deployment was reporting on Voice and Multimedia activity and you decide to suspend collecting Multimedia data), the retention period is calculated from the day that `Job_MaintainGIM` is running ("today"). For simplicity, the labels in Figure 28 use "Today."

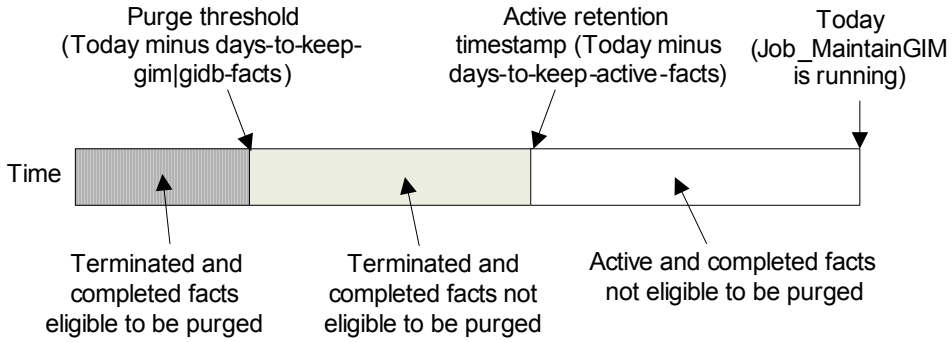

Figure 28: Active Retention Threshold Smaller Than Purge Retention Threshold

Figure 29 on page 407 illustrates the retention period for GIDB for the situation in which `days-to-keep-active-facts` is bigger than `days-to-keep-gidb-facts`. Previous comments about the starting point for calculating retention periods apply.

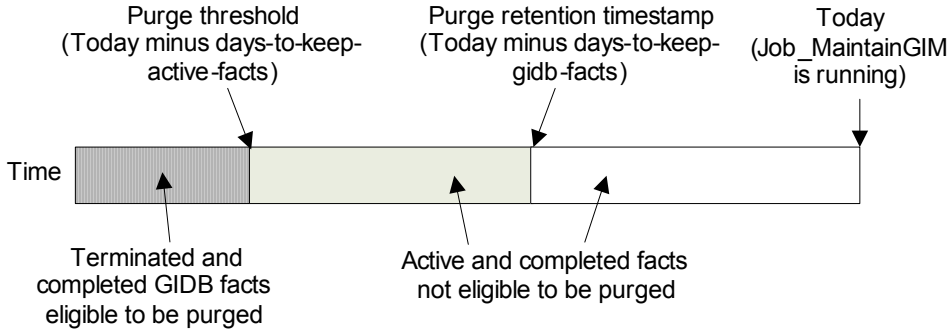

Figure 29: Active Retention Threshold Greater Than Purge Retention Threshold for GIDB Facts

For more information about data-retention policies, see the descriptions of the retention-policy options starting on page 281.

Voice and Multimedia Interaction Data

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

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

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

Retention Period for Discarded Operational Data

The retention period for records in the discard, audit, and history tables is calculated from the time that the ETL process discarded the data. The discard tables are Staging tables that store operational data that the transformation job was unable to process—for example, voice interaction data with unresolved IS-Links, or Configuration facts (configuration relationship records) with missing configuration objects. The audit and history tables are Control tables that store information about data lineage and about ETL-processing activity. Information in the Discard tables is useful for troubleshooting. Purging of discarded data is controlled by the `days-to-keep-discards-and-job-history` option in the `[gim-etl]` section (see page 283).

Purging Rules

The following rules apply to purging:

- Completed facts are evaluated for purging based on the start time of the fact. Similarly, active facts are evaluated for artificial termination based on the start time of the fact.
- Only completed facts or active facts that have been artificially terminated are eligible for purging, but there are different retention periods, controlled by separate configuration options, for completed fact data (before it is eligible to be purged) and for active facts (before they are eligible to be terminated). Similarly, the retention periods for completed facts in the dimensional model and in GIDB are different.

  When it determines the purge thresholds, Genesys Info Mart considers only the date part of timestamps; also, the transition from one day to the next is considered to occur at 00:00 (midnight) GMT. For example, if the starting point from which Genesys Info Mart calculates the retention period is 2010-12-14 17:07:00 and `days-to-keep-gidb-facts=14`, the purge threshold for completed GIDB facts is 2010-12-01 00:00:00.

  For more information, see “Retention Periods for Facts” on page 406.
- Discarded operational data is evaluated for purging based on the date portion of the timestamp of the ETL-processing event.

- Except for the GROUP_ANNEX and RESOURCE_ANNEX tables, which were introduced in release 8.1.4, dimension table rows are not eligible to be purged in either the dimensional model or GIDB.
- Starting with release 8.1.2, if the transformation job was not executed following the last run of the maintenance job, or if the last run of the transformation job failed, purge is suspended.

Configuration Recommendations

Setting Retention Periods

For a summary of the retention-policy options and default values, see the "Database maintenance" section in Table 15 on page 248. For detailed information about the options that control the various retention periods, see the descriptions of the days-to-keep-* options, starting on page 281.

Optimizing the Transaction Size

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

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

Optimal settings for these options for your environment balance purging execution time against the RDBMS resources that are required to purge the Info Mart database.

Scheduling the Maintenance Job

To schedule the maintenance job to run on a daily basis, you must:

1. Configure the ETL cycle (etl-start-time and etl-end-time) so that there is a window within which the maintenance job can run.
2. Enable the scheduler to launch the jobs (run-scheduler and—for Job_MaintainGIM—run-maintain).
3. Specify the time of day at which you want Job_MaintainGIM to start (maintain-start-time). The specified start time must be outside the ETL cycle.

In PostgreSQL deployments, you must also configure the run-update-stats and update-stats-schedule options so that the supplementary maintenance job, Job_UpdateStats, runs regularly throughout the day, outside the maintenance window.

All of the scheduling-related options are configured in the `[schedule]` section of the Genesys Info Mart application. For more information about the options, see “schedule Section” on page 314.

Chapter

20 High Availability

Genesys Info Mart supports high availability (HA) of all types of reporting data: Configuration details, Voice details, Multimedia details, and Outbound Contact details.

In an HA environment, Genesys Info Mart ensures that, if a component or network outage prevents one of the Interaction Concentrator (ICON) processes from obtaining and storing data from its data source in its Interaction Database (IDB), the data is not lost, as long as another ICON process is able to store the information in its IDB.

This chapter provides information about the following topics that relate to Genesys Info Mart support of HA for reporting:

Note: Genesys Info Mart supports HA redundancy at the Interaction Concentrator level; there is no HA for Genesys Info Mart itself. However, you can deploy a second instance of Genesys Info Mart as a standby, as described in "Standby and Disaster Recovery" on page 72.

Genesys Info Mart HA Model

The HA model that is used in Interaction Concentrator and Genesys Info Mart differs significantly from the Genesys standard HA model that is implemented in a majority of Genesys servers.

Unlike typical Genesys servers (for example, T-Server or OCS), two ICON instances that are deployed as an HA pair do not operate in either primary or

backup mode; nor does their mode switch from backup to primary when the other ICON server fails. Instead, the redundant ICONs operate in parallel as stand-alone servers: They process incoming data independently and store the data in independent IDBs.

The Genesys Info Mart HA model is based on redundant ICONs, each of which stores data in its own IDB. In other words, more than one IDB contains the same type of data from the same data source (Configuration Server, T-Server, Interaction Server, or OCS).

For each extraction domain (Configuration, Voice, Multimedia, or Outbound Contact), Genesys Info Mart extracts data from the IDB that contains the most reliable set of data during time slices in that extraction cycle. There is no limit on the number of redundant IDBs from which Genesys Info Mart can potentially extract data. The only limiting requirement is that data from the `G_IR`, `G_CALL`, and `G_IS_LINK` tables that relates to the same voice interaction must be extracted from the same IDB.

Architecture Diagrams

For illustrations of the Genesys Info Mart HA architecture, see Figure 3 on page 77. See also "Recommendations on Hosting" on page 83.

Time-Based Switchover

As described in "Data-Source Availability" on page 403, the extraction job detects whether HA mode is in effect.

In HA mode, Genesys Info Mart uses session information in the `G_DSS_*_PROVIDER` tables in the IDBs to identify which instance of IDB in an HA set contains the most complete and accurate set of data for a particular timeframe. Genesys Info Mart then extracts data from the IDB that contains the more reliable set of data for that timeframe. Genesys Info Mart switches over from one IDB to another, as required, to extract the best set of data that is available for time slices in each extraction cycle.

To avoid unnecessary or premature switchovers, the HA algorithm also considers values of the `NoData` timestamp in the IDB session-control tables, with allowance for inaccuracies in time synchronization between the ICON hosts in the deployment. The `NoData` timestamp is an indirect ICON heartbeat mechanism that records "no data" from ICON after a period of inactivity on the ICON connection with its data source; the inactivity interval is configurable (see the description of the ICON `dss-no-data-tout` option on page 167). The amount of time to allow as a cushion for time synchronization inaccuracies is also configurable (see the description of the Genesys Info Mart `max-time-deviation` option on page 296).

The switchover mechanism enables Genesys Info Mart to minimize data loss, if the ETL identifies gaps in the data between the previous high-water mark and the timestamps of events in the next extraction cycle.

For more information about data extraction in an HA deployment, see "Extracting Data in an HA Deployment" on page 381.

More Information—Using IDB Session Control Data

- For more information about the session information that is available in IDB, see the chapter about the data-source session-control schema in the *Interaction Concentrator 8.x Physical Data Model* for your RDBMS.
- For a description of the kind of analysis that the Genesys Info Mart ETL performs to detect gaps in the data from a particular data source, see the sections in the *Interaction Concentrator 8.x User's Guide* about determining data availability and reliability and about extracting HA data.

Criteria for Best IDB

The rules that govern how the ETL determines the time slices within an extraction cycle, as well as how it selects the best IDB source for each time slice, are designed to minimize both data loss and the number of switchovers.

- If Genesys Info Mart identifies that there are no gaps in the data from any of the redundant IDBs in an HA set:
  - If geolocation is a consideration (in other words, the `geo-location` option has been configured in the extraction and Info Mart DAPs), the ETL extracts from the local IDB. For more information, see "HA Data Extraction with Geolocation" on page 385.
  - If geolocation is not a consideration, the ETL continues to extract from the IDB that was used in the previous extraction cycle. On startup, it is not possible to predict which IDB the ETL will use.
- If Genesys Info Mart identifies that there are gaps in the data from one or more of the redundant IDBs, the ETL extracts data from:
  - The IDB in which there are no gaps in the data for the current extraction cycle
  - If there are disruptions in all the redundant IDBs, the IDB that had the longest period of uninterrupted data
  - If no IDB had an uninterrupted period longer than the other(s), the IDB that has the fewest gaps
- For an HA pair of IDBs that store Configuration details, Genesys Info Mart gives preference to ICON sessions that started with ICON resynchronization with Configuration Server.

Data in the Genesys Info Mart audit and control tables identifies the IDB(s) from which records were extracted during each extraction cycle, the ICON instance(s) that populated the IDB(s), and the original data-source application(s) (for example, T-Server).

