Oracle® Hyperion Financial Data Quality Management

Oracle® Hyperion Financial Data Quality Management for Hyperion Enterprise

Oracle® Hyperion Financial Data Quality Management Adapter Suite

Oracle® Hyperion Financial Data Quality Management ERP Source Adapter for SAP

Oracle® Hyperion Financial Data Quality Management for Oracle Hyperion Enterprise Planning Suite

Oracle® Hyperion Financial Data Quality Management Adapter for Financial Management

Administration Guide

Release 11.1.2.1

ORACLE

ENTERPRISE PERFORMANCE MANAGEMENT SYSTEM

FDM Administration Guide, 11.1.2.1

Copyright © 1989, 2012, Oracle and/or its affiliates. All rights reserved.

Authors: EPM Information Development Team

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS:

Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Contents

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info or visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs if you are hearing impaired.

1 About the DBA Guide and Component Integration Model

In This Chapter

About the DBA Guide

This guide is intended for use by database administrators for optimizing databases for use with Oracle Hyperion Financial Data Quality Management (FDM). The procedures and recommendations in this guide have been thoroughly tested to provide the best database performance with FDM. Failure to follow the guidelines outlined here may result in poor performance.

About Component Integration Model

Component Integration Model (CIM) is a conceptual framework that provides a modular approach to performing the complex data integration tasks that are inherent to analytical applications.

Because processes and data structures are standardized, you can create manageable projects that meet specific needs and that provide scalable and reliable platforms that can integrate into any enterprise-level data model.

Characteristics Common to CIMs

- Standard schema and file-system storage (CIM repository)
- Integrated ETL capability
- Integrated data-cleansing capability
- Integrated data verification capability
- Integrated data transformation engine
- Integrated task scheduling service

- User interface
- Process workflow
- Complete process transparency and audit ability
- Audit, activity, and performance monitoring reports
- Standard upward data certification process
- Push integration for executing calculations and evaluating data quality
- Pull integration for enabling other systems to consume data

These characteristics enable multiple CIM repositories to be combined, used as the building blocks of virtual warehouses, and linked into existing data warehouses. Because the data stored in CIM repositories is of the highest quality, data quality is measurable and sustainable.

CIM repositories are the perfect data source for analytical reporting solutions. Business analysts and information technology professionals can build independent integration solutions that meet their most detailed requirements, are easy to maintain, and fulfill their enterprise-level data integration goals. Data can be collected consistently and transferred across an organization, regardless of the business process or data flow involved in the transfer.

CIM Repository

The Hyperion CIM repository contains a standard, relational-database directory structure. The repository, used for storing documents, reports, and application server files, is referred to as an FDM application.

There is usually a one-to-one relationship between the number of CIM repositories (FDM applications) and the number of target systems. Therefore, a transformation rule set is required for each target system.

When applications are created:

- A CIM relational database is created.
- A CIM directory structure is created.
- FDM Application Manager relates the database and the directory structure and stores them in an XML configuration file.

CIM Relational-Database Areas

- Work area
- Data mart
- Pull area

Work Areas

Work areas are used by the transformation engine to stage, cleanse, and transform incoming data. Objects created in work areas are temporary. Each object name includes the prefix `TW` (for temporary working tables).

Data Mart

The data mart contains cleansed and transformed external data, metadata, log data, push-integration instruction sets, and non-partition application data. The transformation engine posts transformed data from the work area to the data mart. Within the data mart, data is calculated and re-transformed. The most recent data is pulled from the data mart into the work area, transformed and refreshed, processes are completed and refreshed, and then posted back to the data mart.

Pull Area

Pull areas contain sets of views that provide access to the cleansed data that resides in the data mart. The views consist of `Union All` statements that assemble the partitioned tables into one table.

In addition, fact table views (`vDataFacts`) provide access to all incoming and transformed values. The views relate the most commonly used tables and, thereby, provide a standard method for accessing data.

