



# **Oracle® Business Intelligence Data Warehouse Administration Console Guide**

Version 7.9

February 2007

Part Number: B31978-01

Copyright © 2007, Oracle. All rights reserved.

The Programs (which include both the software and documentation) contain proprietary information; they are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright, patent, and other intellectual and industrial property laws. Reverse engineering, disassembly, or decompilation of the Programs, except to the extent required to obtain interoperability with other independently created software or as specified by law, is prohibited.

The information contained in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. This document is not warranted to be error-free. Except as may be expressly permitted in your license agreement for these Programs, no part of these Programs may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose.

If the Programs are delivered to the United States Government or anyone licensing or using the Programs on behalf of the United States 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, use, duplication, disclosure, modification, and adaptation of the Programs, including documentation and technical data, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement, and, to the extent applicable, the additional rights set forth in FAR 52.227-19, Commercial Computer Software--Restricted Rights (June 1987). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

The Programs are not intended for use in any nuclear, aviation, mass transit, medical, or other inherently dangerous applications. It shall be the licensee's responsibility to take all appropriate fail-safe, backup, redundancy and other measures to ensure the safe use of such applications if the Programs are used for such purposes, and we disclaim liability for any damages caused by such use of the Programs.

Oracle, JD Edwards, PeopleSoft, and Siebel are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

The Programs may provide links to Web sites and access to content, products, and services from third parties. Oracle is not responsible for the availability of, or any content provided on, third-party Web sites. You bear all risks associated with the use of such content. If you choose to purchase any products or services from a third party, the relationship is directly between you and the third party. Oracle is not responsible for: (a) the quality of third-party products or services; or (b) fulfilling any of the terms of the agreement with the third party, including delivery of products or services and warranty obligations related to purchased products or services. Oracle is not responsible for any loss or damage of any sort that you may incur from dealing with any third party.

# Contents

## **Chapter 1: What's New in This Release**

## **Chapter 2: Overview of Oracle Business Analytics Warehouse**

- Oracle Business Analytics Warehouse Overview 9
- Oracle Business Analytics Warehouse Architecture 10
  - Oracle Business Analytics Warehouse Architecture Components 11
- About the Data Warehouse Administration Console 12
- About Source System Containers 15
  - About DAC Repository Objects 16
  - About the DAC Process Life Cycle 17

## **Chapter 3: DAC Quick Start**

## **Chapter 4: Overview of the DAC Interface**

- Navigating the DAC Interface 22
- The DAC's Menu Bar 23
- The DAC's Views 26
- The DAC's Top Pane Toolbar 27
- The DAC's Right-Click Menus 28
- The DAC's Server Monitor Icons 33
- The DAC's Navigation Tree 34
- The DAC's Editable Lists 34
- Using the DAC Query Functionality 35
  - DAC Query Commands and Operators 35
  - DAC Query Examples 36
  - Common DAC Query Procedures 36
  - Using Flat Views Querying 37
- About Object Ownership in the DAC 38

## **Chapter 5: Customizing, Designing, Executing and Monitoring ETL Processes**

Creating or Copying a Source System Container	41
About Customizing the Data Warehouse	42
Adding a New Table and Columns to the Data Warehouse	44
Adding an Index to the Data Warehouse	46
Importing New Data Warehouse Objects into the Informatica Repository	46
Creating Informatica Mappings and Workflows	46
Creating Tasks in the DAC for New or Modified Informatica Workflows	47
Setting a Task Phase Dependency	48
Creating a Task Group	49
Setting Source System Parameters	49
Setting Task Level Parameters	50
Working with Configuration Tags	51
Considerations in Designing a Subject Area	54
Creating a Subject Area	56
Building and Running an Execution Plan with the DAC	57
Creating a Micro ETL Execution Plan	58
Scheduling an Execution Plan	60
About Refresh Dates	60
Monitoring Execution Plan Processes	61

## **Chapter 6: Common Tasks Performed in the DAC**

Importing DAC Metadata	63
Exporting DAC Metadata	64
Distributing DAC Metadata	65
Running the DAC Server Automatically	65
Command Line Access to the DAC Server	66
Setting Up Command Line Access to the DAC Server	67
Using the Command Line to Access the DAC Server	68
DAC Repository Command Line Options	70
Replacing an Informatica Workflow with a Custom SQL File	72

Determining the Informatica Server Maximum Sessions Parameter Setting	73
Determining the Number of Transactional and Data Warehouse Database Connections	75
Running Two DAC Servers on the Same Machine	75
Customizing Index and Analyze Table Syntaxes	76
Using SQL Files as an Execution Type in the DAC	77
Overview of Change Capture Process (Siebel Sources Only)	78
Initial Data Capture	79
Change Capture Mechanisms	79
Using the Change Capture Filter	84
Tracking Deleted Records	84
Handling ETL Failures with the DAC	86

## Chapter 7: DAC Functional Reference

Common Elements of DAC Interface Tabs	89
About the DAC Design View	89
About the DAC Subject Areas Tab	90
About the DAC Tables Tab	92
About the DAC Indices Tab	94
About Advanced Custom Index Management	96
About the DAC Tasks Tab	97
About the DAC Task Groups Tab	102
About the DAC Configuration Tags Tab	105
About the DAC Source System Parameters Tab	106
About the DAC Source System Folders Tab	106
About the DAC Setup View	106
About the DAC System Properties Tab	107
About the Informatica Servers Tab	111
About the Physical Data Sources Tab	112
About the Email Recipients Tab	114
About the DAC Execute View	114
About the DAC Execution Plans Tab	115
About the DAC Current Run Tab	119
About the DAC Run History Tab	122
About the DAC Scheduler Tab	122

## Index



# 1

## What's New in This Release

Oracle Business Intelligence Applications consists of components that were formerly available from Siebel Systems as Siebel Business Analytics Applications (both CRM and Enterprise) with a number of significant enhancements.

The Oracle Business Intelligence Data Warehouse Administration Console Guide contains information about using the Data Warehouse Administration Console (DAC), a centralized console for schema management as well as configuration, administration, loading, and monitoring of the Oracle Business Analytics Warehouse.

Oracle recommends reading the Oracle Business Intelligence Applications Release Notes before installing, using, or upgrading Oracle Business Intelligence Applications. The Oracle Business Intelligence Applications Release Notes are available:

- On the Oracle Business Intelligence Applications CD-ROM.
- On the Oracle Technology Network at [http://www.oracle.com/technology/documentation/bi\\_apps.html](http://www.oracle.com/technology/documentation/bi_apps.html) (to register for a free account on the Oracle Technology Network, go to <http://www.oracle.com/technology/about/index.html>).

### **What's New in Oracle Business Intelligence Data Warehouse Administration Console Guide, Version 7.9**

The Oracle Business Intelligence Data Warehouse Administration Console Guide is a new guide in Release 7.9. Some of the information in this guide was previously published in the *Siebel Business Analytics Applications Installation and Administration Guide*.





# 2

## Overview of Oracle Business Analytics Warehouse

This chapter provides an overview of the Oracle Business Analytics Warehouse and the Data Warehouse Administration Console (DAC). It includes the following topics:

- [Oracle Business Analytics Warehouse Overview on page 9](#)
- [Oracle Business Analytics Warehouse Architecture on page 10](#)
- [About the Data Warehouse Administration Console on page 12](#)
- [About Source System Containers on page 15](#)

### Oracle Business Analytics Warehouse Overview

The Oracle Business Analytics Warehouse is a unified data repository for all customer-centric data. The purpose of the Oracle Business Analytics Warehouse is to support the analytical requirements of Oracle Business Intelligence Applications.

The Oracle Business Analytics Warehouse includes the following:

- A data integration engine that combines data from multiple source systems to build a data warehouse.
- An open architecture to allow organizations to use third-party analytical tools in conjunction with the Oracle Business Analytics Warehouse using the Oracle Business Intelligence Server.
- Prebuilt data extractors to incorporate data from external applications into the Oracle Business Analytics Warehouse.
- A set of ETL (extract-transform-load) processes that takes data from multiple source systems and creates the Oracle Business Analytics Warehouse tables.
- The DAC, a centralized console for schema management as well as configuration, administration, loading, and monitoring of the Oracle Business Analytics Warehouse.

Figure 1 provides an overview of the Oracle Business Analytics Warehouse.

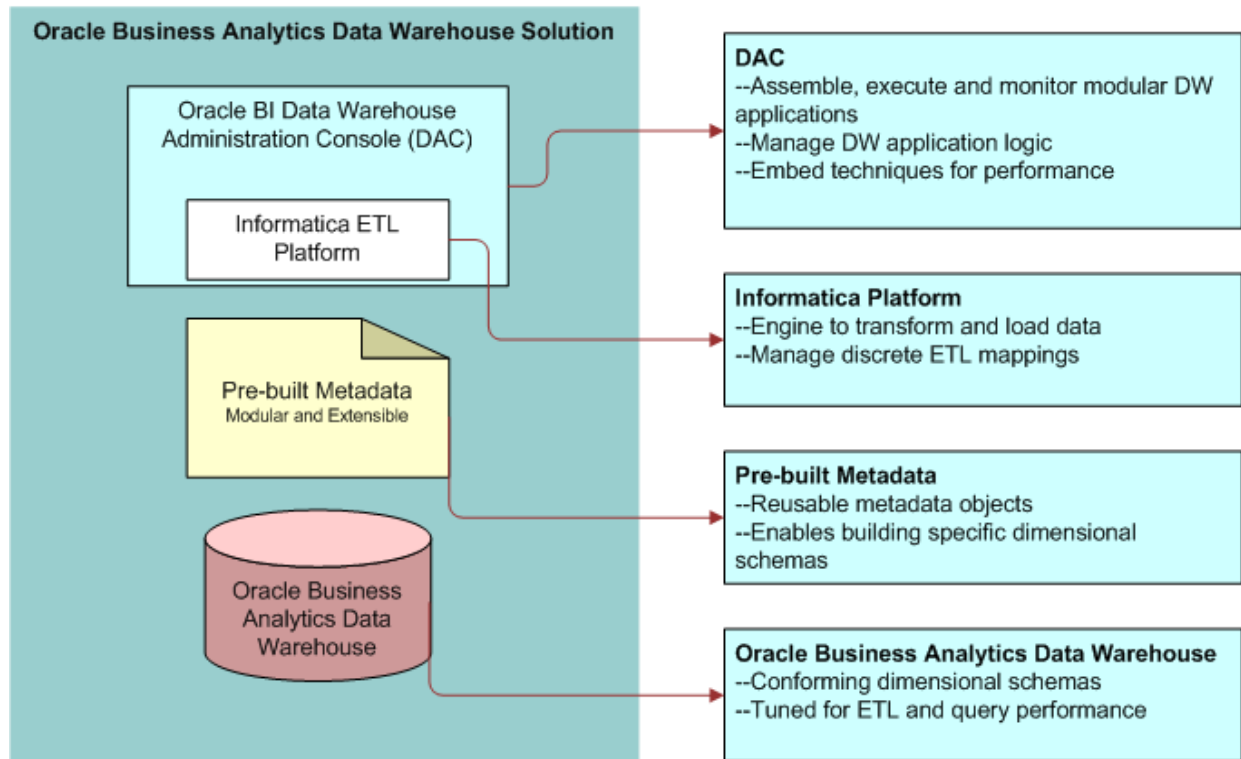


Figure 1. Data Warehouse Overview

## Oracle Business Analytics Warehouse Architecture

High-level analytical queries, like those commonly used in Oracle Business Analytics Warehouse, scan and analyze large volumes of data using complex formulas. This process can take a long time when querying a transactional database, which impacts overall system performance.

For this reason, the Oracle Business Analytics Warehouse was constructed using dimensional modeling techniques to allow for fast access to information required for decision making. The Oracle Business Analytics Warehouse derives its data from operational applications, and uses Informatica's data integration technology to extract, transform, and load data from transactional databases into the Oracle Business Analytics Warehouse.

Figure 2 illustrates how the Oracle Business Analytics Warehouse interacts with the other components of Oracle BI Applications.