Points of Failure

The session-control information in each IDB enables Genesys Info Mart to minimize data loss or data delays (which might result in data loss) from the following sources of failure, which might arise from either network or component outages:

- Disconnection of an ICON from its data source (Configuration Server, T-Server, Interaction Server, or OCS).
- Disconnection of an ICON from its IDB—The session-control tables in IDB provide sufficient information for Genesys Info Mart to distinguish between (a) the situation in which new data is not being stored in IDB, because there is no activity (ICON has not received any new data within the "no data" interval), and (b) failure of ICON or IDB.
- Disconnection of Genesys Info Mart from an IDB.

Consequences of Failures

As described in "Failure Scenarios and Recommendations" on page 394, dropped server connections and other failures that interrupt data might result in missing data or other inconsistencies in IDB that are propagated to Genesys Info Mart.

Additional data inconsistencies that might arise from the implementation of HA in Genesys Info Mart are:

- Skipped data (data is available in an IDB but Genesys Info Mart never extracts it). Data will be skipped in the following circumstances:
  - IDB data becomes available only after the stuck threshold for the next extraction window has expired (see Figure 22 on page 383).
  - Voice calls were created but not ended before the failure, and the redundant ICON(s) were not available when the calls were created. In this situation, the data is skipped, because Genesys Info Mart extracts only completed calls, while ICON does not store information about calls that were in progress when the connection to T-Server is established. For an illustration of this scenario, see Figure 23 on page 384 and Figure 24 on page 385.
- Incorrect or incomplete Info Mart data, as a result of missing or skipped IDB data—for example:
  - Incorrect agent states
  - Incomplete attached data
  - Incorrect Resource Roles or Technical Descriptors

    For example, for a voice interaction, ICON-1 fails while a consultation call is in progress, and ICON-2 was not available when the main call, which spawned the consultation, started. In this case, ICON-2 will report the consultation call as an independent call; Genesys Info Mart

might report a corresponding resource role of `RECEIVED`, instead of `RECEIVED_CONSULT`.

For more information about the data inconsistencies that might arise in Genesys Info Mart, see the section about data issues in a partially monitored environment in the *Genesys Info Mart 8.1 User's Guide*.

For related information, see "HA for Configuration Details" on page 415.

HA for Configuration Details

Every object in your contact center is mirrored as a configuration object in the Genesys environment, and the correctness and completeness of configuration data is vital to reporting on contact center activity. Therefore, Genesys recommends that you set up HA architecture for configuration data. Without HA, some configuration data might be lost, affecting your reporting data.

Resynchronizing IDB and Configuration Server Data

In a non-HA environment, with a single ICON application that captures Configuration details, you can use functionality that is provided by ICON to resynchronize configuration data between IDB and Configuration Server. However, Genesys does not recommend that you rely solely on resynchronization to ensure that your configuration data is up to date.

Limitations of Relying on Resynchronization Only

The Genesys Configuration Database does not contain historical data about configuration changes. When you request the resynchronization of configuration data in a non-HA environment, ICON retrieves current information from the Configuration Database, and the resynchronization might not retrieve a missing configuration object or relationship from the past. For example, if an agent was added and then deleted and ICON missed both events, there will be no configuration data about the agent in IDB. Subsequent resynchronization with Configuration Server will not provide the missing data, because the agent has been deleted from the Configuration Database. Similarly, if an agent was added to and then removed from an agent group, and ICON missed both events, resynchronization will not supply the missing relationship records.

In a non-HA environment, there is no redundant IDB with Configuration details from which Genesys Info Mart could extract data about missing objects or relationships. To compensate for the missing data, if the ETL detects missing configuration objects during transformation, it automatically creates them in the Info Mart database, based on the information that is available in interaction data.

- If the corresponding configuration objects are extracted from IDB at a later point, Genesys Info Mart updates the automatically created configuration objects with the extracted data.

- Otherwise, the Info Mart database stores the automatically created configuration objects. This could result in incomplete configuration data, because interaction data has only a limited set of configuration object attributes available.

Genesys Info Mart does not detect if any configuration object relationships:

- Are missing from the source (IDB).
- Have unreliable start and end times after the relationships have been recovered during the ICON resynchronization of configuration data.

Missing relationships or relationships that have unreliable timestamps result in data-quality issues, such as incorrect data for agent groups or queue groups, or incorrect indications of whether requested skills matched the skills of the handling agent.

Advantages of HA

An HA environment greatly reduces the likelihood of these issues, because two ICON processes are actively processing the source events from Configuration Server and are storing the data in their separate IDBs. If one IDB is missing a configuration object or relationship, Genesys Info Mart can extract the necessary data from the other IDB.

HA and Resynchronization

At the same time, HA architecture for configuration data does not remove the need for the on-demand resynchronization of IDB configuration data. This functionality helps to maintain data consistency between Configuration Database and each IDB in the HA pair, when you suspect that the data no longer is synchronized. Genesys strongly recommends that you utilize ICON functionality to resynchronize ICON configuration data automatically (see the description of the `cfg-auto-resync` option on page 166) or that you initiate resynchronization manually when a configuration data-source session is restored after a failure.

For each ICON that has the `cfg` role, Genesys Info Mart determines whether the ICON session started with a resynchronization with Configuration Server. Genesys Info Mart switches to the resynchonized ICON's session to extract updates.

Configuring for HA

To enable HA in your Genesys Info Mart deployment, you must store data from each data source in more than one IDB, configure DAPs for Genesys Info Mart to access those IDBs, and add the DAPs and all the ICON `Applications` that might populate the IDBs to the `Connections` tab in the Genesys Info Mart

`Application` object. Genesys Info Mart does not perform cross-checks to enforce the correct configuration.

Note: Genesys recommends that you also configure HA at the data-source level (for example, deploy a primary and backup OCS). Otherwise, with both ICONs monitoring the same data source, if the data source goes down, the data in both ICONs will be compromised.

Beyond ensuring that you configure identical redundant ICON applications that populate redundant IDBs through redundant DAPs, there are no special configuration requirements to enable HA in your Genesys Info Mart deployment. Genesys Info Mart automatically switches to HA mode when it detects that the DAPs are providing redundant data (in other words, there is more than one copy of data from the same Configuration Server, T-Server, Interaction Server, or OCS data source).

For more information about how to configure Genesys Info Mart to extract ICON details from redundant HA IDBs, see “Enabling High Availability” on page 152.

Part

5 Appendixes

Part 5 of this document provides the following appendixes:

- Appendix A, “Installation Worksheets,” on page 421
- Appendix B, “Sample ICON Attached Data Specification,” on page 437
- Appendix C, “IDB Tables Accessed by Genesys Info Mart,” on page 441
- Appendix D, “Info Mart Tables Purged by the Maintenance Job,” on page 445
- Appendix E, “Purge Functionality in Genesys Info Mart 8.1.0,” on page 451

Genesys

Appendix

A Installation Worksheets

This appendix contains worksheets that you can use to note configuration information for your environment, including relational database management system (RDBMS)–specific database-connection parameters for each database that Genesys Info Mart accesses, and the mapping of Outbound Contact Server (OCS) record fields to Info Mart tables. Keep this information, so that you can refer to it during deployment and when you need to re-install or upgrade Genesys Info Mart.

This appendix contains the following worksheets:

Worksheet for Oracle Databases

Oracle Databases	
For Oracle, the Database Name is the name that the database client software uses to connect to the database—for example, TNS name or Oracle Name Server name. The Owner ID and Owner Password specify the schema in which the database tables reside. The User ID and User Password are used to connect to the database.	
Interaction Concentrator Database	
Use to run SQL scripts during deployment of the Interaction Database (IDB), and to configure a DAP to extract Interaction Concentrator data. If you have multiple IDBs, note the connection information for each database.	
Database Name	
System ID (SID)	
Service Name	
Host Name	
Communication Port	
Owner ID	
Owner Password	
User ID	
User Password	
Genesys Info Mart Database	
Use to run SQL scripts during deployment of the Info Mart database, and to configure a DAP to connect to the Genesys Info Mart database.	
Database Name	
System ID (SID)	
Service Name	
Host Name	
Communication Port	
Owner ID	
Owner Password	
User ID	
User Password	

Tenant Admin	
Use to run SQL scripts during deployment of the Info Mart database to create read-only views of Genesys Info Mart tables.	
Database Name	
Owner ID	
Owner Password	
Tenant User	
The read-only user for the tenant view. If you have multiple tenants, note the connection information separately for each tenant.	
Database Name	
Owner ID	
Owner Password	

Worksheet for Microsoft SQL Databases

Microsoft SQL Databases	
For Microsoft SQL databases, the Database Name is the name of the Microsoft SQL Server database. The Owner ID and Owner Password specify the owner of the tables in the database. (The Owner ID is either the login that created the database tables or `dbo`, if the login that created the tables created the database that contains the tables or if it is a member of the `System Administrators` server role.) When in doubt, use the SQL Server Enterprise Manager to verify the owner of the tables. The User ID and User Password are used to connect to the database server.	
Interaction Concentrator Database	
Use to run SQL scripts during deployment of the Interaction Database (IDB), and to configure a DAP to extract Interaction Concentrator data. If you have multiple IDBs, note the connection information for each database.	
Database Name	
Host Name	
Communication Port	
Owner ID	
Owner Password	
User ID	
User Password	
Genesys Info Mart Database	
Use to run SQL scripts during deployment of the Info Mart database, and to configure a DAP to connect to the Genesys Info Mart database.	
Database Name	
Host Name	
Communication Port	
Owner ID	
Owner Password	
User ID	
User Password	

Tenant Admin	
Use to run SQL scripts during deployment of the Info Mart database to create read-only views of Genesys Info Mart tables.	
Database Name	
Owner ID	
Owner Password	
Tenant User	
The read-only user for the tenant view. If you have multiple tenants, note the connection information separately for each tenant.	
Database Name	
Owner ID	
Owner Password	

Worksheet for PostgreSQL Databases

PostgreSQL Databases	
For PostgreSQL databases, the Database Name is the name of the PostgreSQL database. The Owner ID and Owner Password specify the owner of the tables in the database. The Owner ID is the login that created the database tables. The User ID and User Password are used to connect to the database server. The Schema is a separate database object in PostgreSQL; it is used to logically group database objects. If a new schema is not created for the Info Mart database or Interaction Database (IDB), the default schema named `public` will be used.	
Interaction Concentrator Database	
Use to run SQL scripts during deployment of the Interaction Database (IDB), and to configure a DAP to extract Interaction Concentrator data. If you have multiple IDBs, note the connection information for each database.	
Database Name	
Host Name	
Communication Port	
Owner ID	
Owner Password	
User ID	
User Password	
Schema	
Genesys Info Mart Database	
Use to run SQL scripts during deployment of the Info Mart database, and to configure a DAP to connect to the Genesys Info Mart database.	
Database Name	
Host Name	
Communication Port	
Owner ID	
Owner Password	
User ID	
User Password	
Schema	

Tenant Admin	
Use to run SQL scripts during deployment of the Info Mart database to create read-only views of Genesys Info Mart tables.	
Database Name	
Owner ID	
Owner Password	
Tenant User	
The read-only user for the tenant view. If you have multiple tenants, note the connection information separately for each tenant. Starting with release 8.1.4, the name of the Tenant User schema must be the same as the name of the corresponding tenant user.	
Database Name	
Owner ID	
Owner Password	

Worksheet for Mapping User Data

Use the following worksheet to map key-value pairs (KVPs) to predefined or custom user-data fact or dimension tables in the Info Mart schema.