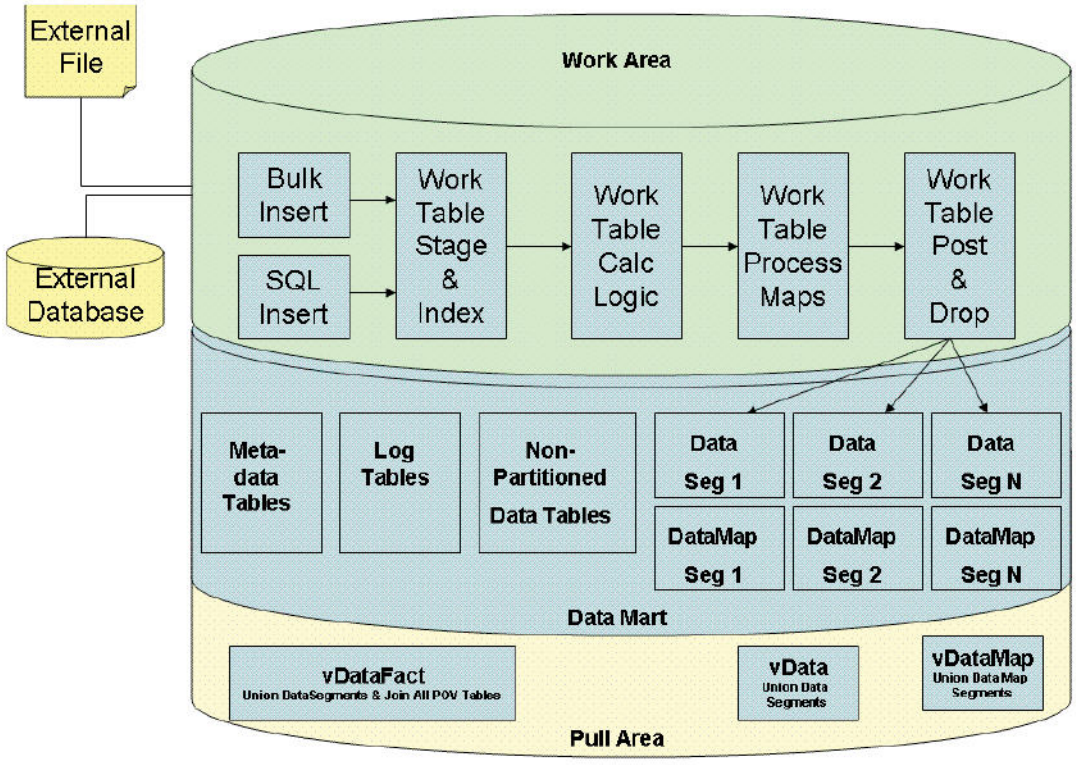

CIM Directory Structure

CIM directory structures are used to archive documents, store custom scripts, and provide space for reports-format files and application server processing (inbox, outbox, templates, and logs) and report format files. Each structure must be created on a file server that can be accessed by all application servers of the application server cluster. For SQL Server, the data server service must be located in the `Inbox` directory.

The following diagram illustrates the CIM directory structure. The `FDMDATA` directory is a user-created directory that stores FDM applications. `ABCCorp` is a system-generated directory, created when an FDM application is created, and named after the FDM application.

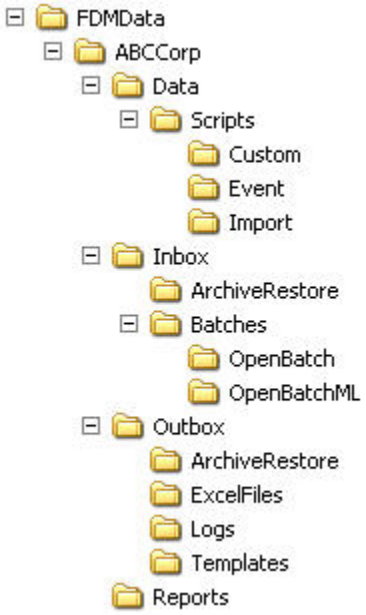

CIM Subdirectories

The data folder contains several subdirectories:

- **Data**—Document archive files (`.ua1` and `.ua2` extensions)
- **Scripts (Custom, Event, Import)**—Visual Basic (VB) script files that are accessed by the transformation engine
- **Inbox**—Incoming files
- **Archive Restore**—Incoming documents restored from the document archive location
- **Batches (OpenBatch)**—Time-stamped folders for file-harvesting or batch-processing tasks
- **Batches (OpenBatchML**—Time-stamped folders for file-harvesting or multiload batch-processing tasks
- **Outbox**—Outgoing files
- **Archive Restore**—Outgoing documents restored from the document archive location (temporary storage)

- **ExcelFiles**—Information exported to Excel
- **Logs**—Processing and error log files
- **Templates**—Excel or other document templates
- **Reports**—Report-format files

CIM Transformation Engine

The CIM transformation engine, the nucleus of the Hyperion Component Integration Model, is a comprehensive set of software libraries that are used for data staging, cleansing, and transformation. The transformation engine delivers highly reliable analytical data sources as standardized data marts.

Transformation Engine Processes

Data transformation tasks:

- Database-level integration (OLEDB/ADO cursor)
- File-based import (bulk insert or SQL Insert)

Note: Data transformation tasks, the most resource-intensive tasks executed during the load process, are the most likely processes to cause resource problems.

Database Level-Integration (OLEDB/ADO Cursor)

Sequence	Task	I/O Location	Active Server
1	A user-specific temporary table is created.	Data server (work area)	Data
2	An integration script executes a SQL `Select` statement that populates ADO record sets with source-system values. Cursor is iterated to write all source records to the user-specific temporary table.	Data server (work area)	Data, or data and application
3	An integration script is added to the document archive directory.	Data directory	Application
4	Indexes are added to the user-specific temporary table.	Data server (work area)	Data
5	The transformation engine executes all calculations and data transformation rules.	Data server (work area)	Data or data and application
6	If data is replacing data, a delete action is executed against the active data mart data-segment table.	Data server (DataMart)	Data
7	The clean and transformed user-specific temporary table data is posted into the data mart data-segment table.	Data server (work area and DataMart)	Data
8	The user-specific temporary table is deleted.	Data server (work area)	Data

File-Based Import (Bulk Insert or SQL Insert)

Sequence	Task	I/O Location	Active Server
1	A file is transferred from the Web server to the application server.	`Inbox` directory	Application
2	The transformation engine stages the source file into a clean, delimited text file which is then copied to the `Inbox` directory.	Application server `Temp` directory and `Inbox` directory	Application
3	The source file is added to the document archive directory.	`Data` directory	Application
4	A user-specific temporary table is created.	Data server (work area)	Data
5	For bulk insert, a SQL Server `Bulk Insert` statement is called, and Oracle SQL Loader is launched on the application server.	`Inbox` directory	Data (for the statement) and application and data (for Oracle SQL Loader)
	For SQL insert, the clean, delimited text file runs `SQL Insert` statements in batches of 100 statements.	`Inbox` directory	Data
6	Indexes are added to the user-specific temporary table.	Data server (work area)	Data
7	The transformation engine executes all calculations and data transformation rules.	Data server (work area)	Data, application, or data and application
8	If data is replacing data, a delete action is executed against the active data mart data-segment table.	Data server (data mart)	Data
9	The cleaned and transformed data from the user-specific temporary table is posted into the data mart data-segment table.	Data server (work area and data mart)	Data
10	The user-specific temporary table is deleted.	Data server (work area)	Data

Component Integration Model Push-Pull Integration (CIMppi)

CIM supports two types of integration techniques—push integration and pull integration.

Push Integration

Push integration involves loading data into target systems, executing calculations, and verifying the quality of target system information (by extracting and evaluating loaded and calculated values).