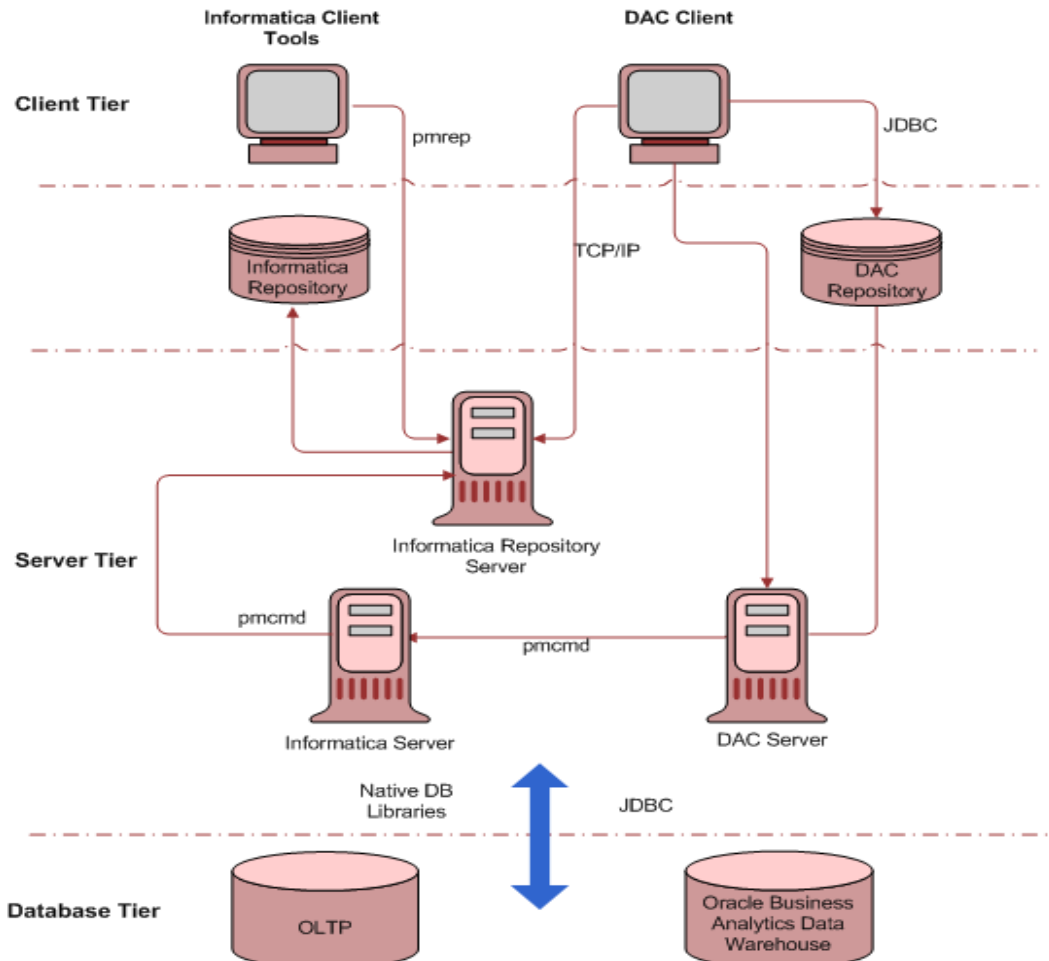


Figure 2. Oracle Business Intelligence Applications Architecture

## Oracle Business Analytics Warehouse Architecture Components

The Oracle Business Analytics Warehouse architecture comprises the following components:

- **DAC client.** A command and control interface for the data warehouse to allow for schema management, and configuration, administration, and monitoring of data warehouse processes. It also allows you to design subject areas and build execution plans.

- **DAC server.** Executes the instructions from the DAC client. The DAC server manages data warehouse processes, including loading of the ETL and scheduling execution plans. It dynamically adjusts its actions based on information in the DAC repository. Depending on your business needs, you might incrementally refresh the Oracle Business Analytics Warehouse once a day, once a week, once a month, or on another similar schedule.
- **DAC repository.** Stores the metadata (semantics of the Oracle Business Analytics Warehouse) that represents the data warehouse processes.
- **Informatica Server.** Loads and refreshes the Oracle Business Analytics Warehouse.
- **Informatica Repository Server.** Manages the Informatica repository.
- **Informatica Repository.** Stores the metadata related to Informatica workflows.
- **Informatica client utilities.** Tools that allow you to create and manage the Informatica repository.

## About the Data Warehouse Administration Console

The DAC provides a framework for the entire life cycle of data warehouse implementations. It allows you to create, configure, execute, and monitor modular data warehouse applications in a parallel, high-performing environment. For information about the DAC process life cycle, see [“About the DAC Process Life Cycle” on page 17](#).

The DAC complements the Informatica ETL platform. It provides *application-specific* capabilities that are not prebuilt into ETL platforms. For example, ETL platforms are not aware of the semantics of the subject areas being populated in the data warehouse nor the method in which they are populated. The DAC provides these application capabilities at a layer of abstraction above the ETL execution platform, as illustrated in Figure 3.

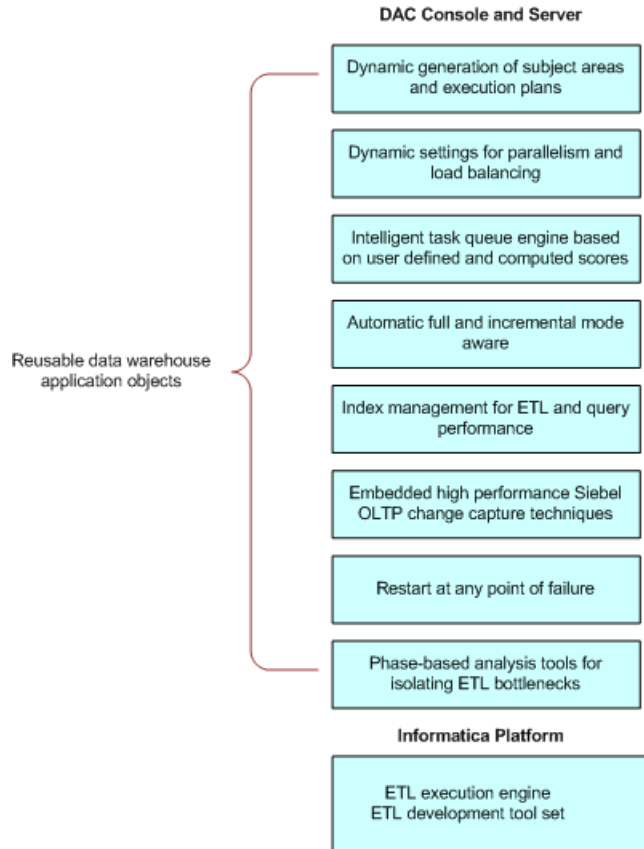


Figure 3. DAC Application-Specific Capabilities

## Important DAC Features

Important DAC features allow you to do the following:

### Minimize install, setup, and configuration time

- Create physical data model in the data warehouse
- Set language, currency, and other settings
- Design subject areas and build execution plans

### Manage metadata driven dependencies and relationships

- Generate custom ETL execution plans
- Automate change capture for the Siebel transactional database
- Capture deleted records
- Assist in index management
- Perform dry runs and test runs of execution plans

### Provide reporting and monitoring to isolate bottlenecks

- Perform error monitoring and email alerting
- Perform structured ETL analysis and reporting

### Utilize performance execution techniques

- Automate full and incremental mode optimization rules
- Set the level of Informatica session concurrency
- Load balance across multiple Informatica servers
- Restart from point of failure

- Queue execution tasks for performance (See [Figure 4.](#))

The DAC manages the task execution queue based on metadata driven priorities and scores computed at runtime. This combination allows for flexible and optimized execution.

**Metadata Driven Priority Assignment**




Figure 4. Task Execution Queue

## About Source System Containers

Source system containers hold repository objects that correspond to a specific source system. For information about the different kinds of repository objects, see [“About DAC Repository Objects” on page 16.](#)

You can use the preconfigured source system containers to create your own source system container. You cannot modify objects in the preconfigured source system containers. You must make a copy of a preconfigured container in order to make any changes to it.

For instructions on creating a new source system container or copying an existing container, see [“Creating or Copying a Source System Container” on page 41.](#)

## About DAC Repository Objects

All DAC repository objects are associated with a source system container. For more information about source system containers, see [“About Source System Containers” on page 15](#) and [“About Object Ownership in the DAC” on page 38](#).

The DAC repository stores application objects in a hierarchical framework that defines a data warehouse application. The DAC allows you to view the repository application objects based on the source system container you specify. The source system container holds the metadata that corresponds to the source system with which you are working.

A data warehouse application comprises the following repository objects:

- **Subject area.** A logical grouping of tables related to a particular subject or application context, as well as the tasks that are associated with the tables. Subject areas are assigned to execution plans, which can be scheduled for full or incremental loads. A subject area also includes the tasks required to load the subject area tables.
- **Tables.** Physical database tables defined in the database schema. Can be transactional database tables or data warehouse tables. Table types can be fact, dimension, hierarchy, aggregate, and so on, as well as flat files that can be sources or targets.
- **Task.** A unit of work for loading one or more tables. A task comprises the following: source and target tables, phase, execution type, truncate properties, and commands for full or incremental loads. When you assemble a subject area, the DAC automatically assigns tasks to it. Tasks that are automatically assigned to the subject area by the DAC are indicated by the Autogenerated flag in the Tasks child tab of the Subject Areas tab.
- **Task Groups.** A group of tasks that you define because you want to impose a specific order of execution. A task group is considered to be a “special task.”
- **Execution plan.** A data transformation plan defined on subject areas that needs to be transformed at certain frequencies of time. An execution plan is defined based on business requirements for when the data warehouse needs to be loaded. An execution plan comprises the following: ordered tasks, indices, tags, parameters, source system folders, and phases.
- **Schedule.** A schedule specifies when and how often an execution plan runs. An execution plan can be scheduled for different frequencies or recurrences by defining multiple schedules.



## About the DAC Process Life Cycle

The DAC is used by different user groups to design, execute, monitor, and diagnose execution plans. These phases together make up the DAC process life cycle, as shown in [Figure 5](#).

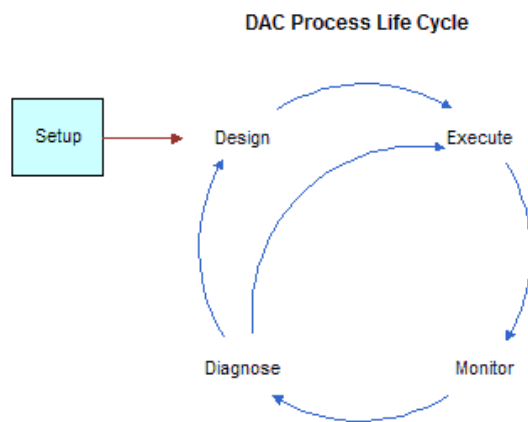


Figure 5. DAC Process Life Cycle

The phases of the process and the actions associated with them are as follows:

- Setup
  - Set up database connections
  - Set up ETL processes (Informatica)
  - Set up email recipients
- Design
  - Define application objects
  - Design execution plans
- Execute
  - Define scheduling parameters to run execution plans
  - Access runtime controls to restart or stop currently running schedules
- Monitor
  - Monitor runtime execution of data warehouse applications
  - Monitor users, DAC repository, and application maintenance jobs



# 3 DAC Quick Start

In order to start the DAC client, you must have completed the steps show in [Figure 6](#).

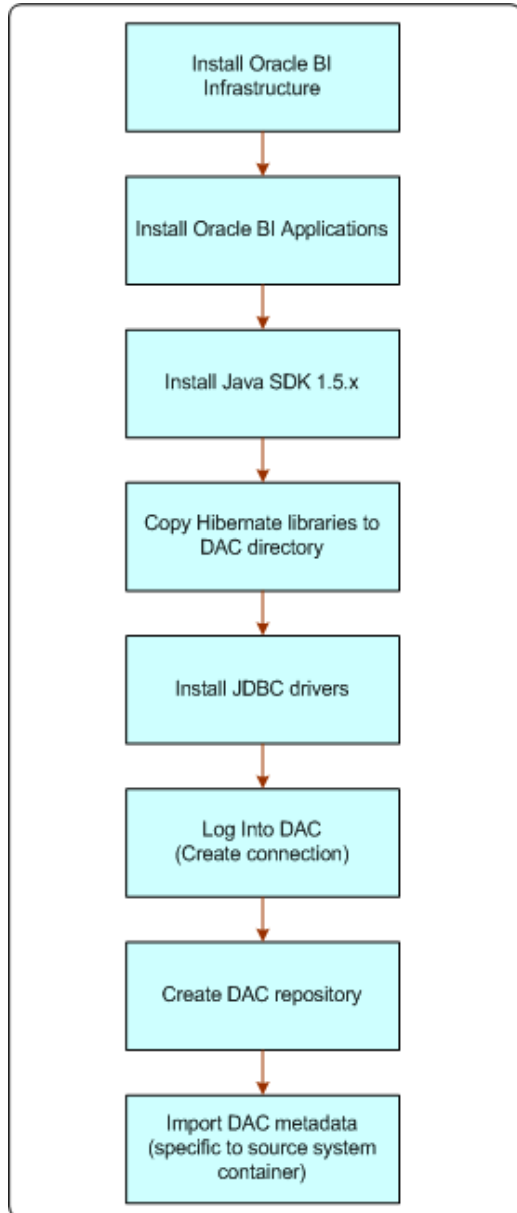


Figure 6. DAC Quick Start Flow

For instructions on installing the Oracle BI Infrastructure, see the *Oracle Business Intelligence Infrastructure Installation and Configuration Guide*. For instructions on completing the remaining steps, see the *Oracle Business Intelligence Applications Installation and Configuration Guide*.