In the following worksheet:

- All the KVP names that are identified as Genesys-defined are predefined names for data that is attached by Genesys solutions—such as Enterprise Routing and Outbound Contact—and should *not* be changed.
- For the KVPs that are listed in the "Genesys-defined KVPs that are mapped by default" section, the mapping to Genesys Info Mart user-data tables is specified by default.
- For the KVPs that are listed in the "Genesys-defined KVPs that are not mapped by default" section, you must specify the mapping to user-data fact or dimension tables if you want Genesys Info Mart to store these KVPs for your reporting.
- You can add other KVPs that are generated in your contact center environment if you require them for your reports—either instead of, or in addition to, the listed KVPs. The "Other KVPs that are not mapped by default" section provides room for you to record custom mappings. The placeholder entries in this section have been provided to help you to match information to equivalent placeholders in the user-data template script, which you use to customize the Info Mart database schema (see Procedure: Customizing the user-data template script, on page 185).
- Genesys Info Mart does not require the mapping of KVPs to user-data tables to be one-to-one. You can supplement default mappings or further customize custom mappings by mapping a particular KVP to more than one user-data fact or dimension table.
- Because Genesys Info Mart uses special logic to process certain KVPs that are attached by Universal Routing Server (URS), a separate section groups the relevant keys ("Genesys-defined Routing KVPs").

Notes:

- If you provide customized mapping of KVPs to user-data dimension tables, be aware of possible RDBMS-related limitations regarding the length of KVP values. For more information, see "RDBMS Considerations" on page 124.
- You must configure Interaction Concentrator to store the `IApplication` KVP, even if you do not choose to map `IApplication` to a user-data table for your reporting purposes.

- In your deployment, if the value of the `Business Result` KVP can be changed after the interaction is completed, change the `Business Result` propagation rule that is stored in the Info Mart database to `PARTY`. Otherwise, Genesys Info Mart will ignore `Business Result` values that are defined during after call work (ACW).
  For more information about the propagation rules, see "Propagation Rules" on page 118.

KVP Name	Propagation Rule	Info Mart Database Target			
		Table Name	Column Name	Data Type	Default Value
Genesys-defined KVPs that are mapped by default					
CaseID	CALL	IRF_USER_DATA_GEN_1	CASE_ID	Character	DEFAULT_CASE_ID
CustomerID	CALL		CUSTOMER_ID	Character	DEFAULT_CUSTOMER_ID
GSW_CALL_ATTEMPT_GUID	CALL		GSW_CALL_ATTEMPT_GUID	Character	Null
IPurpose	IRF		IPURPOSE	Character	Null
Revenue	CALL		REVENUE	Character	Null
Satisfaction	CALL		SATISFACTION	Character	Null
ServiceObjective	CALL		SERVICE_OBJECTIVE	Character	Null
Business Result	CALL	INTERACTION_DESCRIPTOR	BUSINESS_RESULT	Character	DEFAULT_BUSINESS_RESULT
CustomerSegment	CALL		CUSTOMER_SEGMENT	Character	DEFAULT_CUSTOMER_SEGMENT
ServiceType	CALL		SERVICE_TYPE	Character	DEFAULT_SERVICE_TYPE
ServiceSubType	CALL		SERVICE_SUBTYPE	Character	DEFAULT_SERVICE_SUBTYPE
Genesys-defined Routing KVPs					
RRequestedSkillCombination		REQUESTED_SKILL (all columns), REQUESTED_SKILL_COMBINATION	SKILL_COMBINATION_STRING	Character	Unspecified

KVP Name	Propagation Rule	Info Mart Database Target			
		Table Name	Column Name	Data Type	Default Value
RStrategyName		STRATEGY	STRATEGY_NAME	Character	Unknown
RTargetObjectSelected		ROUTING_TARGET	AGENT_GROUP_NAME	Character	UNSPECIFIED
RTargetObjectSelected			PLACE_GROUP_NAME	Character	UNKNOWN
RTargetObjectSelected			SKILL_EXPRESSION	Character	UNSPECIFIED
RTargetObjectSelected			TARGET_OBJECT_SELECTED	Character	UNSPECIFIED
RTargetTypeSelected			ROUTING_TARGET_TYPE	Character	Unspecified
Genesys-defined KVPs that are not mapped by default					
IApplication				Character	
IResult				Character	
IResultReason				Character	
ISpeechRecognition				Character	
ITextToSpeech				Character	
Other KVPs that are not mapped by default					
Custom User Data Fact Table (for high-cardinality KVPs): <IRF_USER_DATA_CUST_1> User Data Columns: <CUSTOM_DATA_*N*> **Note:** Data types can be character, numeric, or date/time.					
<CustomData*N*>					

KVP Name	Propagation Rule	Info Mart Database Target			
		Table Name	Column Name	Data Type	Default Value
Custom User Data Dimension Table (for low-cardinality KVPs): <USER_DATA_CUST_DIM_1> User Data Columns: <DIM_ATTRIBUTE_*N*> Primary Key: <ID> Foreign Key: <CUSTOM_KEY_1>					
<CustomAttribute*N*>				Character	
				Character	
				Character	
				Character	
				Character	

Worksheet for Mapping OCS Record Fields

Use the three checklists in this section to plan how to map Outbound Contact Server (OCS) record fields to the following Info Mart database tables:

- CONTACT_ATTEMPT_FACT
- RECORD_FIELD_GROUP_1
- RECORD_FIELD_GROUP_2

OCS Record Fields to CONTACT_ATTEMPT_FACT Table Mapping			
Genesys Info Mart Column Name	**Field Object Name**	**Tenant Name (if applicable)**	**Is Secure**
Genesys Info Mart stores nonmandatory record field data that is defined in OCS calling lists in the RECORD_FIELD_1 through RECORD_FIELD_40 columns of the CONTACT_ATTEMPT_FACT table. Use the following rows to list the nonmandatory fields that you need to map from OCS to the CONTACT_ATTEMPT_FACT table. The first row contains an example for RECORD_FIELD_1. In the OCS calling list, the name of the record field is Region. This field maps to a column that is named RECORD_FIELD_1 in the CONTACT_ATTEMPT_FACT table. **Note:** Make sure that the data type of the field matches the data type of the target CONTACT_ATTEMPT_FACT column. Unless otherwise indicated, the data type of the columns is *string* (varchar).			
Example: *RECORD_FIELD_1 (numeric)*	Region	HQ	✓
RECORD_FIELD_1 (numeric)			
RECORD_FIELD_2 (numeric)			
RECORD_FIELD_3 (numeric)			
RECORD_FIELD_4 (numeric)			
RECORD_FIELD_5 (numeric)			
RECORD_FIELD_6 (numeric)			
RECORD_FIELD_7 (numeric)			
RECORD_FIELD_8 (numeric)			

OCS Record Fields to CONTACT_ATTEMPT_FACT Table Mapping (Continued)			
Genesys Info Mart Column Name	**Field Object Name**	**Tenant Name (if applicable)**	**Is Secure**
RECORD_FIELD_9 (numeric)			
RECORD_FIELD_10 (numeric)			
RECORD_FIELD_11 (int)			
RECORD_FIELD_12 (int)			
RECORD_FIELD_13 (int)			
RECORD_FIELD_14 (int)			
RECORD_FIELD_15 (int)			
RECORD_FIELD_16 (int)			
RECORD_FIELD_17 (int)			
RECORD_FIELD_18 (int)			
RECORD_FIELD_19 (int)			
RECORD_FIELD_20 (int)			
RECORD_FIELD_21 (int)			
RECORD_FIELD_22 (int)			
RECORD_FIELD_23 (int)			
RECORD_FIELD_24 (int)			
RECORD_FIELD_25 (int)			

OCS Record Fields to CONTACT_ATTEMPT_FACT Table Mapping (Continued)			
Genesys Info Mart Column Name	**Field Object Name**	**Tenant Name (if applicable)**	**Is Secure**
RECORD_FIELD_26 (int)			
RECORD_FIELD_27 (int)			
RECORD_FIELD_28 (int)			
RECORD_FIELD_29 (int)			
RECORD_FIELD_30 (int)			
RECORD_FIELD_31			
RECORD_FIELD_32			
RECORD_FIELD_33			
RECORD_FIELD_34			
RECORD_FIELD_35			
RECORD_FIELD_36			
RECORD_FIELD_37			
RECORD_FIELD_38			
RECORD_FIELD_39			
RECORD_FIELD_40			

OCS Record Fields to RECORD_FIELD_GROUP_1 Table Mapping	
Genesys Info Mart Column Name	**Field Object Name**
Genesys Info Mart stores nonmandatory record field data that is defined in OCS calling lists in the RECORD_FIELD_1_STRING_1 through RECORD_FIELD_1_STRING_10 columns of the RECORD_FIELD_GROUP_1 table. Use the following rows to list the nonmandatory fields that you need to map from OCS to the RECORD_FIELD_GROUP_1 table. **Note:** Make sure that the data type of the field matches the data type of the target RECORD_FIELD_GROUP_1 column.	
RECORD_FIELD_1_STRING_1	
RECORD_FIELD_1_STRING_2	

OCS Record Fields to RECORD_FIELD_GROUP_1 Table Mapping (Continued)	
Genesys Info Mart Column Name	**Field Object Name**
RECORD_FIELD_1_STRING_3	
RECORD_FIELD_1_STRING_4	
RECORD_FIELD_1_STRING_5	
RECORD_FIELD_1_STRING_6	
RECORD_FIELD_1_STRING_7	
RECORD_FIELD_1_STRING_8	
RECORD_FIELD_1_STRING_9	
RECORD_FIELD_1_STRING_10	

OCS Record Fields to RECORD_FIELD_GROUP_2 Table Mapping	
Genesys Info Mart Column Name	**Field Object Name**
Genesys Info Mart stores nonmandatory record field data that is defined in OCS calling lists in the RECORD_FIELD_2_STRING_1 through RECORD_FIELD_2_STRING_10 columns of the RECORD_FIELD_GROUP_2 table. Use the following rows to list the nonmandatory fields that you need to map from OCS to the RECORD_FIELD_GROUP_2 table. **Note:** Make sure that the data type of the field matches the data type of the target RECORD_FIELD_GROUP_2 column.	
RECORD_FIELD_2_STRING_1	
RECORD_FIELD_2_STRING_2	
RECORD_FIELD_2_STRING_3	
RECORD_FIELD_2_STRING_4	
RECORD_FIELD_2_STRING_5	
RECORD_FIELD_2_STRING_6	
RECORD_FIELD_2_STRING_7	
RECORD_FIELD_2_STRING_8	
RECORD_FIELD_2_STRING_9	
RECORD_FIELD_2_STRING_10	

Appendix

B Sample ICON Attached Data Specification

This appendix provides an example of the `ccon_adata_spec_GIM_example.xml` file that is included in the `sql_scripts` folder in your Genesys Info Mart 8.1.x installation package. This file is also available in the `sql_scripts` folder on the Genesys Info Mart product CD.

For the convenience of Genesys Interactive Insights (GI2) customers who might want to include social media metrics—for example, Sentiment, Actionability, and Influence—the script includes commented-out lines for specifying the key names that GI2 requires.

Starting with release 8.1.2, you can use the User Data Assistant, which is briefly described on page 148, to help you customize the attached-data specification file.