Because CIM push integration is that the integration instruction sets used for interacting with target systems are stored in the CIM repository, the CIM transformation engine can use the CIM repository to interact with, and remotely control, any target system.

Storing integration instruction sets in the relational repository and, thereby, enabling sophisticated integration with loosely coupled analytical systems, is a key feature of CIM .

Pull Integration

Pull integration, a more common type of integration, is implemented by allowing a target system to consume the data stored in the CIM relational repository. The standard data views defined in the pull area of the CIM relational repository simplify the process of consuming data from an FDM data mart. Transformed and cleansed data, and data-quality and workflow-status information, is readily accessible through the standard views.

2 Application Settings

In This Chapter

Data-Load Methods

Each location (data transformation profile) within a CIM repository can use one of two methods —bulk insert or SQL insert—to insert data into the work area.

Bulk Inserts

Selecting the bulk insert method enables the CIM transformation engine to engage the bulk insert utility of the RDMS. These utilities provide very fast insert capabilities but may not be able to use the disk space of the tables to which they are writing or appending. The CIM transformation engine uses bulk inserts only within work areas and only to insert into temporary tables that are subsequently deleted. The disk resource used for the work area should be monitored over time.

Considerations Specific to Oracle

The transformation engine uses the Oracle SQL Loader utility to execute an `Unrecoverable Direct-Path Insert`. This process inserts data after the high water mark on the table over time, within the work-area tablespace, disk space may be consumed.

Considerations Specific to SQL Server

When using `Bulk Insert` statements to SQL Server, the transformation engine is limited to one statement per processor on the data server. On a data server, high concurrency combined with low processor count can result in a queue of bulk-insert requests. Therefore, locations that do not import a high volume of data should be switched to use the `SQL Insert` statement load method.

For successful execution of `Bulk Insert` statements against SQL Server, the SQL Server service account must be able to read data from the file-system repository.

SQL Inserts

The SQL insert method enables the CIM transformation engine to create batches of `SQL Insert` statements. The SQL insert process is not as efficient as bulk loading, but, because transactions are smaller, it generally provides better throughput. This method also creates increased network activity between the CIM engine application servers and the database server.

Transformation Rules

The types of data transformation rules defined in FDM applications impact the distribution of the work load between the application and data servers.

Complex Logic or Derived Value Rules

In general, transformation rules that require complex logic evaluation or immediate derivation of the target value from the source value require the use of a client-side cursor. These types of rules place a greater burden on the application server and place only update responsibilities on the data server.

One-to-One, Range, or Wildcard Rules

Transformation rules that can be formulated into a `SQL Update` statement are packed by the application and sent to the data server for processing. Because of the inherent performance benefit, these types of rules are most widely used.

3 General Setup and Tuning Options

In This Chapter

Data Partitioning

Each CIM relational repository uses table partitioning to optimize data-table throughput. Because the primary duty of a CIM relational repository is to process many batch-insert processes simultaneously, table contention can become an issue. This issue is solved by horizontally partitioning the data mart tables that are subject to batch inserts and batch deletes.

Partitioned Tables

Partitioned tables are assigned a prefix of `tDataSeg` or `tDataMapSeg` and a numeric value (the partition ID number). Each location (data transformation profile) that is configured in a CIM relational repository is assigned to a data segment. The data segment identifies which data-segment tables the location uses within the data mart. When locations are created, the CIM transformation engine assigns data segment key values to them.

Data-Segment Count

You can adjust the number of data segments by changing the configuration option "Total No. of Data Segments." This option is set, by default, at 50 segments, the optimal value based on stress testing of 500 concurrent data loads of 20,000 records. At this high level of concurrent batch loads, 50 segments provided good throughput and no deadlocking.

After a CIM relational repository is used to load data, the "Total No. of Data Segments" configuration options can only be increased. To decrease the data-segment count, all segment tables must be dropped and recreated. This process results in the loss of all segment data.

RDMS Disk I/O Optimization

Each CIM relational repository can use up to five RDMS disk I/O resources, two in the work area and three in the data mart.

Note: See Chapter 4, "Working with Oracle Server" and Chapter 5, "Working with SQL Server" for detail on the options for RDMS.

Work-Area Disk Resources

During the data-staging process, the work area supports the use of two disk resources, the first for the staging tables and the second for the indexes created against the staging tables. However, stress testing indicates that splitting the table and index I/O to different resources may increase overhead. Therefore, using one disk resource for work tables and work table indexes is recommended.

Server	Option Key — Work-Area Resource	Default Value
Oracle	ora_WorkTableSpace	Users
	ora_WorkIXTableSpace	Users
SQL Server	`FileGroupWorkTable`	Primary
	`FileGroupWorkTableIndex`	Primary

DataMart Disk Resources

The data mart supports the use of three disk resources: one for the main data-segment tables; one for the data-map-segment tables; and one, the default, for all other tables and indexes. When the CIM repository is created, all objects are created on the default disk resource. To optimize the use of disk resources, you must change the default options and drop and re-create the data-segment tables.

Server	Option Key — DataMart Disk Resources	Default Value
Oracle	ora_DSegTableSpace	Users
	ora_DMSegTableSpace	Users
SQL Server	`FileGroupDataSeg`	Primary
	`FileGroupDataMapSeg`	Primary

4 Working with Oracle Server

In This Chapter

Recommendations for Oracle Configurations

- Allocate a minimum of 1 GB of memory for the database instance.
- Configure Oracle initialization parameters:
  - `log_buffer`
  - `open_cursors`
  - `cursor_sharing`
- Separate redo logs from other database files.
- Create separate tablespaces for Data Seg, Data Map Seg, and work tables.
- Configure the work table tablespace with no logging.
- As a minimum, use a tablespace size of 1 GB. The tablespace size requirement is dependant on the amount of data and the size of the applications.
- Set the Configuration Settings for the Oracle tablespaces as described in "Oracle Initialization Parameters" on page 20.
- Set the work table bitmap index switch to Off for Oracle 10g and Oracle 11g.