# 4

## Overview of the DAC Interface

This chapter contains the following topics:

- [Navigating the DAC Interface on page 22](#)
- [The DAC's Menu Bar on page 23](#)
- [The DAC's Views on page 26](#)
- [The DAC's Top Pane Toolbar on page 27](#)
- [The DAC's Right-Click Menus on page 28](#)
- [The DAC's Server Monitor Icons on page 33](#)
- [The DAC's Navigation Tree on page 34](#)
- [The DAC's Editable Lists on page 34](#)
- [Using the DAC Query Functionality on page 35](#)
- [About Object Ownership in the DAC on page 38](#)

# Navigating the DAC Interface

Figure 7 shows the main elements of the DAC window.

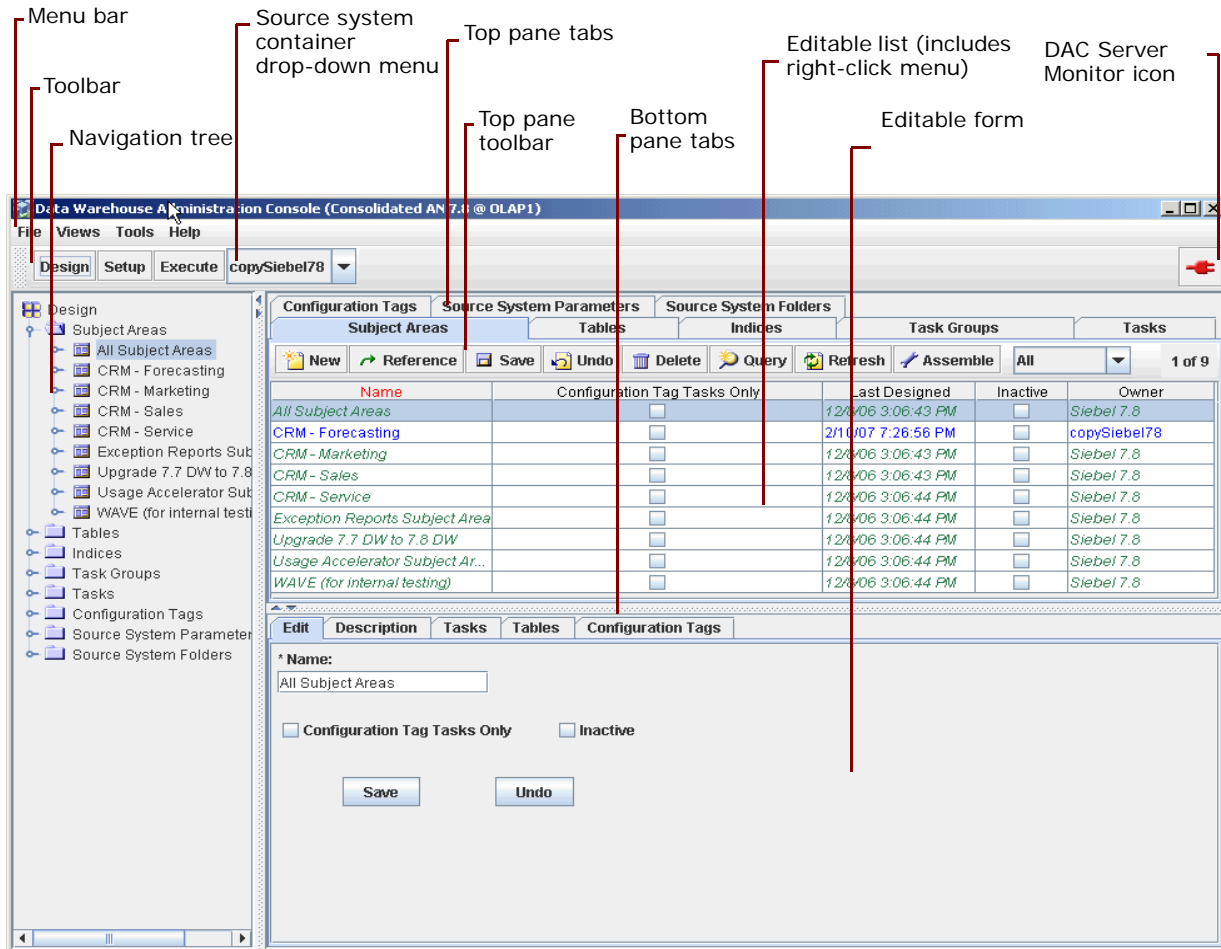


Figure 7. Example of DAC Window

## The DAC's Menu Bar

Table 1 provides a description of the DAC's menu bar options:

Table 1. DAC Menu Bar Options

File	The File menu contains options to close the DAC client and to create, copy, or delete source system containers. For instructions on creating or copying a source system container, see <a href="#">"Creating or Copying a Source System Container"</a> on page 41.
Views	The Views menu allows you to navigate to the various tabs in the top pane window.
Tools	The Tools menu provides access to functionality related to the DAC and Informatica repositories. <a href="#">Table 2</a> provides a description of the Tools menu commands.
Help	The Help menu provides details about the current DAC login as well as the version of DAC metadata and software. There is no online help for the DAC.

## Tools Menu Options

Table 2 provides a description of the Tools menu commands.

Table 2. DAC Tools Menu Commands

Tools Menu Command	Description
DAC Repository Management > Export	<p>Allows you to export the DAC metadata, in XML format, based on the source system container, in order to back up the metadata or to reproduce the environment elsewhere. In the Export dialog box, you can specify a directory in which to store the XML file or accept the default directory, which is DAC\export.</p> <p>In the Export dialog box, you can select the following category options:</p> <ul style="list-style-type: none"> <li>■ <b>Logical.</b> Exports all information contained in the Design view and database connection information.</li> <li>■ <b>Run Time.</b> Exports all information contained in the Execute view.</li> <li>■ <b>System.</b> Exports all information contained in the Setup view, except passwords for servers and database connections.</li> </ul>
DAC Repository Management > Import	<p>Allows you to import the DAC metadata for the source system containers you specify.</p> <p>In the Import dialog box, you can specify the following:</p> <ul style="list-style-type: none"> <li>■ <b>Import/Export folder.</b> A directory from which to import the data. The default directory is DAC\export.</li> <li>■ <b>Truncate repository tables.</b> Indicates whether you want to truncate the repository tables. If you select this option, the existing metadata is overwritten.</li> <li>■ <b>Enable batch mode.</b> Indicates whether batch mode is enabled, which inserts the imported metadata into the repository as an array insert.</li> </ul> <p>In the Import dialog box, you can select the following category options:</p> <ul style="list-style-type: none"> <li>■ <b>Logical.</b> Imports all information contained in the Design view and database connection information.</li> <li>■ <b>Run Time.</b> Imports all information contained in the Execute view.</li> <li>■ <b>System.</b> Imports all information contained in the Setup view, except passwords for servers and database connections.</li> </ul>



Table 2. DAC Tools Menu Commands

Tools Menu Command	Description
DAC Repository Management > Create Repository Report	<p>Allows you to generate a DAC repository report based on the following criteria:</p> <ul style="list-style-type: none"> <li>■ Table Row Counts</li> <li>■ Object References by Entity</li> <li>■ Ownerless Objects</li> <li>■ Unreferenced Objects</li> <li>■ Dead References</li> </ul> <p>The Clean Up command removes unused referenced objects.</p>
DAC Repository Management > Purge Run Details	<p>Allows you to purge completed runs from the run history. You can purge all runs (except the last run) or specify particular runs to be purged. The last run cannot be purged.</p> <p>In the Purging Runs... dialog box, the following options are available:</p> <ul style="list-style-type: none"> <li>■ <b>All.</b> Purges all completed runs except for the last run.</li> <li>■ <b>By Execution Plan.</b> Allows you to select an execution plan whose associated runs you want to purge.</li> <li>■ <b>By Run Name.</b> Allows you to select an individual run for purging.</li> <li>■ <b>Before Specified Date.</b> Allows you to select a date before which all runs except the last run will be purged.</li> <li>■ <b>Details Only.</b> Purges all related information about a run but leaves the run header information.</li> </ul>
DAC Repository Management > Analyze Repository Tables	Allows you to run analyze table commands for all the DAC repository tables.
DAC Repository Management > Default Index Properties	Allows you to specify which databases will be associated with newly created indices.
DAC Repository Management > Drop DAC Repository	Allows you to drop all the DAC repository tables. This action deletes all data in the repository.
DAC Server Management > Get Server Log	When the DAC server is running an ETL process, this command opens a text box that displays streaming data related to the process.
DAC Server Management > DAC Server Setup	Allows you to configure the DAC server connections and server email settings. This action should be performed on the machine where the DAC server is running.

Table 2. DAC Tools Menu Commands

Tools Menu Command	Description
ETL Management > Configure	Opens the Data Warehouse Configuration wizard, which allows you to create and drop data warehouse tables and to create delete triggers.
ETL Management > Reset Data Warehouse	Clears the refresh dates for all source and target tables. This action forces a full load to occur.
Seed Data > Task Phases	Allows you to add, edit, or delete task phases.
Seed Data > Task Folders	Allows you to add, edit, or delete task folders.
Seed Data > Logical Data Sources	Allows you to add, edit, or delete logical data sources.
UI Styles > Windows (MFC)	Changes the user interface to the Windows style.
UI Styles > UNIX (MOTIF)	Changes the user interface to the UNIX style.
UI Styles > Java (METAL)	Changes the user interface to the Java style.

## The DAC's Views

The DAC's View buttons are located directly under the menu bar. [Table 3](#) provides a description of the different DAC Views.

Table 3. DAC Views

Design	<p>The Design view provides access to functionality related to creating and managing subject areas. For more information, see <a href="#">"About the DAC Design View" on page 89</a>.</p> <p><b>Source System Container Drop-Down List</b></p> <p>When the Design view is active, the Source System Container drop-down list appears to the right of the View buttons. It allows you to select the source system container that holds the metadata corresponding to a source system.</p>
Setup	<p>The Setup View provides access to functionality related to setting up DAC system properties, Informatica servers, database connections, and email notification. For more information, see <a href="#">"About the DAC Setup View" on page 106</a>.</p>
Execute	<p>The Execute view provides access to functionality related to setting up, running, monitoring, and scheduling execution plans. For more information, see <a href="#">"About the DAC Execute View" on page 114</a>.</p>

## The DAC's Top Pane Toolbar

Table 4 describes the commands available in the top pane toolbar.

Table 4. DAC Top Pane Toolbar

Command	Description
New	Creates a placeholder for a new record in the selected list.
Save	Saves the current record.
Undo	Undoes changes made to the current record after the last save.
Delete	<p>Deletes the selected record. If you delete a parent record, the child records are also deleted.</p> <p><b>NOTE:</b> When you delete a column from a table, the column is not automatically deleted from the index.</p> <p><b>NOTE:</b> The DAC does not display deleted objects. You must look at the database to figure out what objects were deleted.</p>
Query	Opens a blank query.
Refresh	Retrieves the data from the repository with the last used query.
Reference	Design view only. Opens the Reference dialog box, which allows you to copy objects from one container to another. For more information about referencing objects, see <a href="#">"About Object Ownership in the DAC" on page 38</a> .
Assemble	Design view only. Assembles a subject area, with dimension and related tables as well as tasks.
Drop-down list	Design view only. Allows you to filter the source system container objects that appear in the top pane list.
Run Now	Execute view, Execution Plans tab only. Starts a new ETL process.
Start	Execute view, Current Run and Run History tabs only. Restarts the selected ETL, after the ETL has failed, stopped, or been aborted.
Stop	Execute view, Current Run and Run History tabs only. Stops an ETL in progress. All currently running tasks will complete, and queued tasks will stop. The status of the ETL changes to Stopped.
Abort	Execute view, Current Run and Run History tabs only. Causes an ETL in progress to abort. All currently running tasks will be aborted. The status of queued tasks and the ETL itself will change to Stopped.
Auto Refresh	Execute view, Current Run tab only. Allows you to turn on and off the automatic screen refresh functionality and set the refresh interval.

## The DAC's Right-Click Menus

The commands available in the right-click menus depend on the tab that is active. For descriptions of the commands, see the following topics:

- [Common Right-Click Menu Commands on page 29](#)
- [Design View Right-Click Menu Commands on page 30](#)
- [Setup View Right-Click Menu Commands on page 31](#)
- [Execute View Right-Click Menu Commands on page 32](#)

## Common Right-Click Menu Commands

Table 5. Common Right-Click Menu Commands

Command	Description
Copy String	Copies the contents of a cell (editable and read-only) to the clipboard
Paste String	Pastes a string from the clipboard into a selected cell that supports a string data type.
Copy Record	<p>Creates a copy of the selected record, with a unique record ID. The new record is committed to the database when you click the Save button or click outside the cell.</p> <p>In the Design view tabs (except for the Indices tab), Copy Record copies the selected record and the record's child records. When you copy a subject area, the tables are also copied but the tasks are not copied. You need to use the Assemble command to reassemble the subject area and add tasks to it.</p> <p>In the Design view Indices tab and Setup and Execute views, Copy Record copies only the selected record.</p>
Delete	<p>Deletes the selected record. If you delete a parent record, the child records are also deleted.</p> <p><b>NOTE:</b> When you delete a column from a table, the column is not automatically deleted from the index. You must manually delete columns from indices that were deleted from a table or else the ETL process will fail.</p> <p><b>NOTE:</b> The DAC does not display deleted objects. You must look at the database to figure out what objects were deleted.</p>
Output to File	Outputs to a text file in the DAC root folder the contents of the current tab's record list.
Record Info	Displays the record's unique ID, object type, current source system, owner source system, and the timestamp for when it was last updated. It also displays the source system lineage and the source systems that reference the object.
Update Records	For some columns, allows you to update the column value for each row to a single value.