Note: Genesys Info Mart requires that you specify a value of `all` for the `history` attribute. A value of `all` means that Interaction Concentrator (ICON) will record every change in value for the specified key. Be aware that your reporting results might be affected if you specify a different value for the `history` attribute.

```
<?xml version="1.0" encoding="utf-8" ?>
<!--
   This xml contains the customizable mapping of Attached User Data for use in Genesys
   Info Mart 8.x. The only sections that Genesys Info Mart uses are "public" and
   "secure".  A "key name" can be specified only once, in either the "public" section or
   the "secure" section; if the "key name" is duplicated, ICON will ignore the second
   and subsequent entries.
 -->
<adata_spec>
    <public>
```

```
     <!--
         Predefined keys mapped to the CUSTOMER_SEGMENT, SERVICE_TYPE,
SERVICE_SUBTYPE, and BUSINESS_RESULT columns, respectively, in the
INTERACTION_DESCRIPTOR table.
      -->
     <key name="CustomerSegment"       source="userdata" history="all"/>
     <key name="ServiceType"           source="userdata" history="all"/>
     <key name="ServiceSubType"        source="userdata" history="all"/>
     <key name="Business Result"       source="userdata" history="all"/>
    <!--
         Predefined keys mapped to IRF_USER_DATA_GEN_1 table, CASE_ID, CUSTOMER_ID,
SERVICE_OBJECTIVE, REVENUE, and SATISFACTION columns, respectively, in the
IRF_USER_DATA_GEN_1 table.
      -->
     <key name="CaseID"                source="userdata" history="all"/>
     <key name="CustomerID"            source="userdata" history="all"/>
     <key name="ServiceObjective"      source="userdata" history="all"/>
     <key name="Revenue"               source="userdata" history="all"/>
     <key name="Satisfaction"          source="userdata" history="all"/>
     <!--
        Predefined IApplication key,which is used internally by Genesys Info Mart to
create resources of the IVR Application type.
      -->
     <key name="IApplication"          source="userdata" history="all"/>
     <!--
         Predefined IPurpose key, which is used internally by Genesys Info Mart to
process self-service IVRs.
      -->
     <key name="IPurpose"              source="userdata" history="all"/>
     <!--
         Predefined GSW_CALL_ATTEMPT_GUID key, which is used internally by Genesys
Info Mart to process OCS data.
      -->
     <key name="GSW_CALL_ATTEMPT_GUID" source="userdata" history="all"/>

     <!--
         Custom keys for high-cardinality user data (facts).
         For example, for the IRF_USER_DATA_CUST_1 table.
      -->
     <key name="CustomData1"           source="userdata" history="all"/>
     <key name="CustomData2"           source="userdata" history="all"/>
     <key name="CustomData3"           source="userdata" history="all"/>
     <key name="CustomData4"           source="userdata" history="all"/>
     <key name="CustomData5"           source="userdata" history="all"/>
     <key name="CustomData6"           source="userdata" history="all"/>
     <key name="CustomData7"           source="userdata" history="all"/>
     <key name="CustomData8"           source="userdata" history="all"/>
     <key name="CustomData9"           source="userdata" history="all"/>
     <key name="CustomData10"          source="userdata" history="all"/>
     <key name="CustomData11"          source="userdata" history="all"/>
     <key name="CustomData12"          source="userdata" history="all"/>
```

```
      <key name="CustomData13"          source="userdata" history="all"/>
      <key name="CustomData14"          source="userdata" history="all"/>
      <key name="CustomData15"          source="userdata" history="all"/>
      <key name="CustomData16"          source="userdata" history="all"/>

      <!--
          Custom keys for low-cardinality user data (dimensions).
          For example, for the USER_DATA_CUST_DIM_1 table.
       -->
      <key name="CustomAttribute1"      source="userdata" history="all"/>
      <key name="CustomAttribute2"      source="userdata" history="all"/>
      <key name="CustomAttribute3"      source="userdata" history="all"/>
      <key name="CustomAttribute4"      source="userdata" history="all"/>
      <key name="CustomAttribute5"      source="userdata" history="all"/>

<!--
     The keys that are commented out in this section are required for support of
eServices/Social Media metrics and dimensions.

Include (uncomment) these keys if you require reporting on eServices/Social Media.

     <key name="Classify_Actionability_CtgRelevancy" source="userdata"
history="all"/>
     <key name="Classify_Sentiment_CtgRelevancy" source="userdata" history="all"/>
     <key name="KloutScore" source="userdata" history="all"/>
     <key name="CtgName" source="userdata" history="all"/>
     <key name="Screen_Sentiment_CtgName" source="userdata" history="all"/>
     <key name="Screen_Actionability_CtgName" source="userdata" history="all"/>
     <key name="Classify_Actionability_CtgName" source="userdata" history="all"/>
     <key name="Classify_Sentiment_CtgName" source="userdata" history="all"/>
     <key name="desktop_influence" source="userdata" history="all"/>
-->

</public>
 <secure>
 </secure>
 <call>
 </call>
 <call-cust>
 </call-cust>
 <call-cust1>
 </call-cust1>
 <call-cust2>
 </call-cust2>
 <mcr-f>
 </mcr-f>
 <mcr-l>
 </mcr-l
 </adata_spec>
```

Appendix

C IDB Tables Accessed by Genesys Info Mart

This appendix lists the Interaction Database (IDB) tables that the Genesys Info Mart extraction jobs access (see Table 25 on page 442). The type of data that is extracted depends on the role that is configured for the database access point (DAP).

Note: Table 25 lists only those IDB tables from which Genesys Info Mart actually extracts data. Genesys Info Mart does not populate all of the GIDB tables in the Info Mart database schema.

Table 25: IDB Tables Extracted by Genesys Info Mart

DAP Role	Type of Data	Tables Extracted
ICON_CFG	Configuration details Data source: Configuration Server	GCX_AGENT_PLACE GCX_CAMPGROUP_INFO GCX_ENDPOINT_PLACE GCX_GROUP_AGENT GCX_GROUP_PLACE GCX_GROUP_ROUTEDN GCX_LOGIN_INFO GCX_SKILL_LEVEL GC_AGENT GC_ANNEX GC_CALLING_LIST GC_CAMPAIGN GC_ENDPOINT GC_GROUP GC_IVR GC_IVRPORT GC_LOGIN GC_PLACE GC_SCRIPT GC_SKILL GC_SWITCH GC_TENANT GC_TIME_ZONE G_DSS_CFG_PROVIDER

Table 25: IDB Tables Extracted by Genesys Info Mart (Continued)

DAP Role	Type of Data	Tables Extracted
ICON_CORE	Voice details (including voice interaction, voice agent activity, user data, and virtual-queue activity) Data source: T-Server	GSYS_SYSPROCINFO G_IR G_IS_LINK G_IS_LINK_HISTORY G_CALL G_CALL_STAT G_PARTY G_PARTY_HISTORY GX_SESSION_ENDPOINT G_VIRTUAL_QUEUE G_ROUTE_RESULT G_ROUTE_RES_VQ_HIST G_USERDATA_HISTORY G_SECURE_USERDATA_HISTORY G_LOGIN_SESSION G_AGENT_STATE_HISTORY G_AGENT_STATE_RC G_AGENT_STATE_RC_A G_DND_HISTORY G_CUSTOM_DATA_S G_DSS_GCC_PROVIDER G_DSS_GLS_PROVIDER
ICON_OCS	Outbound Contact details Data source: Outbound Contact Server	GO_CHAIN GO_CHAINREC_HIST GO_FIELDHIST GO_SEC_FIELDHIST GO_METRICS GO_CAMPAIGN GO_CAMPAIGNHISTORY GOX_CHAIN_CALL G_DSS_GOS_PROVIDER

Table 25: IDB Tables Extracted by Genesys Info Mart (Continued)

DAP Role	Type of Data	Tables Extracted
ICON_MM	Multimedia details (including multimedia interactions, attached data, virtual queue, interaction queue, interaction workbin, and agent login, and state and reasons activity) Data source: Interaction Server	G_IR G_CALL G_PARTY G_PARTY_HISTORY GX_SESSION_ENDPOINT G_AGENT_STATE_HISTORY G_AGENT_STATE_RC G_AGENT_STATE_RC_A G_VIRTUAL_QUEUE G_ROUTE_RESULT G_ROUTE_RES_VQ_HIST G_USERDATA_HISTORY G_SECURE_USERDATA_HISTORY GM_F_USERDATA GM_L_USERDATA G_LOGIN_SESSION G_DND_HISTORY G_DSS_GCC_PROVIDER G_DSS_GLS_PROVIDER

Appendix

D Info Mart Tables Purged by the Maintenance Job

This appendix lists the tables in the Info Mart database schema that the Genesys Info Mart maintenance job purges. The maintenance job automatically purges all data that it determines is eligible to be purged, based on configurable retention policies. The maintenance job can be scheduled to run daily, or it can be run manually from the Genesys Info Mart management GUI.

This appendix contains the following sections:

For more information about the configuration options that specify the data-retention policies, see the information about the `days-to-keep-*` options in "gim-etl Section" on page 281.

Dimensional Model Tables

Table 26 on page 446 lists the fact tables in the dimensional model that are purged by `Job_MaintainGIM`.

Table 26: Dimensional Model Fact Tables Purged by Job_MaintainGIM

Type of Data	Tables Purged	Purge Eligibility
Voice and Multimedia interactions	INTERACTION_FACT INTERACTION_RESOURCE_FACT IXN_RESOURCE_STATE_FACT IRF_USER_DATA_KEYS and custom high-cardinality user-data tables (IRF_CUST_* in the make_gim_UDE_template.sql or make_gim_UDE_template_partitioned.sql script) MEDIATION_SEGMENT_FACT	days-to-keep-gim-facts
Agent activity	SM_RES_SESSION_FACT SM_RES_STATE_FACT SM_RES_STATE_REASON_FACT	days-to-keep-gim-facts
Outbound Contact activity	CAMPAIGN_GROUP_SESSION_FACT CALLING_LIST_METRIC_FACT CAMPAIGN_GROUP_STATE_FACT CONTACT_ATTEMPT_FACT	days-to-keep-gim-facts

Table 27 lists the dimension tables in the dimensional model that are purged by Job_MaintainGIM.

Table 27: Dimensional Model Dimension Tables Purged by Job_MaintainGIM

Type of Data	Tables Purged	Purge Eligibility
Configuration-related	GROUP_ANNEX RESOURCE_ANNEX	days-to-keep-deleted-annex

GIDB Fact Tables

Table 28 lists the fact tables in GIDB that are purged by `Job_MaintainGIM`.