Note: Because indexing may result in a significant performance decrease, FDM data tables are not indexed.

Oracle Initialization Parameters

Refer to Appendix A, "Oracle Initialization Parameters" for information regarding Oracle 10g and Oracle 11g initialization parameters.

Oracle Database Instance

Multiple FDM application schemas can reside in one database instance.

Size of the Redo Log Buffer

The default value for size of the redo log buffer is operating system-specific, but, for most systems, it is 500Kb. If buffer size is increased, the frequency of I/O operations is reduced, and the performance of the Oracle server is increased.

If the log buffer is too small, the log writer process is excessively busy. In this case, the log writer process is constantly writing to disk and may not have enough space for its redo entries. You are more likely to encounter problems with buffers that are too small than with buffers that are too large.

To set the redo log buffer size manually, you must configure the `log_buffer` initialization parameter of the `init.ora` file. A size between 1MB and 7MB is optimal.

Open Cursors

The `open_cursors` initialization parameter sets the maximum number of cursors that each session can have open. If `open_cursors` is set to 100, for example, each session can have up to 100 cursors open at one time. If a session with 100 open cursors attempts to open a session, it encounters `ORA-1000 error – "Maximum open cursors exceeded."`

The default value for `open_cursors` is 50, but Oracle recommends that, for most applications, the value be set to a minimum of 500.

Cursor Sharing

Because FDM does not use bind variables, Oracle recommends setting the `cursor_sharing` initialization parameter to Similar instead of the default value of Exact. Cursor sharing is an auto-binder. It forces the database to rewrite queries (by using bind variables) before parsing queries.

Note: Under some instances, the default value of `cursor_sharing` may be more appropriate. Consult Oracle Support if you are unsure about the best cursor sharing setting for your installation.

Optimizing RDMS Disk I/O

When a FDM application is created, if no tablespace is specified, all FDM Oracle objects are created in the `Users` tablespace. When large amounts of data are processed, use of one tablespace may hinder I/O performance.

Redo Logs

An easy way to ensure redo logs do not cause I/O performance issues is to separate them from other database files. For example, because the I/O activity of redo logs is typically sequential and the I/O activity of data files is random, separating redo logs from data files improves performance. When you separate redo logs you should place them on the fastest devices available.

If redo logs are too few or too small, relative to the DML activity of the database, the archiver process must work extensively to archive the filled redo log files. Therefore, you must ensure that the redo logs are large enough to avoid additional checkpointing. Oracle recommends that you size the log files so that files switch every twenty minutes.

Working with Data-Segment Table Tablespaces

To minimize disk contention, you can create a tablespace for the data-segment tables and store the data files for the tablespace on a separate physical disk. After creating the tablespace, you must change the Oracle Data Map Seg TableSpace Name and Oracle Data Seg TableSpace name to the new tablespace name.

➤ To rename tablespaces:

1. Launch Workbench, and log on to the FDM application.
2. Select **Tools > Configuration Settings**.
3. Select **Options > Oracle Data Map Seg TableSpace Name**.
4. In **Name**, enter the name of the new tablespace, and click **Save**.
5. Select **Options > Oracle Data Seg TableSpace Name**.
6. In **Name**, enter the name that you entered in step 4, and click **Save**.
7. Click **Close**.

Note: You can separate the data-map and data-seg tables into their own tablespaces, but no significant increase in performance can be expected.

After the tablespace names are specified, delete the data-map and data-seg tables from the `Users` tablespace and recreate them in the new tablespace.

Note: You do not need to delete and re-create the data-map and data-seg tables if the new tablespace names were specified on the DB Options dialog box when the application was created.

Caution! Deleting and recreating the data-map and data-seg tables truncates all table data. You should change tablespace names and re-create the tables after the application is created and before data is loaded.

➤ To re-create the data-map and data-seg tables:

1. **Launch Workbench, and log on to the FDM application.**
2. **Select Tools > Manage Data Segments > Delete, Recreate, and Reassign All Segments.**

   The Recreate Segments screen is displayed.
3. **Select the number of segments to create (default is 50) and click Save.**
4. **Click Yes to verify that all data in the tables should be deleted.**

After the data-map and data-seg tables are re-created, they are located in the tablespaces specified under Configuration Settings.

Working with Tablespaces for Work Tables and Work Table Indexes

To minimize disk contention and logging, you can create a tablespace with NoLogging for the work tables and indexes and store the data files for the tablespace on a separate physical disk. For example, consider the following command, which creates the `HyperionWORK` tablespace with NoLogging:

```
CREATE TABLESPACE HyperionWORK DATAFILE
 'H:\ORACLE\ORADATA\HyperionWORK.ORA ' SIZE 5120M AUTOEXTEND ON NEXT 100M MAXSIZE
UNLIMITED
NOLOGGING
ONLINE
PERMANENT
EXTENT MANAGEMENT LOCAL UNIFORM SIZE 1M
BLOCKSIZE 8Kb
SEGMENT SPACE MANAGEMENT AUTO;
```

Because work tables are created and dropped during data processing, creating a tablespace without logging for work tables and work table indexes can improve performance. After a tablespace without logging is created, `Oracle Work TableSpaceName` and `Oracle Work Table Index TableSpaceName` must be changed to the new tablespace name.

➤ To modify the Oracle Work `TablespaceName` and Oracle Work Table Index `TableSpaceName` configuration settings:

1. **Launch Workbench, and log on to the FDM application.**

2 Select **Tools > Configuration Settings.**

3 Select **Options > Oracle Work TableSpaceName**.

4 In **Name**, enter the name of the tablespace, and click **Save**.

5 Select **Options > Oracle Work Table Index TableSpaceName**.

6 In **Name**, enter the name you entered in step 4, and click **Save**.

7 Select **Options > Oracle Work Table Bitmap Index Switch,** and set the value to Off for Oracle 10g and Oracle 11g.

8 Click **Save**.

9 Click **Close**.

All work tables and indexes that are created and dropped during data processing are now located in the new tablespace.