## Design View Right-Click Menu Commands

Table 6. Design View Right-Click Menu Commands

Command	Description
Ownership	<ul style="list-style-type: none"> <li>■ <b>Reference.</b> Opens the Reference dialog box, which allows you to reference objects from one container to another. The reference function works like a symbolic link or shortcut.</li> <li>■ <b>Re-Reference.</b> If an object is a referenced object, that is, a reference to an object in another container and a change is made to the original object's child objects, you can use this command to import the changes to the referenced object.</li> <li>■ <b>Push to References.</b> If an original object is changed, you can use this command to export the changes to all referenced objects' child objects.</li> <li>■ <b>De-Clone.</b> When you make changes to a referenced object, the new object is called a <i>clone</i>. This command allows you to revert a cloned object back to its state as a reference.</li> <li>■ <b>Re-Assign Record.</b> This command allows you to reassign an objects ownership.</li> </ul> <p>For more information about the ownership of objects, see <a href="#">"About Object Ownership in the DAC" on page 38</a>.</p>
Assemble	Assembles a subject area, with dimension and related tables as well as tasks.
Generate Index Scripts	Generates drop index, create index, and analyze table scripts for all tables that participate in the ETL process. The results are stored in the log\scripts directory.
Change Capture Scripts	<ul style="list-style-type: none"> <li>■ <b>Image and Trigger Scripts.</b> Generates change capture scripts for tables with defined image suffixes. The scripts may include delete triggers, create and drop statements for delete triggers, and image tables and their indices.</li> <li>■ <b>View Scripts.</b> Generates change capture view scripts for full or incremental mode for tables that participate in the change capture process. This command can be used for unit testing.</li> <li>■ <b>Change Capture SQL.</b> Generates change capture SQL scripts for full or incremental mode for tables that participate in the change capture process. This command can be used for unit testing.</li> </ul>

Table 6. Design View Right-Click Menu Commands

Command	Description
Import from Database	<ul style="list-style-type: none"> <li>■ <b>Import Database Tables.</b> Allows you to import table definitions from a selected database. This action does not import columns.</li> <li>■ <b>Import Indices.</b> Allows you to import index definitions from a selected database for one or more tables as listed in the result of the query.</li> <li>■ <b>Import Database Columns.</b> Allows you to import column definitions from a selected database.</li> </ul>
Filter Indices	Allows you to filter by database type the indices that are displayed in the top pane list.
Output Task Description	Saves to an HTML file the description for a selected task or for all tasks.
Synchronize Tasks	Synchronizes the information the DAC has for a task's source and target tables with the information in the Informatica repository.
Flat Views	<p>Opens a dialog box that allows you to query for various objects, modify data, and do mass updates.</p> <p>You can query for the following objects:</p> <p><b>Tables tab:</b></p> <ul style="list-style-type: none"> <li>■ Related Tables</li> <li>■ Table Columns</li> </ul> <p><b>Indices tab:</b> Index columns</p> <p><b>Tasks tab:</b></p> <ul style="list-style-type: none"> <li>■ Task Source Tables</li> <li>■ Task Target Tables</li> <li>■ Task Conditional Tables</li> <li>■ Task Phase Dependencies</li> <li>■ Task Parameters</li> </ul>

## Setup View Right-Click Menu Commands

Table 7. Setup View Right-Click Menu Commands

Command	Description
Test Connection	<p>In the Physical Data Sources tab, it allows you to test the database connection.</p> <p>In the Informatica Servers tab, it allows you to test the connection to the Informatica Server and Repository Server.</p> <p>The DAC server performs this command if the DAC client is connected to a server. If the DAC client is not connected to a DAC server, then the DAC client performs the command.</p>

## Execute View Right-Click Menu Commands

Table 8. Execute View Right-Click Menu Commands




Command	Description
Mark as Completed	Changes the status of a stopped or failed ETL to Completed. In the audit trail for this ETL, the status is Marked as Completed. Use this command with caution; it can cause the data warehouse to be inconsistent.
Get Run Information > Get Log File	Fetches the log file for this run from the DAC server and saves it in the ServerLog folder.
Get Run Information > Analyze Run	Saves a description of the run as an HTML file in the Log/Statistics folder.
Get Run Information > Get Chart	Displays a chart showing changes in task statuses over time in a separate window.
Get Run Information > Get Graph	Displays a graph showing changes in task statuses over time in a separate window.



## The DAC's Server Monitor Icons

The Server Monitor is located in the upper-right corner of the DAC client. Its color and shape change based on the DAC server status. When the DAC client cannot establish a connection to the DAC server, the Server Monitor icon resembles a red electrical plug. When the client is connected to the server and the server is idle, the icon resembles an orange electrical plug in a socket. Finally, if the client is connected to a server that is running an ETL process, the icon resembles a green electrical plug with a lightning sign superimposed on it. In addition, clicking on the icon when there is a connection to the server opens a text box that displays data related to the ETL process.

Table 9. DAC Server Monitor Icons

Icon	Description
	DAC client cannot establish a connection to the DAC server.
	DAC client is connected to the server, and the server is idle.
	DAC client is connected to a server that is running an ETL process.

## The DAC's Navigation Tree

The navigation tree appears on the left side of the DAC window, as shown in [Figure 8](#). The tree root nodes correspond to the tabs in the top pane of the DAC window. When a plus sign (+) appears before a node, you can expand the node to view the records belonging to the node. You can double-click a record in the tree to have it display in the top pane in a single-record mode (New, Delete, Copy Record, and Query commands are unavailable), and double-click the root node to return to the list mode.

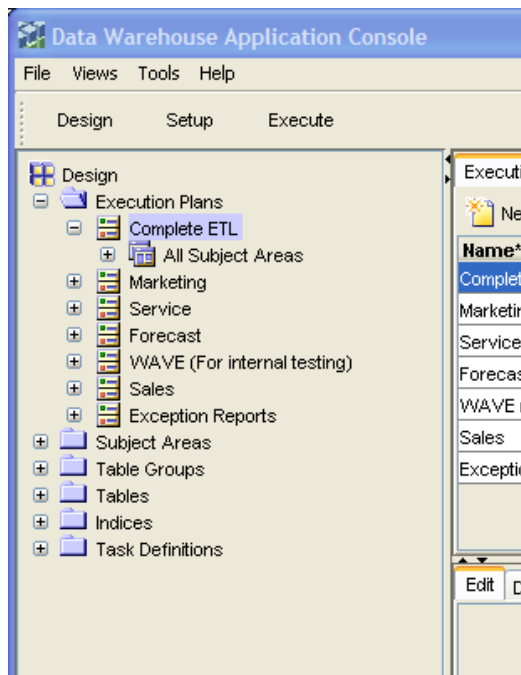


Figure 8. DAC Navigation Tree

## The DAC's Editable Lists

The top and bottom panes of the DAC window display records in a list format. Some of the columns in the list are editable, and others are read-only. The toolbar at the top of each pane allows you to perform various tasks associated with a selected record in the list. For a description of the toolbar commands, see [“The DAC's Top Pane Toolbar” on page 27](#).

A right-click menu is also accessible from the lists in both the top and bottom panes. For a description of these commands, see [“The DAC's Right-Click Menus” on page 28](#).

The list format allows you to do the following:

- Edit the data in place and save the record by either clicking another record in the list or clicking the Save button.
- Reorder the columns.
- Sort the data in the list by clicking on the column name.

- Select predefined values from picklists.
- For fields that refer to values from other entities, use the query functionality in pop-up dialog boxes.
- Use Ctrl+C to copy an editable string to the clipboard (not available for read-only strings).
- Ctrl+V to paste a string from the clipboard into a selected cell that supports a string data type.

## Using the DAC Query Functionality

Querying is a way to locate one or more records that meet your specified criteria. Query functionality is available in every DAC screen. When you enter query mode, the Edit and Description child tabs in the bottom pane are not available.

This section includes the following topics:

- [DAC Query Commands and Operators on page 35](#)
- [DAC Query Examples on page 36](#)
- [Common DAC Query Procedures on page 36](#)
- [Using Flat Views Querying on page 37](#)

## DAC Query Commands and Operators

Table 10 describes the query commands and operators you can use to define your query criteria.

Table 10. DAC Query Commands and Operators

Operator	Description
=	Placed before a value, returns records containing a value equal to the query value.
<	Placed before a value, returns records containing a value less than the query value.
>	Placed before a value, returns records containing a value greater than the query value.
<>	Placed before a value, returns records containing a value that is not equal to the query value.
<=	Placed before a value, returns records containing a value less than or equal to the query value.
>=	Placed before a value, returns records containing a value greater than or equal to the query value.
*	Wildcard that can be placed in the middle, or at the beginning or end of a text string.
!	Used for negation.

Table 10. DAC Query Commands and Operators

Operator	Description
""	Surrounds a string that, unless modified by a wildcard, must be matched exactly.
\	Escape symbol is used when double quotes should not be processed as a special symbol. For example, <code>!("*null text"or("\*\\"))</code> is a value expression for a text field. The query returns values that do not end with a string <code>null text</code> and that are not surrounded by double quotes.
()	Surrounds the values and operators that will be processed first.
NULL	Returns records for which the query field is blank.
AND	Placed between values, returns only records for which all the given conditions are true. (Not case sensitive.)
OR	Placed between values, returns records for which at least one condition is true. (Not case sensitive.)

## DAC Query Examples

The following examples show different ways you can query on the Name column of the Tasks tab.

- `Extract*` lists all tasks whose name starts with `Extract`.
- `*Extract*` lists all tasks whose name contains the word `Extract`.
- `!Extract*` lists all tasks whose name does not start with the word `Extract`.
- `!null` lists all tasks whose name is not null.
- `Extract* or Aggregate*` lists all tasks whose name starts with `Extract` or `Aggregate`.
- `Load* and *Aggregate*` lists all tasks whose name starts with `Load` and also contains the word `Aggregate`.
- `"Extract for Wave Dimension" or "Load into Wave Dimension"` lists tasks whose name is either `Extract for Wave Dimension` or `Load into Wave Dimension`.

**NOTE:** When using spaces within strings, you need to surround the string with quotes (`"`).

## Common DAC Query Procedures

This section includes instructions for common query procedures.

### To create and execute a query in the DAC

- 1 In the top or bottom pane of the DAC, click Query on the toolbar or in right-click menu. A blank row in a list appears.
- 2 Enter the query criteria in the appropriate fields.

- 3 Click Run Query on the toolbar.

The query is executed and the records appear.

### *To enter a query value in a date field*

- 1 In the date field, click the calendar icon on the right side of the cell.

The Date dialog box appears.

- 2 Enter the date and time for which you want to search, and select the appropriate query condition.

## Using Flat Views Querying

You can use the Flat Views query feature to query for various objects, modify data, and do mass updates. This feature is available in the right-click menu in the Tables, Indices, and Tasks tabs of the Design view. The Flat Views right-click command is context-sensitive and allows you to query only on certain columns.

You can modify individual records in the query results window, or you can use the Update Records right-click command to update multiple records.

### *To update multiple records using the Flat Views query feature*

- 1 In the DAC, right-click in the Tables, Tasks or Indices tab.
- 2 Select Flat Views, and then select a context-sensitive column on which you want to query.
- 3 In the query dialog box, enter search criteria, and click Go.
- 4 In the query results dialog box, right-click and select Update Records.
- 5 In the Update Record Set dialog box, select the column you want to update, and then click Set Value.
- 6 Enter a value for the column.
- 7 To update records that are referenced objects, select Update Referenced Records.

If you select this check box, referenced objects as well as original and cloned objects will be updated. The referenced objects will become clones, and the ownership column for these records will be updated to reflect the new ownership.

If you do not select this check box, only the columns in records that are original or cloned objects (objects owned by the source system container) will be modified.

- 8 Click OK.
- 9 Click Yes when asked if you want to proceed.

An informational message tells you which records were updated.

- 10 Click OK to close the window.

## About Object Ownership in the DAC

The source system container in which an object originates is the *owner* container. The tabs in the DAC Design view display the owner of the various repository objects. You can reuse an object among different source system containers by *referencing* the object. A reference works like a symbolic link or shortcut. You can use the referenced object just as you would an original object, but the object's ownership remains unchanged.

For example, W\_INVOICE\_F is a fact table whose owner is the Siebel 7.8 source system container. You can reuse W\_INVOICE\_F in any other container by referencing it.