Table 28: GIDB Tables Purged by Job_MaintainGIM

Type of Data	Tables Purged	Purge Eligibility

Voice interactions	`GIDB_G_IR_V`
	`GIDB_G_IR_HISTORY_V`
	`GIDB_G_CALL_V`
	`GIDB_G_VIRTUAL_QUEUE_V`
	`GIDB_G_CALL_HISTORY_V`
	`GIDB_G_CALL_STAT_V`
	`GIDB_G_ROUTE_RES_VQ_HIST_V`
	`GIDB_G_ROUTE_RESULT_V`
	`GIDB_G_SECURE_UD_HISTORY_V`
	`GIDB_G_USERDATA_HISTORY_V`
	`GIDB_G_CUSTOM_DATA_S_V`
	`GIDB_G_IS_LINK_V`
	`GIDB_G_IS_LINK_HISTORY_V`
	`GIDB_G_PARTY_V`
	`GIDB_G_PARTY_HISTORY_V`
		`days-to-keep-gidb-facts`

Multimedia interactions	`GIDB_G_IR_MM`
	`GIDB_G_IR_HISTORY_MM`
	`GIDB_G_CALL_MM`
	`GIDB_G_VIRTUAL_QUEUE_MM`
	`GIDB_G_CALL_HISTORY_MM`
	`GIDB_G_ROUTE_RES_VQ_HIST_MM`
	`GIDB_G_ROUTE_RESULT_MM`
	`GIDB_G_SECURE_UD_HISTORY_MM`
	`GIDB_G_USERDATA_HISTORY_MM`
	`GIDB_G_CUSTOM_DATA_S_MM`
	`GIDB_GM_F_USERDATA`
	`GIDB_GM_L_USERDATA`
	`GIDB_G_PARTY_MM`
	`GIDB_G_PARTY_HISTORY_MM`
		`days-to-keep-active-facts` or `days-to-keep-gidb-facts`, whichever is greater

Table 28: GIDB Tables Purged by Job_MaintainGIM (Continued)

Type of Data	Tables Purged	Purge Eligibility
Agent activity	GIDB_G_AGENT_STATE_HISTORY_V GIDB_G_AGENT_STATE_RC_V GIDB_G_DND_HISTORY_V GIDB_G_LOGIN_SESSION_V GIDB_GX_SESSION_ENDPOINT_V GIDB_G_AGENT_STATE_HISTORY_MM GIDB_G_AGENT_STATE_RC_MM GIDB_G_DND_HISTORY_MM GIDB_G_LOGIN_SESSION_MM GIDB_GX_SESSION_ENDPOINT_MM	days-to-keep-gidb-facts
Outbound Contact activity	GIDB_GO_CAMPAIGN GIDB_GO_CAMPAIGNHISTORY GIDB_GO_METRICS GIDB_GO_CHAIN GIDB_GO_CHAINREC_HIST GIDB_GO_FIELDHIST GIDB_GO_SEC_FIELDHIST GIDB_GOX_CHAIN_CALL	days-to-keep-gidb-facts

Discard and Control Tables

- Job_MaintainGIM purges discarded operational data from the following Staging tables:
  - STG_TRANSFORM_DISCARDS
  - STG_IDB_FK_VIOLATION
- Job_MaintainGIM purges data lineage and ETL history data from the following Control tables:
  - CTL_AUDIT_LOG
  - CTL_ETL_HISTORY
  - CTL_EXTRACT_HISTORY
  - CTL_PURGE_HISTORY
  - CTL_TRANSFORM_HISTORY

- Purge eligibility is determined by the `days-to-keep-discards-and-job-history` configuration option, calculated from the timestamp of the ETL processing event—`ETL_TS` for the Staging tables and `CREATED_TS` for the Control tables.

Genesys

Appendix

E Purge Functionality in Genesys Info Mart 8.1.0

Purge behavior—in particular, the role of the `days-to-keep-active-facts` configuration option—changed between Genesys Info Mart release 8.1.0 and release 8.1.1. For the benefit of Genesys Info Mart 8.1.0 customers, this appendix reproduces previous documentation to describe purge behavior in release 8.1.0.

This appendix contains the following sections:

Purge-Related Configuration Options

[Information in this section is the 8.1.0 equivalent of descriptions from Chapter 14 on page 271.]

days-to-keep-active-facts

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

Specifies the maximum number of days to retain active data in GIDB and the dimensional model, including certain Staging tables. This option applies only to long-living entities.

- For the purposes of purging GIDB, long-living entities are multimedia interactions and Outbound Contact campaigns.

- For the purposes of purging the dimensional model, long-living entities are voice interactions, multimedia interactions, and Outbound Contact campaigns.

This option sets the lower boundary of the period within which both active and completed long-living facts are eligible to be purged. The earliest active fact that started between the time that is specified by this option and the applicable completed facts option (`days-to-keep-gidb-facts` or `days-to-keep-gim-facts`) determines the actual purge threshold. All facts that started earlier than the purge threshold in a group of related tables are purged, but the purge threshold itself is based on the start timestamp in a root table.

The purging algorithm uses the following logic:

1. Determine the retention timestamp for completed facts (`retentionTS`):
   - For GIDB:
     `retentionTS` = (`ExtractHWM` – `days-to-keep-gidb-facts`).
   - For the dimensional model:
     `retentionTS` = (`ExtractHWM` – `days-to-keep-gim-facts`).
2. Determine the retention timestamp for active facts (`activeTS`):
   `activeTS` = (`ExtractHWM` – `days-to-keep-active-facts`).
3. Determine the purge threshold (`purgeTS`):
   `purgeTS` = (start time of earliest active entity in the root table), where the start time is between `retentionTS` and `activeTS`.
4. For each table in the group, purge from the table in which (Start timestamp) < `purgeTS`.

   In determining the purge thresholds, Genesys Info Mart considers only the date part of timestamps, and the transition from one day to the next is considered to occur at 00:00 (midnight) GMT. For example, if the ExtractHWM is 2010-12-14 17:07:00 and `days-to-keep-gidb-facts=14`, the purge threshold for completed GIDB facts is 2010-12-01 00:00:00.

Active interactions can significantly affect when completed facts are purged. Tables 29 and 30, starting on page 455, provide examples of different scenarios for GIDB and the dimensional model, respectively.

For a list of the tables in the dimensional model that contain records for long-living entities (as defined for purging purposes), including the root tables for each group of facts, see "Info Mart Tables Purged by the Maintenance Job" on page 465.

For information about purging rules and guidelines to consider when you are setting the value of this option, see the sections about purging rules and purging strategies, starting on page 459.

days-to-keep-gidb-facts

Default Value: `14`
Valid Values: Any positive integer

Dependencies: None
Changes Take Effect: At the next run of `Job_MaintainGIM`

Specifies the number of days to retain terminated data in GIDB.

- For short-living entities in GIDB (voice interactions and agent session activity), this option directly specifies the retention period. Facts that have a start time that is earlier than the retention period will be purged. The purge condition is:

  (Start timestamp) < (`ExtractHWM` – `days-to-keep-gidb-facts`)

- For long-living entities in GIDB (multimedia interactions and Outbound Contact campaigns), this option specifies the minimum number of days after which terminated facts might be eligible to be purged. The eligibility criterion is based on the start timestamp in a root table. However, eligible completed facts will not actually be purged until all other active facts that have the same start timestamps or later also become eligible to be purged, as described in `days-to-keep-active-facts`. If there are no active facts, then `days-to-keep-gidb-facts` specifies the retention period for long-living entities.

For a list of the GIDB tables that contain short- and long-living entities (as defined for purging purposes), including the root tables for each category of data, see "Info Mart Tables Purged by the Maintenance Job" on page 465.

For information about purging rules and guidelines to consider when setting the value of this option, see the sections about purging rules and purging strategies, starting on page 459.

days-to-keep-gim-facts

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

Specifies the number of days to retain terminated data in the dimensional model, including certain Staging tables. `Job_MaintainGIM` does not purge dimension data or aggregate tables (if aggregation is enabled).

- For short-living entities in the dimensional model (agent session activity), this option directly specifies the retention period. Facts that have a start time that is earlier than the retention period will be purged. The purge condition is:

  (Start timestamp) < (`ExtractHWM` – `days-to-keep-gim-facts`)

- For long-living entities in the dimensional model (voice and multimedia interactions, and Outbound Contact campaigns), this option specifies the minimum number of days after which terminated facts are eligible to be purged. The eligibility criterion is based on the start timestamp in a root table. However, eligible completed facts will not actually be purged until all other active facts that have the same start timestamps or later also

become eligible to be purged, as described in `days-to-keep-active-facts` (see page 451). If there are no active facts, then `days-to-keep-gim-facts` specifies the retention period for long-living entities.

Note: By definition, all voice interaction data in the dimensional model relates to completed interactions. However, as Table 30 on page 456 shows, active multimedia interactions can significantly delay the purging of completed interactions. If this is not acceptable for your circumstances, and if you are not required to combine voice and multimedia for reporting purposes, consider installing separate instances of the Genesys Info Mart Server application and Info Mart database schema for Voice and Multimedia details.

For a list of the tables in the dimensional model that contain short- and long-living entities (as defined for purging purposes), including the root tables for each category of data, see "Info Mart Tables Purged by the Maintenance Job" on page 465.

For information about purging rules and guidelines to consider when setting the value of this option, see the sections about purging rules and purging strategies, starting on page 459.

How Purge Options Work

To illustrate how various `days-to-keep-*` options combine to determine when data will be purged, Tables 29 and 30, starting on page 455, provide examples of different scenarios for GIDB and the dimensional model, respectively. For related information, see "Purging Fact Data from GIDB" on page 462 and "Purging Fact Data from the Dimensional Model" on page 460.

Table 29: Example of GIDB Purging Options

Scenario		Start Time of Facts to Be Purged	
		Short-Living Entities	Long-Living Entities
• `days-to-keep-gidb-facts=14`. • `days-to-keep-active-facts=600`. • `Job_MaintainGIM` is running on June 4, 2012, after the last ETL cycle for the day (in other words, `days-to-keep-gidb-facts` threshold is May 21, 2012; `days-to-keep-active-facts` threshold is October 12, 2010).			
1	The earliest active multimedia fact relates to an e-mail interaction that started on March 20, 2011. There are no active Outbound Contact facts.	Voice interactions: May 21, 2012, or earlier Agent activity: May 21, 2012, or earlier	Multimedia interactions: March 20, 2011, or earlier Outbound Contact facts: May 21, 2012, or earlier
2	An e-mail interaction that started on October 12, 2010, or earlier remains active in a workbin.		Multimedia interactions: October 12, 2010, or earlier
3	The earliest active fact relates to an e-mail interaction that started on May 27, 2012.		Multimedia interactions: May 21, 2012, or earlier

Table 30: Example of GIM (Dimensional Model) Purging Options

Scenario		Start Time of Facts to Be Purged	
		Short-Living Entities	Long-Living Entities
• days-to-keep-gim-facts=400. • days-to-keep-active-facts=600. • Job_MaintainGIM is running on June 4, 2012, after the last ETL cycle for the day (in other words, days-to-keep-gim-facts threshold is May 1, 2011; days-to-keep-active-facts threshold is October 12, 2010).			
1	The earliest active fact relates to an e-mail interaction that started on March 20, 2011. There are no active Outbound Contact facts.	Agent activity: May 1, 2011, or earlier	Voice and Multimedia interactions: March 20, 2011, or earlier Outbound Contact facts: May 1, 2011, or earlier
2	An e-mail interaction that started on October 12, 2010, or earlier remains active in a workbin.		Voice and Multimedia interactions: October 12, 2010, or earlier
3	The earliest active fact relates to an e-mail interaction that started on June 20, 2011.		Voice and Multimedia interactions: May 1, 2011, or earlier

Purging Info Mart Data

[Information in this section is the 8.1.0 equivalent of material from Chapter 19 on page 401.]

Job_MaintainGIM purges:

- Completed and active fact data from GIDB, in accordance with data-retention policies that are specified by the days-to-keep-gidb-facts and days-to-keep-active-facts options (see pages 452 and 451)
- Completed and active fact data from the dimensional model, including certain Staging tables, in accordance with data-retention policies that are specified by the days-to-keep-gim-facts and days-to-keep-active-facts options (see pages 453 and 451)

- Discarded operational data from discard tables, as well as data lineage and ETL job history data from various Control tables, in accordance with a data-retention policy that is specified by the `days-to-keep-discards-and-job-history` option (see page 283)

Notes:

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

Retention Periods

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

For the configuration options that determine the retention periods, the values that you choose must allow the Info Mart database to retain the data long enough for you to complete deployment-specific tasks (such as 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 retention period for Info Mart fact data to 380 days. This allows for the extra day in a leap year, plus two extra weeks for deployment-specific tasks.