Note: You can separate work tables and indexes into their own tablespaces, but no significant increase in performance can be expected.

Optimizing Other Tables

All other Oracle objects created and used by FDM are stored in the `Users` tablespace. To improve performance and reduce disk contention, you can separate the `Users` tablespace from other database files and move it to a separate disk.

Account Permissions for Oracle Server

FDM uses the FDM Oracle account to access the FDM Oracle database. FDM can use Windows Integrated Security or the Oracle account that you specify. If FDM is accessed from the Web and Windows Integrated Security is used to access the Oracle database, the Application Server account is used to log on to the Oracle database. If Workbench is used, the user name that you used to log on to Workbench is used to log on to the Oracle database.

You can connect through Windows Integrated Security only if Oracle is configured to enable such connections. By default, the `sqlnet.ora` file contains the entry that enables operating system authentication. The `SQLNET.AUTHENTICATION_SERVICES= (NTS)` entry enables authentication by the operating system.

To create an Oracle account that can connect using Windows Integrated Security, you must know the value of the `os_authent_prefix` parameter. Oracle uses this parameter when it authenticates external users. The value of this parameter is prefixed to the operating system user name. The default value is `OPS$`, but the value may be different on your system. If the value is `OPS$`, the Oracle account is formatted as `OPS$hostname\username`, where `hostname` is the machine name or domain name, and `username` is the Windows user name.

New FDM accounts must be granted the DBA role or the following system privileges:

- `CREATE PROCEDURE`

- `CREATE SEQUENCE`
- `CREATE SESSION`
- `CREATE TABLE`
- `CREATE TRIGGER`
- `CREATE VIEW`
- CREATE DATABASE LINK (only required when using the ERPI FIN-B adapter)

The default tablespace used by FDM is the Users tablespace. Oracle recommends creating a new tablespace for Default. The account should have an appropriate quota set on each tablespace used to allow for future data growth. If you want to ensure that the user does not exceed a space-used threshold or if you have any questions about the appropriate value for the quota, consult the database administrator.

Client Software Requirements for Oracle

- For application servers, Oracle database utilities (including SQL*Loader) and Oracle Windows interfaces (including Oracle Provider for OLE DB) are required.
- For load balancers, Oracle Windows interfaces (including Oracle Provider for OLE DB) are required.
- Use the 32-bit Oracle Client only (even when using a 64-bit operating system).

NLS_LANG Settings

NLS_LANG is used to indicate to Oracle what character set your client's OS is using. Using this information, Oracle can perform, if necessary, a conversion from the client's character set to the database character set. Setting NLS_LANG to the character set of the database may be correct but is not always correct. Do not assume that NLS_LANG must always be the same as the database character set. Oracle recommends that you use the AL32UTF8 character set. This is based on information found in the Oracle document titled *Doc ID: 158577.1 NLS_LANG Explained (How does Client-Server Character Conversion Work?)*. You cannot change the character set of your client by using a different NLS_LANG setting.

5 Working with SQL Server

In This Chapter

Optimizing RDMS I/O

When a FDM application is created, by default all SQL objects are created in the primary file group. Usually, the primary file group works well, but, when large amounts of data are processed, disk contention may hinder I/O performance.

Working with Data-Segment Tables

To minimize disk contention, you can create a file group for the data-segment tables and store the data files for the group on a separate physical disk. After creating the new group, you must, within the `Data Map Seg Table` file group and `Data Seg Table` file group configuration settings, change from the primary file group name to the new file group name.

➤ To change the name of the `Data Map Seg Table` and `Data Seg Table` file group:

1. Launch Workbench and log on to the FDM application.
2. Select **Tools > Configuration Settings**.
3. From **Options**, select **Data Map Seg Table File.**
4. Enter a name, and click **Save.**
5. Click **Close.**

Note: The data-map and data-seg tables can be separated into two file groups, but, during testing, no significant increase in performance was observed.

After the file group name is specified, the data-map and data-seg tables must be deleted from the primary file group and re-created in the new file group.

Deleting and re-creating the data-map and data-seg tables truncates all data of the tables. After the application is created and before data is loaded, the file group names should be changed and the tables should be re-created.

➤ To recreate the data-map and data-seg tables:

1 Launch Workbench, and log on to the FDM application.

2 Select **Tools > Manage Data Segments > Delete, Recreate, and Reassign All Segments.**

The **Recreate Segments** screen is displayed.

3 Select the number of segments to create (default is 50 segments) and click **Save.**

4 Click **Yes** to verify that all data should be deleted.

The re-created data-map and data-seg tables are located in the file groups specified under **Configuration Settings**.

Working with Work Tables and Work Table Indexes

To minimize disk contention, you can create a file group for the work tables and work table indexes and store the data files for the file group on a separate physical disk. After creating the file group, within the configuration settings, change from the primary file group name to the new file group name.

➤ To change the name of the work table and work table index file group:

1 Launch Workbench, and log on to the FDM application.

2 Select **Tools > Configuration Settings.**

3 Select **Options > Work Table File Group**.

4 In **Name**, enter a name, and click **Save.**

5 Select **Options > Work Table Index File Group**.

6 In **Name**, enter the name that you entered in step 4, and click **Save.**

7 Click **Close.**

All work tables and indexes that are created and dropped during data processing are located in the new file group.

Note: Work tables and indexes can be separated, but no significant increase in performance should be expected.

Account Permissions for SQL Server

To access the SQL Server database, FDM uses the FDM SQL Server account. When accessing the database, FDM can use Windows Integrated Security or a specified SQL Server account.