You can reference an object from its owner container, and you can also reference an object that has already been referenced by another source system container.

If you modify a referenced object, the modified object becomes a *clone* and the ownership changes to the source system container in which you performed the modification.

When you make changes to an original object that has been referenced by other containers, any updates to the original object are immediately reflected in the referenced object. If you delete the original object, all referenced objects are also deleted.

Changes to an original object's child objects are not automatically reflected in the referenced object's child objects. You need to use the right-click command Ownership > Push to References to *push* the changes to the referenced object's child objects. And, conversely, you can import into a referenced object the changes made to an original object; this function is referred to as a *re-reference*.

For a description of the ownership functionality available in the Design view right-click menu, see ["Design View Right-Click Menu Commands" on page 30](#).

## Font Variations of Objects Displayed in the DAC

The different categories of objects are represented in the DAC with differing fonts.

Table 11. Sample table

Object Type	Font
Original object	(System dependent) Black by default, regular style.
Referenced object	Green color, italic style.
Clone	Blue color, regular style.

## Object Behaviors To Consider

- Changes made to parent objects in the owner container are automatically pushed to the parent referenced objects.
- When you add child objects to a parent object, you have to push the changes to the referenced objects. For example, if you add a column to a table registered in the DAC, the new column is not automatically added to the references.

- When you delete a referenced object, only the referenced object is deleted, not the original object.
- If you delete an object from the owner container, the object is deleted and all references are deleted. This is referred to as a *deep delete*. For example, if you delete a table from the owner container, the table and columns are deleted.
- If you delete a column from the owner table, the column is deleted in all the referenced objects.
- If you delete child objects from the owner object, the referenced child objects are automatically deleted.





# 5

## Customizing, Designing, Executing and Monitoring ETL Processes

This chapter provides information about customizing, designing, executing, and monitoring ETL processes.

This section includes the following topics:

- [Creating or Copying a Source System Container on page 41](#)
- [About Customizing the Data Warehouse on page 42](#)
- [Adding a New Table and Columns to the Data Warehouse on page 44](#)
- [Adding an Index to the Data Warehouse on page 46](#)
- [Importing New Data Warehouse Objects into the Informatica Repository on page 46](#)
- [Creating Informatica Mappings and Workflows on page 46](#)
- [Creating Tasks in the DAC for New or Modified Informatica Workflows on page 47](#)
- [Setting a Task Phase Dependency on page 48](#)
- [Creating a Task Group on page 49](#)
- [Setting Source System Parameters on page 49](#)
- [Setting Task Level Parameters on page 50](#)
- [Working with Configuration Tags on page 51](#)
- [Considerations in Designing a Subject Area on page 54](#)
- [Creating a Subject Area on page 56](#)
- [Building and Running an Execution Plan with the DAC on page 57](#)
- [Creating a Micro ETL Execution Plan on page 58](#)
- [Scheduling an Execution Plan on page 60](#)
- [About Refresh Dates on page 60](#)
- [Monitoring Execution Plan Processes on page 61](#)

### Creating or Copying a Source System Container

The metadata for a source system is held in a container. You cannot change the metadata for preconfigured containers. If you want to customize the metadata in a preconfigured container, you must first make a copy of the container. The DAC keeps track of all customizations in the copied container, so that at any time you can find the newly created objects and modified objects, as well as the original objects.

You can also create a new, empty container if you want to build your own container with customized metadata.

### *To create a new container or copy an existing container*

- 1 In the DAC's menu bar, click File > New Source System Container.
- 2 Enter an ID and a name for the container.  
The ID and Name fields are alphanumeric. The Name can contain spaces but the ID cannot.
- 3 Select one of the following:
  - Create Empty New Source System Container
  - Create as a Copy of Source System Container
- 4 If you are creating an empty, new container, click OK.
- 5 If you are making a copy of an existing container, select the existing container from the drop-down list, and then click OK.

## About Customizing the Data Warehouse

You can add tables, columns, and indices to the data warehouse, and you can modify these existing objects. Customizing the data warehouse in this way requires using the DAC and Informatica client tools. For more information about using Informatica client tools to customize the Oracle Business Analytics Warehouse, see the *Oracle Business Intelligence Applications Installation and Configuration Guide*.

Figure 9 shows the major steps required for adding a new object to the data warehouse or modifying existing objects. As shown in Figure 9, you can begin the customization process by adding or modifying the new data warehouse object in the DAC and then using the DAC's Data Warehouse Configurator to create or update the object in the data warehouse. Alternatively, you can add or modify the object directly in the data warehouse database and then use the DAC's Import from Database command to add the new object in the DAC.

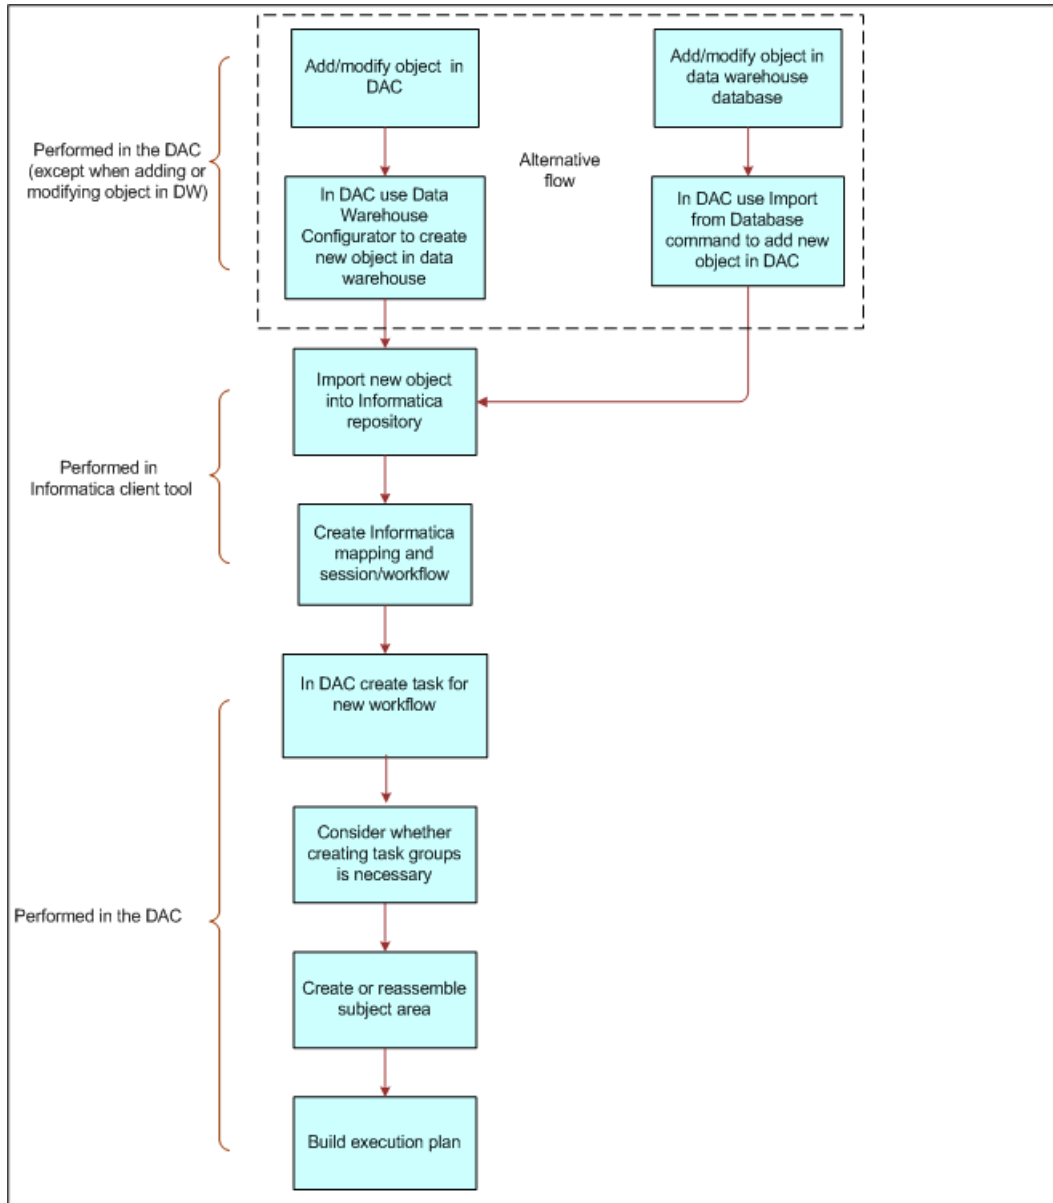


Figure 9. Process Flow to Add New Object to Data Warehouse

## Adding a New Table and Columns to the Data Warehouse

As shown in [Figure 9](#), there are two alternative process flows for adding a new object to the data warehouse. You can enter the table and column definitions in the DAC and then use the DAC's Data Warehouse Configurator to create the table and columns in the data warehouse database; for this method, follow the procedure, "[To add a new table and columns to the data warehouse using the DAC's Data Warehouse Configurator](#)" on page 44.

Alternatively, you can add the new table and column definitions directly in the data warehouse database and then use the DAC's Import from Database command to add the new table and columns in the DAC; for this method, follow the procedure, "[To add a new table and columns using the DAC's Import command](#)" on page 45.

### *To add a new table and columns to the data warehouse using the DAC's Data Warehouse Configurator*

- 1 In the DAC toolbar, select the appropriate source system container from the drop-down list in the toolbar.
  - 2 In the Menu bar, click Views > Design > Tables.
  - 3 Create the new table.
    - a In the Tables tab, click New.
    - b In the Edit child tab, enter the appropriate information about the table, and click Save.  
For a description of the fields in this tab, see "[About the DAC Tables Tab](#)" on page 92.
  - 4 Add the columns for the new table.
    - a In the Columns child tab, click New.
    - b Enter the appropriate column information for each column you want to add to the table, and click Save.
    - c Enter the appropriate foreign key table and column information.
- NOTE:** For performance purposes, it is recommended that you do not enter more than 254 columns to a dimension or fact table.
- 5 Create the new tables and columns in the data warehouse database.
    - a Choose Tools > ETL Management > Configure.
    - b Select the appropriate Source and Target database platforms, and then click OK.
    - c In the Data Warehouse Configuration Wizard, select Create Data Warehouse Tables, and then click Next.
    - d Enter the required information, and then click Start.

An informational message reports whether the process was successful. For information about the process, you can review the createwtables.log file in the OracleBI\DAC\log\config folder.

**To add a new table and columns using the DAC's Import command**

- 1 Add the new table and column definitions into the data warehouse database.
- 2 In the DAC toolbar, select the appropriate source system container from the drop-down list in the toolbar.
- 3 In the Menu bar, click Views > Design > Tables.
- 4 Import the new table definition.
  - a Right-click and select Import from Database > Import Database Tables.
  - b In the Import Tables dialog box, select DataWarehouse.
  - c Optionally, enter filter criteria to identify the table name you entered in [Step 1](#).  
See ["DAC Query Commands and Operators"](#) on page 35 for available filter commands and operators.
  - d Click Read Tables.
  - e In the list of tables displayed, select the Import check box for the tables you want to import.
  - f Click Import Tables.

An informational message indicates whether the process was successful.

- 5 Import the new column definitions.
  - a In the Tables tab, query for the table you imported in [Step 4](#).
  - b With the table highlighted, right-click and select Import from Database > Import Database Columns.
  - c In the Importing Columns... dialog box, select Selected Record Only, and then click OK.
  - d In the Import Table Columns dialog box, click Read Columns.

The Changes column displays a description of column changes, which are explained below:

Change	Explanation
The object was added to the database.	The column is in the database but not the DAC repository. Importing it will add the column to the DAC repository.
The object was added to the repository.	The column is in the DAC repository but not in the database. Importing it will delete it from the DAC repository.
The object was modified.	The column definition in the database doesn't

- e In the list of columns displayed, select the Import check box for the columns you want to import.
- f Click Import Columns.

An informational message indicates whether the process was successful.

## Adding an Index to the Data Warehouse

Follow this procedure to add a new index to the data warehouse.

### *To add a new index to the data warehouse*

- 1 Add the new index definition into the data warehouse database.
- 2 In the DAC toolbar, select the appropriate source system container from the drop-down list in the toolbar.
- 3 In the Menu bar, click Views > Design > Tables.
- 4 Query for the table for which you want to import index definitions.
- 5 Right-click and select Import from Database > Import Indices.
- 6 Choose to import indices for a selected table or for all the records retrieved in your query, and click OK.
- 7 In the Import Indices dialog box, select DataWarehouse from the Data Sources drop-down list.
- 8 Click Read Indices.
  - a In the list of indices displayed, select the Import check box for the indices you want to import.
  - b Click Import Indices.

An informational message indicates whether the process was successful.

## Importing New Data Warehouse Objects into the Informatica Repository

This step requires using Informatica client tools to import new data warehouse objects into the Informatica repository. For instructions on this step of customizing the data warehouse, see *Oracle Business Intelligence Applications Installation and Configuration Guide*.