Retention Periods for Facts

The subjects of fact records are entities such as interactions, agent login sessions, or Outbound Contact campaigns.

Short-Living and Long-Living Entities

The purging algorithms divide dimensional model and GIDB fact tables into two categories:

- Tables that contain data about *short-living entities*—Entities whose lifespans are well defined and of short durations, typically less than a day (for example, agent login sessions). `Job_MaintainGIM` purges records relating to short-living entities based on the start time of facts in the applicable tables. The `days-to-keep-gim-facts` and `days-to-keep-gidb-facts` options in the `[gim-etl]` section of the Genesys Info Mart configuration directly specify the retention period for short-living entities.
- Tables that contain data about *long-living entities*—Entities whose lifespans are indefinite and potentially of long durations (for example, e-mail interactions). `Job_MaintainGIM` purges long-living entities based on the start time of the earliest active fact in the interval that is defined by the

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

- If there are active facts within the interval, the start time of the earliest active fact determines the purge eligibility threshold (in other words, the retention period) for both active and completed long-living facts. All completed long-living facts that started earlier than the purge threshold in a group of related tables are purged, but the purge eligibility threshold itself is based on the start timestamp in a root table.
- If there are no active facts within the interval, the `days-to-keep-gim-facts` and `days-to-keep-gidb-facts` options determine the retention period for completed long-living facts.

For a list of the tables in GIDB and in the dimensional model that contain records for short- and long-living entities (as defined for purging purposes), including the root tables for each category of data, see "Info Mart Tables Purged by the Maintenance Job" on page 465.

Figure 30 illustrates the purging strategy for long-living entities.

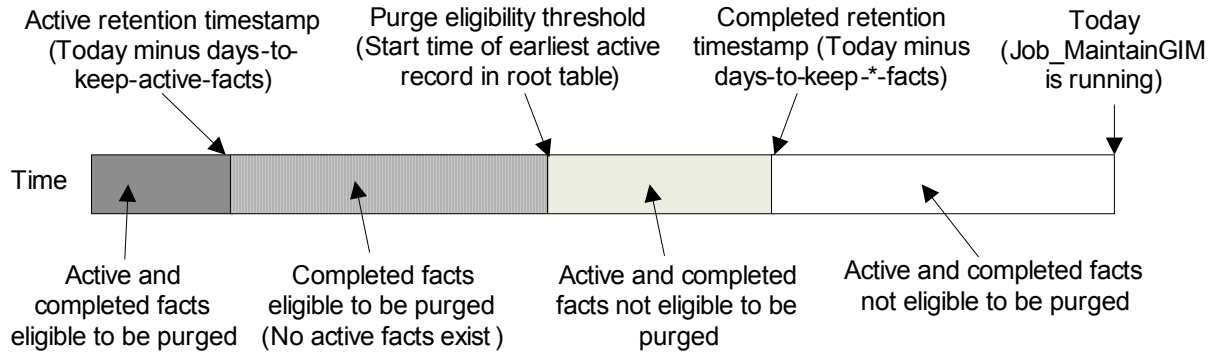

Figure 30: Purge Threshold for Long-Living Entities

For more information about data-retention policies, see the descriptions of the retention-policy options starting on page 451. For more detailed examples of how the retention thresholds define purge eligibility, see "Purge Strategies" on page 460.

Voice and Multimedia Interaction Data

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

In GIDB, voice interaction data and multimedia interaction data are stored in separate, media-specific tables. For the purposes of purging GIDB, voice

interactions are considered to be short-living entities, and multimedia interactions are considered to be long-living entities.

In the dimensional model, voice interaction data and multimedia interaction data are stored in common tables. For the purposes of purging the dimensional model, both voice interactions and multimedia interactions are considered to be long-living entities.

Notes:

- Active multimedia interactions might delay the purging of voice data from the dimensional model. For example, if `days-to-keep-gim-facts=400` but, on the day that `Job_MaintainGIM` is running, the earliest active interaction is an e-mail that has been left in a workbin for 550 days, then voice interactions will be retained for 550 days.
- In partitioned databases, data is purged by partition. Therefore, partitioning might further delay the purging of multimedia interactions in GIDB and the purging of voice and multimedia interactions in the dimensional model, because a partition will be dropped only when all data in that partition is eligible to be purged.

Retention Period for Discarded Operational Data

The retention period for records in the discard, audit, and history tables is calculated from the time that the ETL process discarded the data. The discard tables are Staging tables that store operational data that the transformation job was unable to process—for example, voice interaction data with unresolved IS-Links, or Configuration facts (configuration relationship records) with missing configuration objects. The audit and history tables are Control tables that store information about data lineage and about ETL-processing activity. Information in the Discard tables is useful for troubleshooting. Purging of discarded data is controlled by the `days-to-keep-discards-and-job-history` option in the `[gim-etl]` section (see page 283).

Purging Rules

The following rules apply to purging:

- Facts are evaluated for purging based on the start time of the fact.
- Both active and completed facts are eligible for purging, but they have different retention periods, controlled by separate configuration options. Similarly, the retention periods for completed facts in the dimensional model and in GIDB are different.
- For both short-living and long-living entities, records are purged independently from the applicable tables, for all records that have a start time earlier than the purge threshold.

For more information, see "Retention Periods for Facts" on page 457 and "Purge Strategies" on page 460.

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

Purge Strategies

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

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

For more information about how `Job_MaintainGIM` applies data-retention policies and defines short-living and long-living entities for purging purposes, see "Retention Periods" on page 457.

For information about how to purge old data from the Interaction Database (IDB), see the chapter about managing data sources in the *Genesys Info Mart 8.1 Operations Guide*.

Purging Fact Data from the Dimensional Model

Figure 31 on page 462 shows a purging eligibility example for Fact tables in the dimensional model. When looking at the example, keep in mind the following:

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

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

- `Job_MaintainGIM` subtracts 400 days from January 21, 2012, which is December 17, 2010. This is the retention timestamp for short-living entities, and the upper bound of the purge eligibility interval for completed long-living interactions.
- `Job_MaintainGIM` subtracts 600 days from January 21, 2012, which is June 1, 2010. This is the retention timestamp for active long-living entities.
- From the root table in the group of related tables for the particular type of data, `Job_MaintainGIM` calculates the start time of the earliest active fact in the time range between the `days-to-keep-gim-facts` and the `days-to-keep-active-facts` timestamps. For example, for multimedia interactions, the root table is `STG_ACTIVE_IF`.

- The start time of the earliest active fact between December 17, 2010 and June 1, 2010 sets the purge threshold for completed long-living entities in the group of interaction-related tables. In the example shown in Figure 31, this date is August 3, 2010.

  In all related tables, completed facts that started on August 3, 2010 or earlier are eligible to be purged.
- Active facts that started earlier than June 1, 2010 are also eligible to be purged.

In Figure 31, `Job_MaintainGIM` purges facts relating to the following entities:

Long-Living Entities Purged

- `Active Multimedia Interaction 1`, because the start date is earlier than the retention timestamp for active facts.
- `Completed Multimedia Interaction 3`, because the start date is earlier than the calculated purge threshold for completed long-living entities.
- `Voice Interaction 1` (by definition, a completed fact), because the start date is earlier than the calculated purge threshold for completed long-living entities.
- `Outbound Contact Campaign 1`, because, for the group of Outbound Contact–related tables, there are no active facts within the time range between the `days-to-keep-gim-facts` and the `days-to-keep-active-facts` timestamps, and the start date is earlier than the `days-to-keep-gim-facts` timestamp.

Short-Living Entities Purged

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

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

Long-Living Entities Not Purged

- `Active Multimedia Interaction 2`, because the start date is not earlier than the retention timestamp for active facts. However, as the earliest active interaction within the interval defined by the retention-policy options for active and completed facts, `Active Multimedia Interaction 2` sets the purge threshold for completed long-living entities.
- `Completed Multimedia Interaction 4`, because the start date is later than the calculated purge threshold for completed long-living entities.
- `Voice Interaction 2` or `Voice Interaction 3`, because the start date is later than the calculated purge threshold for completed long-living entities.
- `Outbound Contact Campaign 2`, because it is still active and the start date is not earlier than the `days-to-keep-gim-facts` and `days-to-keep-gidb-facts` timestamps.

Short-Living Entities Not Purged

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

Note: Figure 31 does not depict all Fact tables from which `Job_MaintainGIM` purges old data. For a list of the Fact tables in the dimensional model that `Job_MaintainGIM` purges, see "Dimensional Model Fact Tables" on page 465.