When FDM is accessed from the Web, and Windows Integrated Security is used, the FDM Application Server account is used to log on to the SQL Server database. When the Workbench client is used, the user name used to log on to Workbench is used to log on to the SQL Server database.

The account used to create a database must have SQL Server system administrator or database-creator and bulk-insert administrator rights. After the database is created, the account can be limited to bulk-insert administrator and db-owner rights. The account used for running the MSSQLServer Windows service must have read access to the FDM `Inbox` folder.

Client Software Requirements for SQL Server

SQL Server requires the SQL Native Client driver or Microsoft OLE DB provider.

Collation

Oracle Hyperion Financial Data Quality Management supports only non-case sensitive collations.

A Oracle Initialization Parameters

In This Appendix

Oracle 10g and Oracle 11g

This table details the Oracle 10g initialization parameters that were used during product testing. Use these parameters for Oracle 10g and 11g. For 11g, we recommend that the customer uses the Oracle auto tuning.

Name	Value	Description
O7_DICTIONARY_ACCESSIBILITY	FALSE	No Version 7 Dictionary Accessibility Support
active_instance_count		Number of active instances in the cluster database
aq_tm_processes	0	Number of AQ time managers to be started
archive_lag_target	0	Maximum number of seconds of redos that the standby can lose
asm_diskgroups		Disk groups to mount automatically
asm_diskstring		Disk set locations for discovery
asm_power_limit	1	Number of processes for disk rebalancing
audit_file_dest	E:\ORACLE\PRODUCT\10.2.0\ADMIN\STRESS10\ADUMP	Directory in which auditing files reside
audit_sys_operations	FALSE	No system auditing
audit_trail	NONE	No system auditing
background_core_dump	partial	Core Size for Background Processes
background_dump_dest	E:\ORACLE\PRODUCT\10.2.0\ADMIN\STRESS10\BDUMP	Detached process dump directory
backup_tape_io_slaves	FALSE	No backup tape I/O secondaries
bitmap_merge_area_size	1048576	For BITMAP MERGE, maximum amount of memory
blank_trimming	FALSE	Blank trimming semantics not allowed

Name	Value	Description
buffer_pool_keep		Number of database blocks (latches) in keep the buffer pool
buffer_pool_recycle		Number of database blocks (latches) in the recycle buffer pool
circuits		Maximum number of circuits
cluster_database	FALSE	Startup not in cluster database mode
cluster_database_instances	1	Number of instances to use for sizing cluster DB SGA structures
cluster_interconnects		Interconnects for RAC use
commit_point_strength	1	Bias that this node has toward not preparing in a two-phase commit
commit_write		Write behavior of the transaction commit log
compatible	10.2.0.1.0	Database completely compatible with this software version
control_file_record_keep_time	7	Number of days that file records are kept
control_files	E:\ORACLE\PRODUCT\10.2.0\ORADATA\STRESS10\CONTROL01.CTL, E:\ORACLE\PRODUCT\10.2.0\ORADATA\STRESS10\CONTROL02.CTL, E:\ORACLE\PRODUCT\10.2.0\ORADATA\STRESS10\CONTROL03.CTL	List of control file names
core_dump_dest	E:\ORACLE\PRODUCT\10.2.0\ADMIN\STRESS10\CDUMP	Core dump directory
cpu_count	4	Number of CPUs for this instance
create_bitmap_area_size	8388608	For the bitmap index, the size of the create bitmap buffer
create_stored_outlines		For DML statements, stored outlines created
cursor_sharing	SIMILAR	Cursor sharing mode
cursor_space_for_time	FALSE	Use of more memory to obtain faster execution, not allowed
db_16k_cache_size	0	Size of cache for 16Kb buffers
db_2k_cache_size	0	Size of cache for 2Kb buffers
db_32k_cache_size	0	Size of cache for 32Kb buffers
db_4k_cache_size	0	Size of cache for 4Kb buffers

Name	Value	Description
db_8k_cache_size	0	Size of cache for 8Kb buffers
db_block_buffers	0	Number of database blocks cached in memory
db_block_checking	FALSE	No header checking or data and index block checking
db_block_checksum	TRUE	Checksum stored in DB blocks and evaluated during reads
db_block_size	8192	Size of database block in bytes
db_cache_advice	ON	Buffer cache sizing advisory
db_cache_size	0	Size of DEFAULT buffer pool for standard block size buffers
db_create_file_dest		Default database location
db_create_online_log_dest_1		Online log (controlfile) destination #1
db_create_online_log_dest_2		Online log (controlfile) destination #2
db_create_online_log_dest_3		Online log (controlfile) destination #3
db_create_online_log_dest_4		Online log (controlfile) destination #4
db_create_online_log_dest_5		Online log (controlfile) destination #5
db_domain		Directory part of the global database name, stored with CREATE DATABASE
db_file_multiblock_read_count	8	Number of DB blocks to be read each IO
db_file_name_convert		Data file name convert patterns and strings for the standby (clone) DB
db_files	200	Maximum allowable number of DB files
db_flashback_retention_target	1440	Maximum Flashback Database log retention time (in minutes)
db_keep_cache_size	0	Size of KEEP buffer pool for standard block size buffers
db_name	Stress10	Database name specified in CREATE DATABASE
db_recovery_file_dest	E:\oracle\product\10.2.0/flash_recovery_area	Default database recovery file location
db_recovery_file_dest_size	2147483648	Database recovery file size limit
db_recycle_cache_size	0	Size of RECYCLE buffer pool for standard block size buffers
db_unique_name	Stress10	Database unique name