## Creating Informatica Mappings and Workflows

This step requires using Informatica client tools to create new Informatica mappings and workflows for the data warehouse objects that you imported into the Informatica repository. For instructions on this step of customizing the data warehouse, see *Oracle Business Intelligence Applications Installation and Configuration Guide*.

## Creating Tasks in the DAC for New or Modified Informatica Workflows

You need to perform this step for all new workflows you create in Informatica and for all workflows that you modify.

- 1** In the DAC toolbar, select the appropriate source system container from the drop-down list in the toolbar.
- 2** In the DAC, create custom logical and physical task folders for the custom folder you created in the Informatica repository.
  - a** In the DAC, navigate to Tools > Seed Data > Task Folders.
  - b** To create a custom logical folder, click New.
  - c** In the Name field, enter a name for the custom logical folder, for example, Custom Logical.
  - d** In the Type field, select Logical.
  - e** To create a custom physical folder, click New.
  - f** In the Name field, enter a name for the custom physical folder, for example, Custom Physical.
  - g** In the Type field, select Physical.
- 3** Register the folders you created in [Step 2](#) in the Source System Folders tab.
  - a** Navigate to Design > Source System Folders.
  - b** Click New.
  - c** In the Edit child tab, enter the name of the custom logical folder in the Logical Folder field.
  - d** Enter the name of the custom physical folder in the Physical Folder field, and click Save.
- 4** Create new tasks for the workflows.
  - a** Navigate to Design > Tasks, and click New in the top pane toolbar.
  - b** In the Edit child tab, enter the workflow name as it appears in Informatica Workflow Manager.
  - c** Right-click and select Synchronize Tasks.
  - d** Select Selected Record Only, and click OK. Click OK in the informational message box.

This command synchronizes the source and target table information between the DAC and Informatica.
  - e** In the Tasks tab, enter the remaining information required for the task.

For a description of the fields in this tab, see [“About the DAC Tasks Tab” on page 97](#).

The new table is now ready to be associated with a subject area. For information about creating a subject area, see [“Creating a Subject Area” on page 56](#).

## Setting a Task Phase Dependency

A task phase dependency allows you to change the order in which tasks are executed. When you set a task phase dependency, you work with three properties:

- **Action.** The action to be taken in relation to the phase dependency. Possible values are the following:
  - **Wait.** Allows you to set a task to wait for other tasks of a specified phase to be completed before it runs.
  - **Block.** Allows you to block all tasks of a specified phase from being executed until the specified task has been executed.
- **Grain.** Applicable only for blocks. Allows you to specify whether the action you choose affects all tasks of a specified phase or related tasks. (You can view a task's related tasks by navigating to Execution Plans > All Dependencies and viewing the specified task's predecessor tasks.)

Possible values are:

- **All.** Indicates the action will affect *all* tasks.
- **Related.** Indicates the action will affect only *related* tasks. You can view a task's related tasks by navigating to Execution Plans > All Dependencies and viewing the specified task's predecessor tasks.
- **Phase.** The phase that applies to the action you choose.

### *To set a task phase dependency*

- 1 In the DAC toolbar, select the appropriate source system container from the drop-down list in the toolbar.
- 2 In the Menu bar, click Views > Design > Tasks.
- 3 Query for the task for which you want to add a phase dependency, and make sure it is highlighted.
- 4 Click the Phase Dependency child tab.
- 5 Click Add/Remove.
- 6 In the Choose Phases dialog box, select a phase and click Add.
- 7 Click OK in the message box that states the phase was added.
- 8 Select an Action and a Grain, click OK.

The task phase dependency appears in the Phase Dependency child tab.

- 9 Reassemble the appropriate subject area.
  - a In the Subject Areas tab, query for the appropriate subject area.
  - b Click Assemble.
- 10 Rebuild the execution plan.
  - a Navigate to the Execution Plans tab in the Execute view.



- b** Query for the appropriate execution plan.
- c** In the Parameters child tab, click Generate.
- d** In the top pane toolbar, click Build.

## Creating a Task Group

The DAC automatically organizes tasks into a dependency structure based on dependency rules. For information about the DAC's dependency rules, see ["How the DAC Determines the Order of Task Execution within an Execution Plan" on page 55](#). The DAC assigns priority randomly to tasks that have the same properties. You can use the Task Group feature to group such tasks that share the same properties and enforce a priority of your choosing.

This feature can be useful for the following: truncation and restartability purposes; when more than one task with similar properties writes to the same table; and when there is a circular read/write relationship between tables; for example, task 1 reads from table A and writes to table B, and task 2 reads from table B and writes to table A.

### *To create a task group*

- 1** In the DAC, select the appropriate source system container from the drop-down list in the toolbar.
- 2** In the Menu bar, click Views > Design > Task Groups.
- 3** Create a new task group.
  - a** Click New in the top pane toolbar.
  - b** In the Edit child tab, enter a name and select the appropriate properties.
- 4** Click the Child Tasks child tab, and click Add/Remove in the toolbar.
- 5** In the left-hand window of the Choose Child Tasks dialog box, query for the tasks you want to add to the task group.
- 6** Select the tasks, and click Add.
- 7** In the right-hand window, enter an execution order.
- 8** Click Save, and then click OK to close the window.

## Setting Source System Parameters

Source system parameters are specific to source system containers. You can set static or runtime parameters in the Source System Parameters tab of the Design view.

### *To set a source system parameter*

- 1** In the DAC, select the appropriate source system container from the drop-down list in the toolbar.

- 2 In the Menu bar, click Views > Design > Source System Parameters.
- 3 Click New in the top pane toolbar.
- 4 In the Edit child tab, enter a name for the parameter.
- 5 Select one of the following data types:
  - Text
  - Timestamp
  - SQL
- 6 Click in the Value field to open the Enter Parameter Value dialog box.
- 7 If you selected the Text data type:
  - a Select Static or Runtime.
  - b For a static parameter, enter the parameter information in the text box, and click OK.
  - c For a runtime parameter, select a variable, and click OK.
- 8 If you selected the Timestamp data type:
  - a Select Static or Runtime.
  - b Choose a function from the Function drop-down list.
    - If you select Custom Format, enter a format in the Format field.
    - If you select SQL Syntax, select a database connection in the Connection drop-down list.
  - c If you selected Static, click in the Date field, and enter a date and time.
  - d If you selected Runtime, select a variable from the Variable drop-down list.
- 9 If you selected the SQL data type:
  - a Select a logical data source.
  - b Enter the SQL statement in the text box, and click OK.
- 10 Click Save in the top pane toolbar or in the Edit child tab.

## Setting Task Level Parameters

In addition to source system parameters, you can set a parameter at the task level. Task level parameters take precedence over the source system parameters.

### *To set a parameter at the task level*

- Navigate to Design view > Tasks tab > Parameters child tab, and then follow the instructions for [“Setting Source System Parameters”](#) on page 49.

## Working with Configuration Tags

A configuration tag is an object that controls the inclusion of tasks in subject areas. When a task is tagged, it is not eligible to be included in the collection of tasks for any subject area, unless the tag is part of the subject area definition "Include Task" property.

A configuration tag can function in one of the following ways:

### ■ Remove tasks from all subject areas

If you assign a task to a configuration tag, the task will not be eligible to participate in *any* subject area. For instructions, see ["To remove tasks from all subject areas" on page 51](#).

### ■ Reassign autogenerated tasks to a specific subject area

For autogenerated tasks that were removed from participating in a subject area, you can set up the configuration tag to reassign a task to participate in specific subject areas. You do this by associating the configuration tag with the desired subject area. This method only applies to tasks that are autogenerated tasks of a subject area. For instructions, see ["To reassign autogenerated tasks to a subject area" on page 52](#).

### ■ Add non-autogenerated tasks to a subject area

You can set up a configuration tag to add non-autogenerated tasks to a subject area. The non-autogenerated tasks will participate in the subject area along with the subject area's autogenerated tasks. For instructions, see ["To add non-autogenerated tasks to a subject area" on page 52](#).

Indicates whether configuration tag tasks are the only tasks associated with this subject area that will participate in the ETL process. If this check box is selected, only the tasks associated with the configuration tag will be chosen by the DAC when the subject area is assembled.

### ■ Assign only configuration tag tasks to a subject area (excludes the subject area's autogenerated tasks)

You can also set up a configuration tag so that *only* tasks that were assigned to the configuration tag participate in a specific subject area. In this case, the subject area's autogenerated tasks do not participate. For instructions, see

### *To remove tasks from all subject areas*

- 1 In the DAC toolbar, select the appropriate source system container from the drop-down list in the toolbar.
- 2 In the Menu bar, click Views > Design > Configuration Tags.
- 3 Create a new configuration tag.
  - a Click New in the top pane toolbar.
  - b In the Edit child tab, enter a name.
  - c Make sure the Include Tasks check box is not selected.
  - d Click Save.
- 4 Add tasks to the configuration tag.

- a With the new configuration tag highlighted in the top pane, click the Tasks child tab.
- b In the bottom pane toolbar, click Add/Remove.
- c In the Tasks dialog box, query for the tasks you want to add to the configuration tag.
- d Highlight the tasks, and then click Add.  
The tasks appear in the right-hand window.
- e Click Save, and then click OK to close the window.  
These tasks will not be eligible to participate in any subject area.

### ***To reassign autogenerated tasks to a subject area***

- 1 In the DAC toolbar, select the appropriate source system container from the drop-down list in the toolbar.
- 2 In the Menu bar, click Views > Design > Configuration Tags.
- 3 Query for the configuration tag that contains the tasks you want to reassign to a subject area.
- 4 Verify the configuration tag contains the appropriate tasks by clicking the Tasks child tab and reviewing the list of tasks associated with this configuration tag.  
**NOTE:** Only a subject area's autogenerated tasks will be reassigned. If non-autogenerated tasks appear in the list, the DAC will ignore them.
- 5 Associate the configuration tag with the subject areas to which you want to reassign the tasks.
  - a With the configuration tag highlighted in the top pane, click the Subject Areas child tab.
  - b Click Add/Remove in the bottom pane toolbar.
  - c In the Subject Areas dialog box, query for one or more subject areas to which you want to reassign the task or tasks.
  - d Highlight the appropriate subject areas, and click Add.
  - e Click Save, and then click OK to close the window.
- 6 Reassemble the subject area.
  - a In the Subject Area tab, query for all the subjects areas you added to the configuration tag.
  - b Highlight the subject areas, and click Reassemble.

### ***To add non-autogenerated tasks to a subject area***

- 1 In the DAC toolbar, select the appropriate source system container from the drop-down list in the toolbar.
- 2 In the Menu bar, click Views > Design > Configuration Tags.
- 3 Create a new configuration tag.
  - a Click New in the top pane toolbar.
  - b In the Edit child tab, enter a name.

- c Select the Include Tasks check box.
        - d Click Save.
- 4 Add the non-autogenerated tasks to the configuration tag.
  - a With the new configuration tag highlighted in the top pane, click the Tasks child tab.
  - b In the bottom pane toolbar, click Add/Remove.
  - c In the Tasks dialog box, query for the extraneous tasks you want to add to the configuration tag.
  - d Highlight the tasks, and then click Add.
  - e Click Save, and then click OK to close the window.
- 5 Associate the configuration tag with the subject areas to which you want to add the non-autogenerated tasks.
  - a With the configuration tag highlighted in the top pane, click the Subject Areas child tab.
  - b Click Add/Remove in the bottom pane toolbar.
  - c In the Subject Areas dialog box, query for one or more subject areas to which you want to add the non-autogenerated tasks.
  - d Highlight the appropriate subject areas, and click Add.
  - e Click Save, and then click OK to close the window.
- 6 Reassemble the subject area.
  - a In the Subject Area tab, query for all the subjects areas you added to the configuration tag.
  - b Highlight the subject areas, and click Reassemble.

***To assign only configuration tag tasks to a subject area (excludes the subject area's autogenerated tasks)***

- 1 In the DAC toolbar, select the appropriate source system container from the drop-down list in the toolbar.
- 2 In the Menu bar, click Views > Design > Subject Areas.
- 3 Query for the subject area to which you want to add configuration tag tasks.

**NOTE:** The autogenerated tasks for this subject area will be excluded.
- 4 Select the Configuration Tag Tasks Only check box, and click Save.
- 5 Create a configuration tag.
  - a Navigate to the Configuration Tags tab.
  - b Click New in the top pane toolbar.
  - c In the Edit child tab, enter a name.
  - d Select the Include Tasks check box.
  - e Click Save.
- 6 Add the tasks to the configuration tag.

- a With the new configuration tag highlighted in the top pane, click the Tasks child tab.
  - b In the bottom pane toolbar, click Edit.
  - c In the Tasks dialog box, query for the tasks you want to add to the configuration tag.
  - d Highlight the tasks, and then click Add.
  - e Click Save, and then click OK to close the window.
- 7 Associate the configuration tag with the subject area.
  - a With the configuration tag highlighted in the top pane, click the Subject Areas child tab.
  - b Click Add/Remove in the bottom pane toolbar.
  - c In the Subject Areas dialog box, query for the appropriate subject area.
  - d Highlight the subject area, and click Add.
  - e Click Save, and then click OK to close the window.
- 8 Reassemble the subject area.
  - a In the Subject Area tab, query for all the subjects areas you added to the configuration tag.
  - b Highlight the subject areas, and click Reassemble.