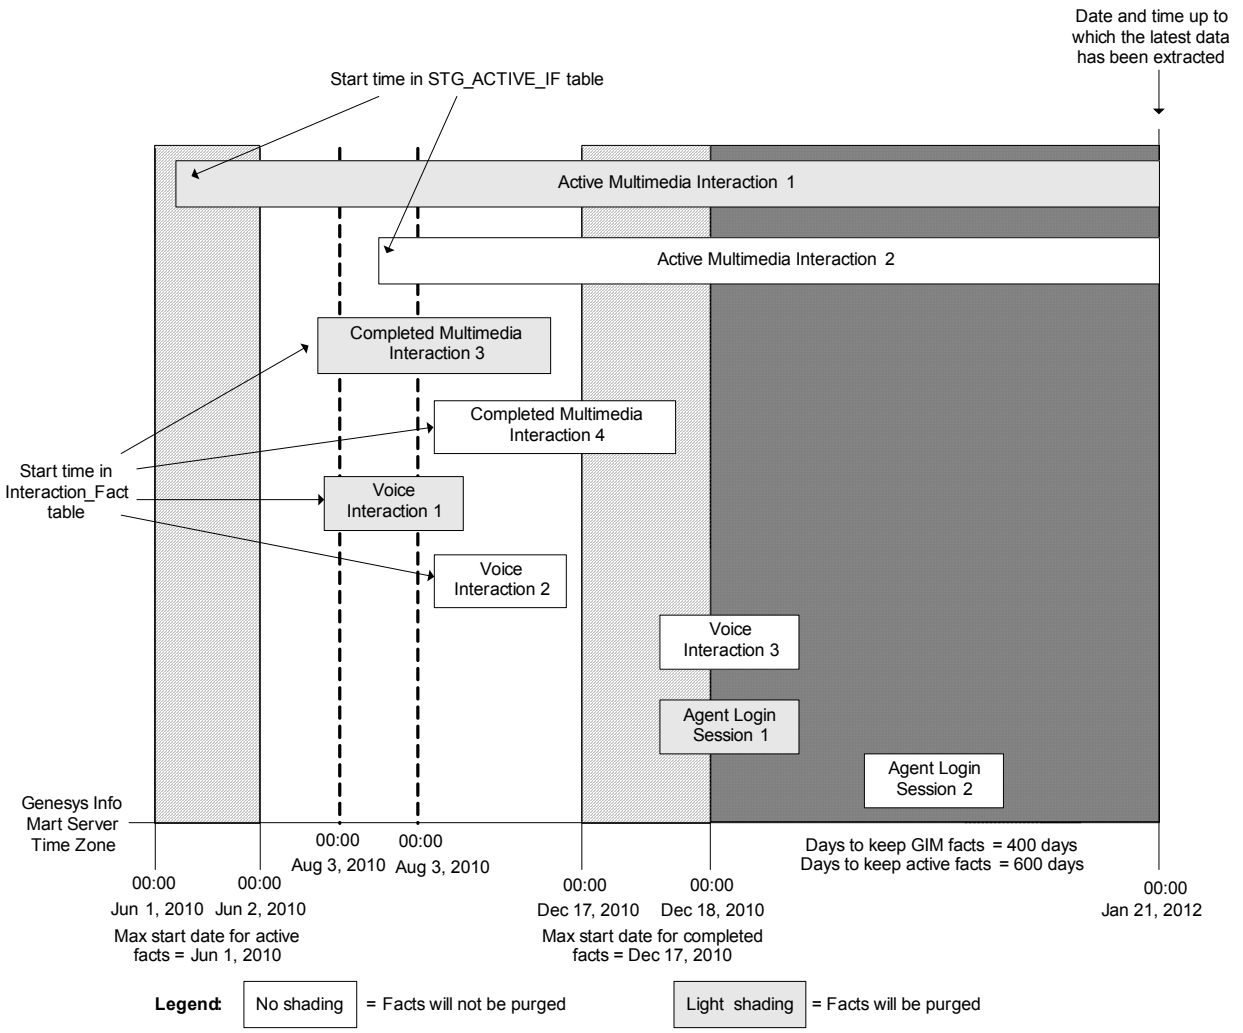

Figure 31: Purging Eligibility Example—Dimensional Model

Purging Fact Data from GIDB

Figure 32 on page 464 shows a purging eligibility example for Fact tables in GIDB. When looking at the example, keep in mind the following:

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

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

- Job_MaintainGIM subtracts 14 days from January 21, 2012, which is January 7, 2012. This is the retention timestamp for short-living entities, and the upper bound of the purge eligibility interval for completed long-living interactions.
- Job_MaintainGIM subtracts 600 days from January 21, 2012, which is June 1, 2010. This is the retention timestamp for active long-living entities.
- From the root table in the group of related tables for the particular type of data, Job_MaintainGIM calculates the start time of the earliest active fact in the time range between the days-to-keep-gidb-facts and the days-to-keep-active-facts timestamps. For example, for multimedia interactions, the root table is STG_ACTIVE_IF.
  - The start time of the earliest active fact between January 7, 2012 and June 1, 2010 sets the purge threshold for completed long-living entities in the group of interaction-related tables. In the example shown in Figure 31, this date is August 3, 2010.

    In all related tables, completed facts that started on August 3, 2010 or earlier are eligible to be purged.
  - Active facts that started earlier than June 1, 2010 are also eligible to be purged.

In Figure 32, Job_MaintainGIM purges facts relating to the following entities:

Long-Living Entities Purged

- Active Multimedia Interaction 1, because the start date is earlier than the retention timestamp for active facts.
- Completed Multimedia Interaction 3, because the start date is earlier than the calculated purge threshold for completed long-living entities.
- Outbound Contact Campaign 1, because, for the group of Outbound Contact–related tables, there are no active facts within the time range between the days-to-keep-gidb-facts and the days-to-keep-active-facts timestamps, and the start date is earlier than the days-to-keep-gidb-facts timestamp.

Short-Living Entities Purged

- Voice Interaction 1, Voice Interaction 2, and Voice Interaction 3 (by definition, completed facts), because the start date is earlier than the calculated purge threshold for completed short-living entities.
- Agent Login Session 1, because the start date is earlier than the retention timestamp for short-living entities.

Job_MaintainGIM does not purge facts relating to the following entities:

Long-Living Entities Not Purged

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

- `Completed Multimedia Interaction 4`, because the start date is later than the calculated purge threshold for completed long-living entities.
- `Outbound Contact Campaign 2`, because it is still active and the start date is not earlier than the `days-to-keep-gim-facts` and `days-to-keep-gidb-facts` timestamps.

Short-Living Entities Not Purged

- `Voice Interaction 4` because the start date is later than the calculated purge threshold for completed short-living entities.
- `Agent Login Session 2`, because the start date is later than the retention timestamp for short-living entities.

Note: Figure 32 does not depict all Fact tables from which `Job_MaintainGIM` purges old data. For a list of the Fact tables in GIDB that `Job_MaintainGIM` purges, see "GIDB Fact Tables" on page 466.

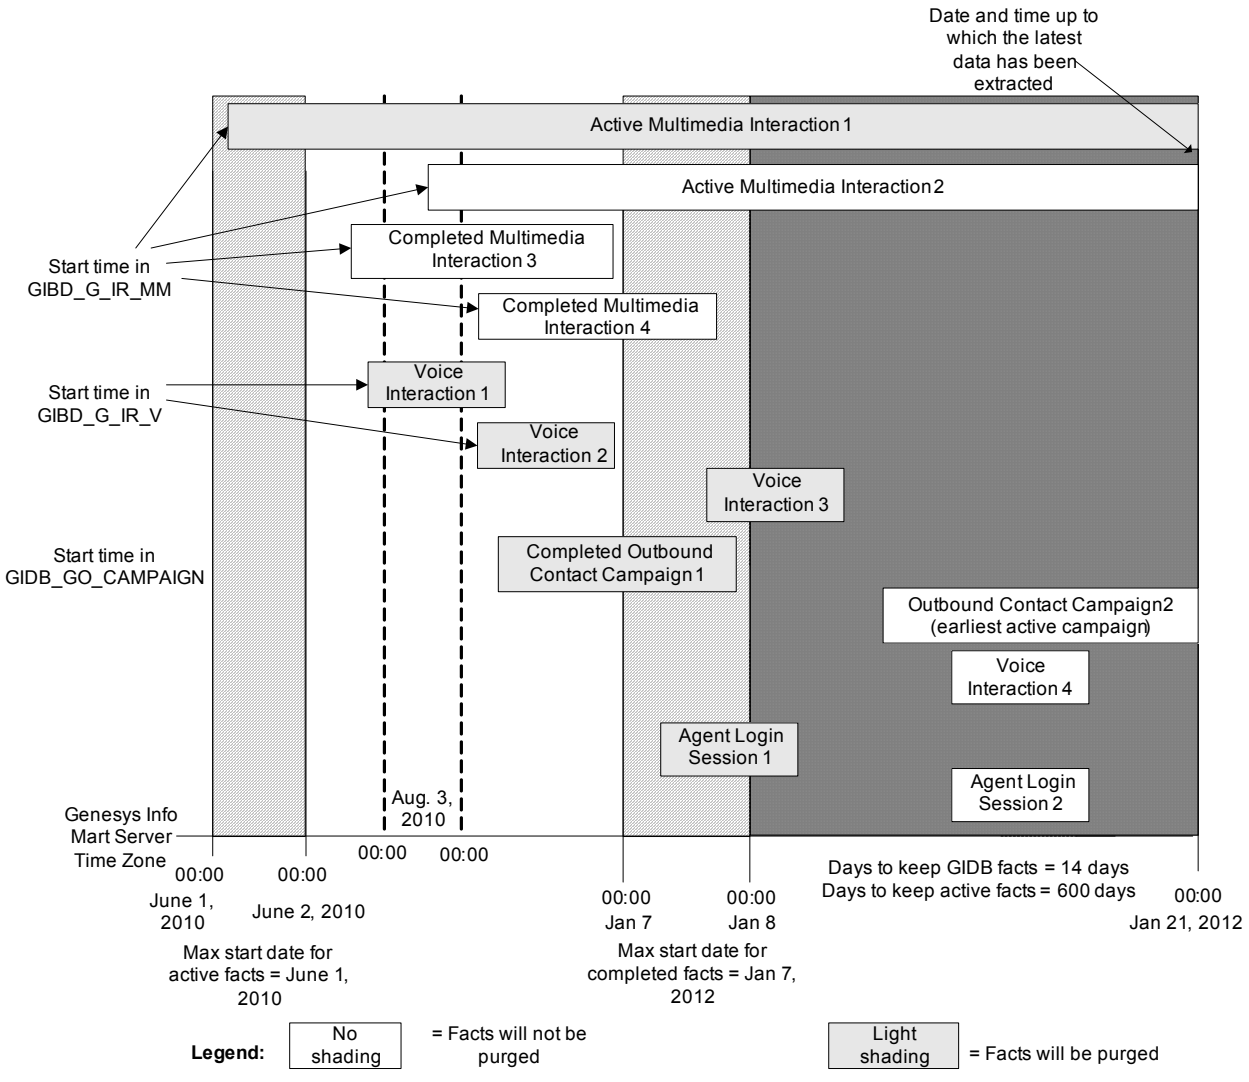

Figure 32: Purging Eligibility Example—GIDB

Info Mart Tables Purged by the Maintenance Job

[Information in this section is the 8.1.0 equivalent of material from Appendix C on page 441.]

Dimensional Model Fact Tables

Table 31 lists the fact tables in the dimensional model that are purged by `Job_MaintainGIM`.

Table 31: Dimensional Model Tables Purged by Job_MaintainGIM

Type of Data	Tables Purged	Purge Condition
Short-Living Entities		
Agent activity	• `SM_RES_SESSION_FACT` • `SM_RES_STATE_FACT` • `SM_RES_STATE_REASON_FACT`	(Start timestamp) < (`ExtractHWM` – `days-to-keep-gim-facts`)
Long-Living Entities		
Voice and Multimedia interactions	• `STG_ACTIVE_IF` • `STG_ACTIVE_IRF` • `STG_ACTIVE_MSF` • `STG_ACTIVE_IRF_REPLIES` • `INTERACTION_FACT` • `INTERACTION_RESOURCE_FACT` • `IXN_RESOURCE_STATE_FACT` • `IRF_USER_DATA_KEYS` and custom high-cardinality user-data tables (`IRF_CUST_*` in the `make_gim_UDE_template.sql` script) • `MEDIATION_SEGMENT_FACT`	(Start timestamp) < `purgeTS`, where `purgeTS` = (`ExtractHWM` – `days-to-keep-active-facts`) < (Start timestamp of the earliest active fact in the `STG_ACTIVE_IF` table) < (`ExtractHWM` – `days-to-keep-gim-facts`) If there are no active facts in the time interval between `days-to-keep-active-facts` and `days-to-keep-gim-facts`, the effective condition is: (Start timestamp) < (`ExtractHWM` – `days-to-keep-gim-facts`)

Table 31: Dimensional Model Tables Purged by Job_MaintainGIM (Continued)

Type of Data	Tables Purged	Purge Condition
Outbound Contact activity	• CAMPAIGN_GROUP_SESSION_FACT • CALLING_LIST_METRIC_FACT • CAMPAIGN_GROUP_STATE_FACT • CONTACT_ATTEMPT_FACT	(Start timestamp) < purgeTS, where purgeTS = (ExtractHWM – days-to-keep-active-facts) < (Start timestamp of the earliest active fact in the CAMPAIGN_GROUP_SESSION_FACT table) < (ExtractHWM – days-to-keep-gim-facts) If there are no active facts in the time interval between days-to-keep-active-facts and days-to-keep-gim-facts, the effective condition is: (Start timestamp) < (ExtractHWM – days-to-keep-gim-facts)

GIDB Fact Tables

Table 32 lists the fact tables in GIDB that are purged by Job_MaintainGIM.