Name	Value	Description
db_writer_processes	2	Number of background database writer processes to start
dbwr_io_slaves	0	DBWR I/O secondaries
ddl_wait_for_locks	FALSE	NOWAIT DML lock acquisitions disabled
dg_broker_config_file1	E:\ORACLE\PRODUCT\10.2.0\DB_1\DATABASE\DR1STRESS10.DAT	Data guard broker configuration file #1
dg_broker_config_file2	E:\ORACLE\PRODUCT\10.2.0\DB_1\DATABASE\DR2STRESS10.DAT	Data guard broker configuration file #2
dg_broker_start	FALSE	Data Guard broker framework (DMON process) not to be started
disk_asynch_io	TRUE	Asynch I/O used for random access devices
dispatchers	(PROTOCOL=TCP) SERVICE=Stress10XDB)	Specifications of dispatchers
distributed_lock_timeout	60	Number of seconds that a distributed transaction waits for a lock
dml_locks	748	DML locks—one for each table modified in a transaction
drs_start	FALSE	DG Broker monitor (DMON process) not to be started
event		Debug event control—default null string
fal_client		FAL client
fal_server		FAL server list
fast_start_io_target	0	Upper bound on recovery reads
fast_start_mttr_target	0	MTTR target of forward crash recovery (in seconds)
fast_start_parallel_rollback	LOW	Maximum number of parallel recovery secondaries that can be used
file_mapping	FALSE	No file mapping
fileio_network_adapters		Network adapters for file I/O
filesystemio_options		IO operations on file system files
fixed_date		Fixed SYSDATE value
gc_files_to_locks		Mapping between file numbers and global cache locks

Name	Value	Description
gcs_server_processes	0	Number of background GCS server processes to start
global_context_pool_size		Global Application Context Pool size (in bytes)
global_names	FALSE	Database links and the remote database not required to have same name
hash_area_size	131072	Size of in-memory hash work area
hi_shared_memory_address	0	SGA starting address (high order 32-bit on 64-bit platforms)
hs_autoregister	TRUE	Automatic server DD updates in HS agent self-registration enabled
ifile		File included in init.ora
instance_groups		List of instance group names
instance_name	stress10	Instance name supported by the instance
instance_number	0	Instance number
instance_type	RDBMS	Type of instance to be executed
java_max_sessionspace_size	0	Maximum allowed size in bytes of a Java sessionspace
java_pool_size	0	Size (in bytes) of Java pool
java_soft_sessionspace_limit	0	Warning limit on size (in bytes) of a Java sessionspace
job_queue_processes	10	Number of job queue secondary processes
large_pool_size	0	Size in bytes of large pool
ldap_directory_access	NONE	No RDBMS LDAP access
license_max_sessions	0	Maximum number of non-system user sessions allowed
license_max_users	0	Maximum number of named users that can be created in the database
license_sessions_warning	0	Warning level for number of non-system user sessions
local_listener		Local listener
lock_name_space		Lock name space used for generating lock names for standby (clone)
lock_sga	FALSE	SGA not locked in physical memory

Name	Value	Description
log_archive_config		Log archive configuration parameter
log_archive_dest		Archival destination text string
log_archive_dest_1		Archival destination #1 text string
log_archive_dest_10		Archival destination #10 text string
log_archive_dest_2		Archival destination #2 text string
log_archive_dest_3		Archival destination #3 text string
log_archive_dest_4		Archival destination #4 text string
log_archive_dest_5		Archival destination #5 text string
log_archive_dest_6		Archival destination #6 text string
log_archive_dest_7		Archival destination #7 text string
log_archive_dest_8		Archival destination #8 text string
log_archive_dest_9		Archival destination #9 text string
log_archive_dest_state_1	enable	Archival destination #1 state text string
log_archive_dest_state_10	enable	Archival destination #10 state text string
log_archive_dest_state_2	enable	Archival destination #2 state text string
log_archive_dest_state_3	enable	Archival destination #3 state text string
log_archive_dest_state_4	enable	Archival destination #4 state text string
log_archive_dest_state_5	enable	Archival destination #5 state text string
log_archive_dest_state_6	enable	Archival destination #6 state text string
log_archive_dest_state_7	enable	Archival destination #7 state text string
log_archive_dest_state_8	enable	Archival destination #8 state text string
log_archive_dest_state_9	enable	Archival destination #9 state text string
log_archive_duplex_dest		Duplex archival destination text string
log_archive_format	ARC%S_%R.%T	Archival destination format
log_archive_local_first	TRUE	EXPEDITE attribute default value established
log_archive_max_processes	2	Maximum number of active ARCH processes
log_archive_min_succeed_dest	1	Minimum number of archive destinations that must succeed
log_archive_start	FALSE	Archival process not to start on SGA initialization

Name	Value	Description
log_archive_trace	0	Archivelog operation tracing level established
log_buffer	7024640	Reset circular buffer size
log_checkpoint_interval	0	For checkpoint threshold, the number redo blocks
log_checkpoint_timeout	1800	Maximum time interval between checkpoints (in seconds)
log_checkpoints_to_alert	FALSE	Checkpoint begin and end not to be logged to the alert file
log_file_name_convert		Log file name convert patterns and strings for standby (clone) DB
logmnr_max_persistent_ sessions	1	Maximum number of threads to mine
max_commit_propagation_delay	0	Maximum age of new snapshot (in .01 second intervals)
max_dispatchers		Maximum number of dispatchers
max_dump_file_size	UNLIMITED	Maximum size (in blocks) of the dump file
max_enabled_roles	150	Maximum number of roles that a user can have enabled
max_shared_servers		Maximum number of shared servers
object_cache_max_size_percent	10	For the user session object cache, percentage by which maximum size can exceed optimal size
object_cache_optimal_size	102400	Optimal size of the user session's object cache in bytes
olap_page_pool_size	0	Size of the OLAP page pool in bytes
open_cursors	300	Maximum number of cursors per session
open_links	4	Maximum number of open links per session
open_links_per_instance	4	Maximum number of open links per instance
optimizer_dynamic_sampling	2	Optimization of dynamic sampling
optimizer_features_enable	10.2.0.1	Optimization of plan compatibility parameter
optimizer_index_caching	0	Optimization of percent index caching
optimizer_index_cost_adj	100	Optimization of index cost adjustment
optimizer_mode	ALL_ROWS	Optimizer mode
optimizer_secure_view_merging	TRUE	Optimization of secure view merging and predicate pushdown (movearound)