## Considerations in Designing a Subject Area

Oracle Business Intelligence Applications provides preconfigured subject areas. You can change these preconfigured subject areas or create new subject areas to correspond to your particular business processes.

**NOTE:** To change a preconfigured subject area or to create a new subject area, you must first make a copy of an existing source system container or create a new container. For instructions, see [“Creating or Copying a Source System Container”](#) on page 41.

### Designing a Subject Area

In designing a subject area, you should consider the following questions:

- **Tables.** Which tables need to be populated for the data warehouse? From which tables does your organization source data? What tables will create the star schemas.
- **Subject areas.** Do the subject areas cover all the relevant tables?
- **Tasks.** Are the tasks that load this table defined?
- **Indices.** Do the target tables have the correct indices defined?

### How the DAC Determines Tasks Required for Subject Areas

A subject area is a collection of tasks. When a subject area is defined, the DAC uses the following logic to assemble the collection of tasks:

- 1** Initial selection of tables.  
Find all the fact tables that belong to the subject areas.
- 2** Recursive selection of related tables.  
Recursively find all the tables directly related through foreign keys and all other logically related tables.
- 3** Initial selection of tasks.  
Find all the tasks that load into the tables selected above, that is, tasks whose target tables are one of the tables identified above.
- 4** Recursive selection of all tasks.  
Depending on the source and target table relationships, recursively figure out the prerequisite tasks.

## How the DAC Determines the Order of Task Execution within an Execution Plan

An execution plan is a collection of subject areas and a unique collection of tasks. A task can have prerequisite tasks that need to be executed before its own execution. The DAC determines the order of tasks based on the following considerations:

- A task's source and target table

The DAC server first looks at a task's source and target table. For example, suppose table A is populated by task T1 by reading from table B, and table B is populated by task T2 by reading from table C. The DAC server would determine task T2 should be executed before T1.

The DAC server next considers the following:

- Task phase

An ETL process typically goes through several phases. An example of a typical order in which phases are executed is as follows:

- a** Extract Dimension
- b** Extract Fact
- c** Load Dimension
- d** Load Fact and Load Hierarchy (executed in parallel)
- e** Load Aggregate tables
- f** Update Dimensions

- A table's Truncate Always properties

The order of execution based on Truncate Always properties is as follows:

- a** Insert
- b** Upsert
- c** Physical data source

**d** Priority

- The DAC randomly organizes tasks that are the same in all properties. If some tasks need to be executed in a particular order, the DAC allows you to create a task group in which you can specify an execution order.

## Creating a Subject Area

When you create a new subject area, you assign one or more fact tables to the subject area. The DAC then determines which dimension and other related tables are required as well as the tasks and their order of execution.

### *To create a new subject area*

- 1 In the DAC, select the appropriate source system container from the drop-down list in the toolbar.
- 2 In the Menu bar, click Views > Design > Subject Areas.
- 3 In the top pane toolbar, click New.
- 4 In the Edit child tab, enter a name for the subject area, and click Save.
- 5 Make sure the new subject area name is highlighted in the top pane, and click the Tables child tab.
- 6 Click Add/Remove in the child tab toolbar.  
The Choose Tables dialog box opens. The left-hand window lists all the tables held in the selected container.
- 7 Query for one or more fact tables.
- 8 Select the fact table (use Shift+click to select more than one table), and click Add.  
The tables are added to the right-hand window, which represents the subject area.
- 9 Click OK to close the Choose Tables dialog box.
- 10 In the top pane toolbar, click Assemble.
- 11 In the Assembling... dialog box, select Selected Record Only.

If you select the option All Records in the List, the DAC will reassemble all the subject areas listed in the top pane.

The DAC assembles the selected subject area by determining what dimensions and other related tables are required and what tasks are needed to load these tables.

You will receive an informational message when the assemble process is completed.



- 12 Click the Tasks tab to view which tasks the DAC has determined are required for this subject area.

Tasks that are automatically assigned to the subject area by the DAC are indicated with the Autogenerated check mark.

You can inactivate a task from participating in the subject area by selecting the Inactive check box. When the Inactive check box is selected, the task remains inactive even if you reassemble the subject area.

You can also remove a task from the subject area using the Add/Remove command, but when you remove a task it is only removed from the subject area until you reassemble the subject area.

## Building and Running an Execution Plan with the DAC

Execution plans are subject areas that are used to execute ETL processes. Before you attempt to run an execution plan, make sure you have completed the following:

- Set database connections to the transactional and data warehouse databases
- Registered the Informatica servers and repository server

### *To build and run an execution plan*

- 1 Navigate to Execute > Execution Plans.
- 2 Create a new execution plan.
  - a In the top pane toolbar, click New.
  - b In the Edit child tab, enter a name for the execution plan and other appropriate information.  
For a description of the fields in this tab, see [“About the DAC Execution Plans Tab” on page 115](#).
  - c Click Save.
- 3 Associate one or more subject areas with the execution plan.
  - a Click the Subject Areas child tab.
  - b Click Add/Remove in the bottom pane toolbar.
  - c In the Choose Subject Areas dialog box, select the appropriate source system container.
  - d Query for the subject area you want to associate with the execution plan.
  - e Select the subject area and click Add.  
You can associate multiple subject areas with an execution plan, but all the subject areas must be from the same source system container.
  - f Click OK to close the window.
- 4 Generate the runtime parameters.
  - a Click the Parameters child tab.

- b** Click Generate in the bottom pane toolbar.
  - c** For each Datasource type, enter the appropriate name in the Value field.
- 5** In the top pane of the Execution Plans tab, make sure the new execution plan is highlighted, and click Build.
- The DAC builds the execution plan.
- 6** To run the execution plan, select the execution plan in the top pane, and click Run Now.
- Once the ETL process starts running you can monitor its progress in the Current Run tab.
- For information about how refresh dates are tracked, see [“About Refresh Dates” on page 60](#).
- To schedule an execution plan, see [“Scheduling an Execution Plan” on page 60](#).

## Creating a Micro ETL Execution Plan

Micro ETL execution plans are ETL processes that you schedule at very frequent intervals, such as hourly or half-hourly. They usually handle small subject areas or subsets of larger subject areas. The DAC tracks refresh dates for tables in micro ETL execution plans separately from other execution plans and uses these refresh dates in the change capture process.

After a micro ETL execution plan runs, the DAC populates refresh date values in the Refresh Dates child tab of the Execution Plans tab. If a subject area is used in a regular execution plan (an execution plan with the Keep Separate Refresh Dates option not selected) as well as a micro ETL execution plan, the DAC maintains refresh dates for the tables in the regular execution plan in the Refresh Dates child tab of the Physical Data Sources tab (Setup view).

In cases of a subject area being used in both a regular and micro ETL execution plan and the micro ETL execution plan is suspended for a few days but the regular execution plan runs nightly, the DAC automatically detects the last refresh date for the tables common to both execution plans and intelligently extracts only the most recent records for the micro ETL execution plan.

**CAUTION:** Micro ETL processes can cause issues with data inconsistencies, data availability, and additional load on the transactional database. Therefore, you should consider the following factors before implementing a micro ETL process:

- For related star schemas, if one schema is omitted from a micro ETL execution plan, the cross-star reports may be inaccurate. For example, if the Person fact table is refreshed more frequently than the Revenue fact table, a report that spans the Person and Revenue dimensional schemas may produce inconsistent results.
- If you omit dimension tables from a micro ETL execution plan, the foreign keys for the fact tables will point to Unspecified rows for the new dimension records. The foreign key references will be resolved when the Complete ETL execution plan is run, but users of the reports should be aware of such inconsistencies.
- If you do not include aggregate tables in micro ETL execution plans, the reports that use data from these tables will be inconsistent with the reports that use data from the detailed fact tables. However, if aggregate tables are included in the micro ETL execution plan, the aggregate calculations are performed for each ETL process, which will take a constant amount of time and may be inefficient to perform at such frequent intervals.

- Hierarchy tables are rebuilt during every ETL execution plan by querying the base dimension tables. This operation takes a constant amount of time. If the base tables are big, this operation may take a long time and may be inefficient if the micro ETL execution plan runs several times a day. However, if you avoid populating the hierarchy tables during micro ETL processes, data inconsistencies will occur.
- With micro ETL execution plans, caching will occur more frequently, which may have performance implications.
- Micro ETL execution plans will put more load on the transactional database because of the frequent extracts.

### *To create a micro ETL execution plan*

- 1 In the DAC toolbar, select the appropriate source system container from the drop-down list in the toolbar.
- 2 In the Menu bar, click Views > Design > Subject Areas.
- 3 In the Subject Areas tab, assemble a small subject area.
- 4 In the Tasks child tab, inactivate all tasks that are not required for the execution plan.
- 5 Create a new execution plan.
  - a Navigate to Execute > Execution Plans.
  - b Enter a name for the execution plan
  - c Select the Keep Separate Refresh Dates check box.
  - d Click Save.
- 6 Associate one or more subject areas with the execution plan.
  - a Click the Subject Areas child tab.
  - b Click Add/Remove in the bottom pane toolbar.
  - c In the Choose Subject Areas dialog box, select the appropriate source system container.
  - d Query for the subject area you want to associate with the execution plan.
  - e Select the subject area and click Add.

You can associate multiple subject areas with an execution plan, but all the subject areas must be from the same source system container.
  - f Click OK to close the window.
- 7 Generate the runtime parameters.
  - a Click the Parameters child tab.
  - b Click Generate in the bottom pane toolbar.
  - c For each Datasource type, enter the appropriate name in the Value field.

- 8 In the top pane of the Execution Plans tab, make sure the new execution plan is highlighted, and click Build.  
The DAC builds the execution plan.
- 9 Navigate to the Ordered Tasks child tab and verify the sequence of tasks.  
The execution plan is now ready to run as a micro ETL execution plan.
- 10 Create a schedule for the micro ETL execution plan. For instructions, see [“Scheduling an Execution Plan” on page 60](#).

## Scheduling an Execution Plan

Use this procedure to schedule an execution plan in the DAC.

### *To schedule an execution plan*

- 1 In the DAC, navigate to the Scheduler tab.  
The current list of schedules appears in the top pane.
- 2 Click New in the top pane toolbar.  
The Edit tab in the bottom pane becomes active.
- 3 Enter a name for the schedule.
- 4 Select an execution plan.
- 5 If you want the schedule to run once, select the Run Only Once check box, and then select a start and end date and time.
- 6 To create a periodic schedule, select a recurrence pattern, and enter the appropriate date and time parameters.
- 7 Click Save.

## About Refresh Dates

Refresh dates refer to the date of the last ETL process (the last time data was extracted from tables in a given database or loaded into tables in a given database). The DAC uses the refresh dates to determine whether to run the incremental load commands or to run full load commands and whether to truncate the target tables.

Refresh dates are tracked only for tables that are either a primary source or a primary target on tasks in a completed run of an execution plan. The DAC runs the full load command for tasks on which a table is a primary source or target if the refresh date against the table is null. When there are multiple primary sources, the earliest of the refresh dates will trigger a full load or an incremental load. If any one of the primary source tables has no refresh date, then the DAC will run the full load command.

Table 12 shows the possible scenarios regarding refresh dates.

Table 12. Refresh Date Scenarios

Scenario	Table Type (in Tasks child tabs)	Refresh Date	Command DAC Will Use	Truncate Target Table?
1	Primary Source	Null	Full Load	Yes
	Primary Target	Null		
2 (See note below)	Primary Source	Null	Full Load	No
	Primary Target	Not Null		
3 (See note below)	Primary Source	Not Null	Full Load	Yes
	Primary Target	Null		
4	Primary Source	Not Null	Incremental Load	No
	Primary Target	Not Null		

- **Scenario 2.** When two or more source tables load into the same target table as separate tasks, the source table in the second task may have refresh date as null while the target may have a refresh date.
- **Scenario 3.** When a source loads into more than one target table in separate tasks, the refresh date may be null for the second target table while the source table may have refresh dates.

## Monitoring Execution Plan Processes

The Current Run tab in the Execute view provides predefined reports that allow you to monitor execution plan processes in order to isolate bottlenecks and enhance performance.

### *To monitor an execution plan process*

- 1 In the DAC, navigate to the Current Run tab.
- 2 Right-click and select Get Run Information.