Table 32: GIDB Tables Purged by Job_MaintainGIM

Type of Data	Tables Purged	Purge Condition
Short-Living Entities		
Voice interactions	• GIDB_G_IR_V • GIDB_G_IR_HISTORY_V • GIDB_G_CALL_V • GIDB_G_VIRTUAL_QUEUE_V • GIDB_G_CALL_HISTORY_V • GIDB_G_CALL_STAT_V • GIDB_G_ROUTE_RESULT_V • GIDB_G_SECURE_UD_HISTORY_V • GIDB_G_USERDATA_HISTORY_V • GIDB_G_CUSTOM_DATA_S_V • GIDB_G_IS_LINK_V • GIDB_G_IS_LINK_HISTORY_V • GIDB_G_PARTY_V • GIDB_G_PARTY_HISTORY_V	(Start timestamp) < (ExtractHWM – days-to-keep-gidb-facts)

Table 32: GIDB Tables Purged by Job_MaintainGIM (Continued)

Type of Data	Tables Purged	Purge Condition
Agent activity	• GIDB_G_AGENT_STATE_HISTORY_V • GIDB_G_AGENT_STATE_RC_V • GIDB_G_DND_HISTORY_V • GIDB_G_LOGIN_SESSION_V • GIDB_GX_SESSION_ENDPOINT_V • GIDB_G_AGENT_STATE_HISTORY_MM • GIDB_G_AGENT_STATE_RC_MM • GIDB_G_DND_HISTORY_MM • GIDB_G_LOGIN_SESSION_MM • GIDB_GX_SESSION_ENDPOINT_MM	(Start timestamp) < (ExtractHWM – days-to-keep-gidb-facts)
Long-Living Entities		
Multimedia interactions	• GIDB_G_IR_MM • GIDB_G_IR_HISTORY_MM • GIDB_G_CALL_MM • GIDB_G_VIRTUAL_QUEUE_MM • GIDB_G_CALL_HISTORY_MM • GIDB_G_ROUTE_RESULT_MM • GIDB_G_SECURE_UD_HISTORY_MM • GIDB_G_USERDATA_HISTORY_MM • GIDB_G_CUSTOM_DATA_S_MM • GIDB_GM_F_USERDATA • GIDB_GM_L_USERDATA • GIDB_G_PARTY_MM • GIDB_G_PARTY_HISTORY_MM	(Start timestamp) < purgeTS, where purgeTS = (ExtractHWM – days-to-keep-active-facts) < (Start timestamp of the earliest active fact in the GIDB_G_IR_MM table) < (ExtractHWM – days-to-keep-gidb-facts) If there are no active facts in the time interval between days-to-keep-active-facts and days-to-keep-gidb-facts, the effective condition is: (Start timestamp) < (ExtractHWM – days-to-keep-gidb-facts)

Table 32: GIDB Tables Purged by Job_MaintainGIM (Continued)

Type of Data	Tables Purged	Purge Condition
Outbound Contact activity	• GIDB_GO_CAMPAIGN • GIDB_GO_CAMPAIGNHISTORY • GIDB_GO_METRICS • GIDB_GO_CHAIN • GIDB_GO_CHAINREC_HIST • GIDB_GO_FIELDHIST • GIDB_GO_RECORD • GIDB_GO_SEC_FIELDHIST • GIDB_GOX_CHAIN_CALL	(Start timestamp) < purgeTS, where purgeTS = (ExtractHWM – days-to-keep-active-facts) < ((Start timestamp of the earliest active fact in the GIDB_GO_CAMPAIGN table) < (ExtractHWM – days-to-keep-gidb-facts) If there are no active facts in the time interval between days-to-keep-active-facts and days-to-keep-gidb-facts, the effective condition is: (Start timestamp) < (ExtractHWM – days-to-keep-gidb-facts)

Related Documentation Resources

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

Framework

- The *Framework 8.x Management Layer User's Guide* provides information about the concepts, terminology, and procedures that apply to this layer of the Genesys Framework.
- The *Framework 8.x Configuration Options Reference Manual* provides information about configuration options for Framework components.
- The *Framework 8.x Configuration Manager Help* provides information about using Configuration Manager in either an enterprise or a multi-tenant environment.
- The *Framework 8.x Deployment Guide* provides information about configuring, installing, starting, and stopping Framework components.
- The *Framework 8.x Combined Log Events Help* describes log events that Genesys server applications generate and that Solution Control Interface displays. The *Framework 8.x Combined Log Events Help* includes descriptions of Genesys Info Mart log events.

Interaction Concentrator

- The *Interaction Concentrator 8.x Deployment Guide* provides information about architecture, configuration requirements, and installation steps for Interaction Concentrator, and it describes how to make data from the Genesys Outbound Contact solution available in Interaction Database (IDB).

- The *Interaction Concentrator 8.x User's Guide* provides basic information about IDB architecture and detailed information about Interaction Concentrator features and functionality, including attached data processing, available stored procedures, and integration with other Genesys products.
- The *Interaction Concentrator 8.x Physical Data Model* for your relational database management system (RDBMS) provides information about the IDB schemas.

Genesys Info Mart

- The *Genesys Info Mart 8.1 Operations Guide* provides information about the Genesys Info Mart jobs. The jobs execute extract, transform, and load (ETL) processes, maintain the Info Mart database, and migrate the database schema as required. The guide also explains how to use the Genesys Info Mart Manager and Genesys Info Mart Administration Console to monitor and administer the jobs.
- The *Genesys Info Mart 8.1 User's Guide* provides information about how to use data that is stored by Genesys Info Mart for contact center historical reporting.
- The *Genesys Info Mart 8.1 Reference Manual* for your RDBMS provides information about the Info Mart database schema.
- The *Genesys Info Mart 8.1 Business Continuity Deployment Guide*, which is available as wiki pages, provides information and procedures that are relevant to Genesys Info Mart deployment in an environment that requires support for Business Continuity.
- The *Genesys Info Mart 8.1 Database Size Estimator* helps you estimate the size of your Info Mart database when you are planning your deployment. The estimator is a Microsoft Office Excel 2007 spreadsheet that is available from the Genesys Documentation website.
- The *Genesys Info Mart Database Compatibility Reference* includes compatibility information for database tables and fields that existed in the Genesys Info Mart database schema in release 7.6. This document, which is available as wiki pages, provides guidelines for mapping Info Mart 7.6 database SQL queries for use with an Info Mart 8.x database.
- Release Notes and Product Advisories for this product, which are available on the Genesys Customer Care website at http://genesys.com/customer-care.

Reporting and Analytics Aggregates

- The *Reporting and Analytics Aggregates 8.1 Deployment Guide* describes how to deploy the Reporting and Analytics Aggregates (RAA) package provided with Genesys Info Mart.

- The *Reporting and Analytics Aggregates 8.1 Reference Manual* describes the aggregate tables that are available to Genesys Info Mart customers with deployment of RAA.
- The *Reporting and Analytics Aggregates 8.1 User's Guide* describes the aggregation process, provides the aggregation hierarchies, and explains how to enable aggregation of user data.

Genesys Interactive Insights

- The *Genesys Interactive Insights 8.1 Deployment Guide* describes how to install Genesys Interactive Insights (GI2) and set up the environment required in order to run the GI2 reports.
- The *Genesys Interactive Insights 8.1 Universe Guide* describes, in detail, the reports and measures that are provided in the GI2 release.
- The *Genesys Interactive Insights 8.1 User's Guide* summarizes how to operate GI2 reports, provides basic instructions for customizing your own reports, and provides information about the report upgrade utility.

Genesys

- The *Genesys Technical Publications Glossary,* which is available on the Genesys Documentation website, provides a comprehensive list of the Genesys and computer-telephony integration (CTI) terminology and acronyms used in this document.
- The *Genesys Migration Guide*, which ships on the Genesys Documentation Library DVD, provides documented migration strategies for Genesys product releases. The Genesys Info Mart 8.x part of the guide includes instructions on how to migrate Genesys Info Mart from release 8.0.x to release 8.1. Contact Genesys Customer Care for more information.

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

- *Genesys Supported Operating Environment Reference Guide*
- *Genesys Supported Media Interfaces Reference Manual*

Consult the following additional resources as necessary:

- The *Genesys Hardware Sizing Guide* provides information about Genesys hardware sizing guidelines for the Genesys 8.x releases.
- The *Genesys Interoperability Guide* provides information on the compatibility of Genesys products with various Configuration Layer Environments; Interoperability of Reporting Templates and Solutions; and *Gplus* Adapters Interoperability.

- The *Genesys Licensing Guide* introduces you to the concepts, terminology, and procedures that are relevant to the Genesys licensing system.
- The *Genesys Database Sizing Estimator 8.x Worksheets* provides a range of expected database sizes for various Genesys products.

For additional system-wide planning tools and information, see the release-specific listings of System-Level Documents on the Genesys Documentation website (docs.genesys.com).

Genesys product documentation is available on the:

- Genesys Customer Care website at http://genesys.com/customer-care.
- Genesys Documentation site at http://docs.genesys.com/.
- Genesys Documentation Library DVD, which you can order by e-mail from Genesys Order Management at orderman@genesys.com.

Document Conventions

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

Document Version Number

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

80fr_ref_06-2008_v8.0.001.00

You will need this number when you are talking with Genesys Customer Care about this product.

Screen Captures Used in This Document

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

Type Styles

Table 33 describes and illustrates the type conventions that are used in this document.

Table 33: Type Styles

Type Style	Used For	Examples
Italic	• Document titles • Emphasis • Definitions of (or first references to) unfamiliar terms • Mathematical variables Also used to indicate placeholder text within code samples or commands, in the special case where angle brackets are a required part of the syntax (see the note about angle brackets on page 474).	Please consult the *Genesys Migration Guide* for more information. Do *not* use this value for this option. A *customary and usual* practice is one that is widely accepted and used within a particular industry or profession. The formula, $x + 1 = 7$ where x stands for . . .

Table 33: Type Styles (Continued)

Type Style	Used For	Examples
Monospace font (Looks like `teletype` or `typewriter text`)	All programming identifiers and GUI elements. This convention includes: • The *names* of directories, files, folders, configuration objects, paths, scripts, dialog boxes, options, fields, text and list boxes, operational modes, all buttons (including radio buttons), check boxes, commands, tabs, CTI events, and error messages. • The values of options. • Logical arguments and command syntax. • Code samples. Also used for any text that users must manually enter during a configuration or installation procedure, or on a command line.	Select the `Show variables on screen` check box. In the `Operand` text box, enter your formula. Click `OK` to exit the `Properties` dialog box. T-Server distributes the error messages in `EventError` events. If you select `true` for the `inbound-bsns-calls` option, all established inbound calls on a local agent are considered business calls. Enter `exit` on the command line.
Square brackets ([ ])	A particular value that is optional within a logical argument, a command, or some programming syntax. That is, the presence of the parameter or value is not required to resolve the argument, command, or block of code. The user decides whether to include this optional information.	`smcp_server -host [/flags]`
Angle brackets (< >)	A placeholder for a value that the user must specify. This might be a DN or a port number specific to your enterprise. **Note:** In some cases, angle brackets are required characters in code syntax (for example, in XML schemas). In these cases, italic text is used for placeholder values.	`smcp_server -host <confighost>`

Index

Symbols

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

Q

R

S

T

U

V

W

X