Name	Value	Description
os_authent_prefix	OPS$	Prefix for auto-logon accounts
os_roles	FALSE	Roles not received from the operating system
parallel_adaptive_multi_user	TRUE	For multiple user streams, adaptive setting of degree enabled
parallel_automatic_tuning	FALSE	For parallel execution parameters, disabled intelligent defaults
parallel_execution_message_size	2148	Message buffer size for parallel execution
parallel_instance_group		Instance group to use for all parallel operations
parallel_max_servers	80	Maximum number of parallel query servers per instance
parallel_min_percent	0	Minimum percent of threads required for parallel query
parallel_min_servers	0	Minimum number of parallel query servers per instance
parallel_server	FALSE	Startup in parallel server mode disabled
parallel_server_instances	1	Number of instances to use for sizing OPS SGA structures
parallel_threads_per_cpu	2	Number of parallel execution threads per CPU
pga_aggregate_target	419430400	Target size for the aggregate PGA memory consumed by the instance
plsql_ccflags		PL/SQL ccflags
plsql_code_type	INTERPRETED	PL/SQL codetype
plsql_compiler_flags	INTERPRETED, NON_DEBUG	PL/SQL compiler flags
plsql_debug	FALSE	PL/SQL debug
plsql_native_library_dir		PL/SQL native library directory
plsql_native_library_subdir_count	0	For PL/SQL native library, the number of subdirectories
plsql_optimize_level	2	PL/SQL optimize level
plsql_v2_compatibility	FALSE	PL/SQL version 2.x not compatible
plsql_warnings	DISABLE:ALL	PL/SQL compiler warnings disabled
pre_page_sga	FALSE	No pre-page SGA for process
processes	150	User processes

Name	Value	Description
query_rewrite_enabled	TRUE	Query rewrites disabled
query_rewrite_integrity	enforced	Requirement to perform rewrite using materialized views with preferred integrity
rdbms_server_dn		RDBMS distinguished name
read_only_open_delayed	FALSE	If TRUE, delay opening of read only files until first access
recovery_parallelism	0	Number of server processes to use for parallel recovery
recyclebin	off	No recycle bin processing
remote_archive_enable	true	Remote archival enabled
remote_dependencies_mode	TIMESTAMP	Remote Procedure Call dependencies mode parameter
remote_listener		Remote listener
remote_login_passwordfile	EXCLUSIVE	Password file usage parameter
remote_os_authent	FALSE	Non-secure remote clients not allowed to use auto-logon accounts
remote_os_roles	FALSE	Non-secure remote clients not allowed to use OS roles
replication_dependency_tracking	TRUE	Dependency for replication parallel propagation tracked
resource_limit	FALSE	No resource limit
resource_manager_plan		Resource manager top plan
resumable_timeout	0	Resumable timeout setting
rollback_segments		Segment list returned to previous state
serial_reuse	disable	Frame segments not reused
service_names	Stress10	Service names supported by the instance
session_cached_cursors	20	Number of cursors cached in session.
session_max_open_files	10	Maximum number of open files allowed per session
sessions	170	Number of user and system sessions
sga_max_size	1660944384	Maximum total SGA size
sga_target	1048576000	Target size of SGA
shadow_core_dump	partial	Core size for shadow processes

Name	Value	Description
shared_memory_address	0	SGA starting address (low order 32-bit on 64-bit platforms)
shared_pool_reserved_size	5452595	Size (in bytes) of the reserved area of shared pool
shared_pool_size	0	Size (in bytes) of shared pool
shared_server_sessions		Maximum number of shared server sessions
shared_servers	1	Number of shared servers to be started
skip_unusable_indexes	TRUE	Unusable indexes skipped
smtp_out_server		utl_smtp server and port configuration parameter
sort_area_retained_size	0	Size of in-memory sort work area that is retained between fetch calls
sort_area_size	65536	Size of in-memory sort work area
spfile	E:\ORACLE\PRODUCT\10.2.0\DB_1\DBS\SPFILESTRESS10.ORA	Server parameter file
sql92_security	FALSE	Select privilege for searched update and delete not required
sql_trace	FALSE	No SQL trace
sql_version	NATIVE	SQL language version parameter for compatibility issues
sqltune_category	DEFAULT	Category qualifier for applying hintsets
standby_archive_dest	%ORACLE_HOME%\RDBMS	Destination text string for the standby database archive log
standby_file_management	MANUAL	On standby, if AUTO then files created or dropped manually
star_transformation_enabled	FALSE	Star transformation not used
statistics_level	TYPICAL	Statistics level
streams_pool_size	0	Size (in bytes) of the streams pool
tape_asynch_io	TRUE	For tape devices, asynch I/O requests used
thread	0	Redo thread to mount
timed_os_statistics	0	Internal OS statistic gathering interval in seconds
timed_statistics	TRUE	Internal timing statistics maintained
trace_enabled	FALSE	No KST tracing

Name	Value	Description
tracefile_identifier		Trace file custom identifier
transactions	187	Maximum. number of concurrent active transactions
transactions_per_rollback_segment	5	Number of active transactions per rollback segment
undo_management	AUTO	Instance run in SMU mode
undo_retention	900	Undo retention (in seconds)
undo_tablespace	UNDOTBS1	Undo tablespace used
use_indirect_data_buffers	FALSE	No indirect data buffers (very large SGA on 32-bit platforms)
user_dump_dest	E:\ORACLE\PRODUCT\10.2.0\ADMIN\STRESS10\UDUMP	User process dump directory
utl_file_dir		utl_file accessible directories list
workarea_size_policy	AUTO	SQL working areas sized automatically