The following options are available:

- Get log file
- Analyze run
- Get chart
- Get phase chart
- Get graph



# 6

## Common Tasks Performed in the DAC

This chapter contains the following topics:

- [Importing DAC Metadata on page 63](#)
- [Exporting DAC Metadata on page 64](#)
- [Distributing DAC Metadata on page 65](#)
- [Running the DAC Server Automatically on page 65](#)
- [Command Line Access to the DAC Server on page 66](#)
- [DAC Repository Command Line Options on page 70](#)
- [Replacing an Informatica Workflow with a Custom SQL File on page 72](#)
- [Determining the Number of Transactional and Data Warehouse Database Connections on page 75](#)
- [Running Two DAC Servers on the Same Machine on page 75](#)
- [Customizing Index and Analyze Table Syntaxes on page 76](#)
- [Using SQL Files as an Execution Type in the DAC on page 77](#)
- [Overview of Change Capture Process \(Siebel Sources Only\) on page 78](#)
- [Using the Change Capture Filter on page 84](#)
- [Tracking Deleted Records on page 84](#)
- [Handling ETL Failures with the DAC on page 86](#)

### Importing DAC Metadata

The DAC's Import/Export feature allows you to import or export source system-specific DAC metadata into or out of the DAC repository. You can use this feature to migrate DAC metadata from one environment to another, such as from a development environment to test or production environments.

#### *To import DAC metadata*

- 1 In the DAC menu bar, select Tools > DAC Repository Management > Import.
- 2 Select the directory from which you want to import DAC metadata, or accept the default directory.
- 3 Select the appropriate source system containers.
- 4 Select the appropriate categories of metadata you want to import:

- **Logical.** Imports all information contained in the Design view and database connection information.
  - **System.** Imports all information contained in the Setup view, except passwords for servers and database connections.
  - **Run Time.** Imports information about ETL runs (contained in the Execute view)
- 5 If you are importing metadata into a blank repository or to completely replace the current metadata in the repository, select Truncate Repository Tables.  

This action overwrites the content in the current repository. Selecting the Truncate Repository Tables option greatly increases the speed of the import process.
  - 6 (Optional) Select Enable Batch Mode to insert the imported metadata into the repository as an array insert.  

This action increases the speed of the import process.
  - 7 Click OK.
  - 8 Verify the import process by reviewing the log file `\OracleBI\DAC\log\import.log`.

## Exporting DAC Metadata

The DAC's Import/Export feature allows you to import or export source system-specific DAC metadata into or out of the DAC repository. You can use this feature to migrate DAC metadata from one environment to another, such as from a development environment to test or production environments.

### *To export DAC metadata*

- 1 In the DAC menu bar, select Tools > DAC Repository Management > Export.
- 2 Select the directory to which you want to export DAC metadata, or accept the default directory.
- 3 Select the appropriate source system containers.
- 4 Select the appropriate categories of metadata you want to export:
  - **Logical.** Exports all information contained in the Design view and database connection information.
  - **System.** Exports all information contained in the Setup view, except passwords for servers and database connections.
  - **Run Time.** Exports information about ETL runs (contained in the Execute view).
- 5 Click OK.
- 6 Verify the export process by reviewing the log file `\OracleBI\DAC\log\export.log`.



## Distributing DAC Metadata

Typically, you may have multiple environments, such as development, QA, production, and so on. When you make changes to the development environment, you test it, and then deliver it, exporting the whole environment and distributing it to the other environments. The data is exported as XML files, which are stored in the DAC\export directory on the client machine where the export is done.

To apply changes from the development environment to any other, you copy all of the XML files into the DAC\export folder and then import the data. To export the DAC metadata, follow the instructions in the procedure, [“Exporting DAC Metadata” on page 64](#). To import the DAC metadata, follow the instructions in the procedure, [“Importing DAC Metadata” on page 63](#).

## Running the DAC Server Automatically

Follow this procedure to set up the DAC server to be run automatically when your machine reboots.

### *To set up the DAC server to run automatically upon rebooting the machine*

- 1 Navigate to Programs > Accessories > System Tools > Scheduled Tasks.
- 2 Double-click Add Scheduled Task.
- 3 In the Scheduled Task Wizard, browse to the startserver.bat file, and click Open.
- 4 Select the option “When my computer starts,” and click Next.
- 5 Enter the domain user account to start the DAC server and a password, and click Finish.  
The startserver task appears in the Scheduled Task window.
- 6 Right-click the task and select Properties.
- 7 In the Settings tab, remove the check from the “Stop the task if it runs for 72 hours” check box.

### *To start the DAC server as a scheduled task*

- 1 Navigate to Programs > Accessories > System Tools > Scheduled Tasks.
- 2 Right-click startserver, and then click Run.

### *To stop the DAC server as a scheduled task*

- 1 Navigate to Programs > Accessories > System Tools > Scheduled Tasks.
- 2 Right-click startserver, and then click End Task.

### *To check if the DAC server is running*

- 1 Navigate to Programs > Accessories > System Tools > Scheduled Tasks.
- 2 Select the startserver task.
- 3 In the Windows menu bar, select View > Details.

## Command Line Access to the DAC Server

This section covers accessing the DAC server through a command line. It includes the following topics:

- [Setting Up Command Line Access to the DAC Server on page 67](#)
- [Using the Command Line to Access the DAC Server on page 68](#)

You can access the DAC server through a command line to start and stop execution plans and to get status information for servers, databases, and execution plans. This feature allows you to access the DAC server using third-party administration or management tools, without using the DAC client.

### Command Line Operations

The command line feature allows you to start an execution plan and stop the operation of a running execution plan.

#### Starting an Execution Plan

When the DAC server receives a request to start an execution plan, it performs a series of checks to verify that the execution plan can be started. It first checks that an execution plan with the requested name exists and that the execution plan is active. Next, it checks the status of the execution plan that last ran. If an execution plan is still running and the DAC server receives a request to start another execution plan, the request will be rejected. If an execution plan failed, a request to run the same execution plan again will be executed; however, a request to run a different execution plan will be rejected. If the execution plan that last ran completed successfully, a request to run a new execution plan will be executed.

When the DAC server receives a request to start an execution plan, it will issue a warning if any of the following conditions are true. (A warning is for informational purposes and does not mean the execution plan will not start.)

- The Generic task concurrency limit (set in the system properties) is not a positive number.
- There are no active Informatica servers in the server list.
- One or more Informatica servers do not have the passwords defined.
- One or more Informatica servers do not have a maximum number of sessions properly set.
- One or more data sources do not have the passwords defined.
- One or more data sources do not have a maximum number of connections properly set.
- One or more data sources do not have a number defined (set in Physical Data Sources tab).

#### Stopping the Operation of a Running Execution Plan

When the DAC server receives a request to stop the operation of a running execution plan, the request will fail in the following cases:

- The name of the execution plan that is running is different from the name in the request.
- There is no execution plan currently running.

## Command Line Status Monitoring Queries

The command line feature allows you to get the following status information:

- Summary of the requested execution plan. If there are multiple instances of the same execution plan, a summary of the instance that last ran is returned. Below is an example of the information contained in the summary.

(c) 2003 Siebel Systems, Inc.

Siebel DAC Server comprising the etl execution-management, scheduler, logger, and network server.

ETL details for the last run:

```
ETL Process Id : 255
ETL Name : Complete ETL
Run Name : DRY RUN OF Complete ETL: ETL Run - 2004-06-17 18:30:13.201
DAC Server : (aqamarD510)
DAC Port : 3141
Status: Stopped
Log File Name: Complete_ETL.255.log
Database Connection(s) Used :
```

```
OLTP jdbc:mscrosft:sql server://vranganaw8:1433;DatabaseName=OLTP
Data Warehouse jdbc:mscrosft:sql server://vranganaw8:1433;DatabaseName=olap
```

Informatica Server(s) Used :

```
InformaticaServer4-vranganaw8: (4)
InformaticaServer2-vranganaw8: (2)
InformaticaServer3-vranganaw8: (3)
InformaticaServer1-vranganaw8: (10)
```

```
Start Time: 2004-06-17 19:00:06.885
Message: ETL was interrupted
Actual Start Time: 2004-06-17 18:30:13.357
End Time: 2004-06-17 19:05:56.781
Total Time Taken: 35 Minutes
```

```
Start Time For This Run: 2004-06-17 19:00:06.885
Total Time Taken For This Run: 5 Minutes
```

```
Total steps: 212
Running steps: 0
Complete steps: 142
Failed/Stopped steps: 70
```

- Summary of connection status to all active databases and Informatica servers.

## Setting Up Command Line Access to the DAC Server

Follow this procedure to set up the command line functionality.

*To set up command line access to the DAC server*

- 1 Make sure you have installed the supported version of the Java SDK.
- 2 Copy the following files from the OracleBI\DAC directory to a local directory:
  - DAWSystem.jar
  - dac.properties
  - dacCmdLine.bat
- 3 In the dacCmdLine.bat file, edit the JAVA\_HOME variable to point to the directory where the Java SDK is installed.  
Make sure there are no spaces in the path reference.
- 4 Edit the DAC\_HOME variable to point to the directory where DAC-REMOTE is installed.
- 5 In the dac.properties file, edit the following parameter values.

Parameter	Value
ServerHost=	Host name of the DAC server.
ServerPort=	Port of the DAC server. The default is 3141.
RepositoryStampVal=	Repository stamp that appears in the DAC client Login Details screen.  To find this value, in the DAC client navigate to Help > Login Details.

Your dac.properties file should look similar to the following:

```
ServerHost=vranganaw8
ServerPort=3141
RepositoryStampVal =851E0677D5E1F6335242B49FCCd6519
```

## Using the Command Line to Access the DAC Server

Follow this procedure to use the command line to access the DAC server.

*To use the command line to access the DAC server*

- At the command prompt, enter the following:

`dacCmdLine <method name> <optional execution plan name>`

where `method name` is one of the following:

Method Name	Description
StartETL	Starts an execution plan. You must specify an execution plan name.
StopETL	Stops the operation of an execution plan. You must specify an execution plan name.
ETLStatus	If you do not specify an execution plan name, the status of the execution plan that last ran is returned. If you specify an execution plan name, the status of the specified execution plan is returned.
DatabaseStatus	Verifies whether the DAC server can connect to all active database connections. You do not need to specify an execution plan name.
InformaticaStatus	Verifies whether the DAC server is able to ping all active Informatica servers.

**NOTE:** The method names are case insensitive. Execution plan names are case sensitive. Also, if the execution plan name contains spaces, place beginning and ending double quotes around the name.

For example:

Command Line	Description
<code>dacCmdLine EtlStatus</code>	Returns the status of the execution plan that last ran.
<code>dacCmdLine EtlStatus Forecast</code>	Returns the status of the last instance of the Forecast execution plan.
<code>dacCmdLine StopEtl Forecast</code>	If the execution plan currently running is Forecast, the operation will be terminated. Otherwise, the request is ignored.
<code>dacCmdLine databasestatus</code>	Returns the health status of all the database connections as defined in the DAC repository from the DAC server.
<code>dacCmdLine InformaticaStatus</code>	Returns the health status of all the Informatica server connections as defined in the DAC repository from the DAC server.

## DAC Repository Command Line Options

This section describes the DAC repository command line parameters that are exposed by the AutomationUtils.bat file, which is located in the OracleBI\DAC folder.

### Import DAC Metadata by Application

The IMPORT option imports DAC metadata into the DAC repository for specified source system containers. The import process truncates all imported tables. You cannot perform an incremental import with this command.

Syntax:

```
IMPORT <fol derName> <contName1> <contName2> . . .
```

where:

fol derName	Full path to the root of the import file structure.
contName	(Optional) Name of the source system container for which you want to import DAC metadata. If no container is named, all containers that are found in the file structure will be imported.

### Export DAC Metadata by Application

The EXPORT option exports DAC metadata from the DAC repository for specified source system containers.

Syntax:

```
EXPORT <fol derName> <contName1> <contName2> . . .
```

where:

fol derName	Full path to the root of the export file structure.
contName	(Optional) Name of the source system container for which you want to export DAC metadata. If no container is named, all containers that are found in the file structure will be exported.

### Import DAC Metadata by Categories

The IMPORTCATEGORY option imports DAC metadata into the DAC repository based on the Logical, Run Time, or System categories. The import process truncates all imported tables. You cannot perform an incremental import with this command.

Syntax:

```
IMPORTCATEGORY <fol derName> <l ogi cal FI ag> <runti meFI ag> <systemFI ag>
```

where:

fol derName	Full path to the root of the import file structure.
-------------	---

logicalFlag	If the value of this parameter is true, all data categorized as logical is imported (information contained in the DAC Design view). Otherwise, this parameter is ignored.
runtimeFlag	If the value of this parameter is true, all data categorized as run time is imported (information contained in the DAC Execute view). Otherwise, this parameter is ignored.
systemFlag	If the value of this parameter is true, all data categorized as run time is imported (information contained in the DAC Setup view). Otherwise, this parameter is ignored.

### Export DAC Metadata by Categories

The EXPORTCATEGORY option exports DAC metadata from the DAC repository based on the Logical, Run Time, or System categories.

Syntax:

```
EXPORTCATEGORY <folderName> <logicalFlag> <runtimeFlag> <systemFlag>
```

where:

folderName	Full path to the root of the import file structure.
logicalFlag	If the value of this parameter is true, all data categorized as logical is exported (information contained in the DAC Design view). Otherwise, this parameter is ignored.
runtimeFlag	If the value of this parameter is true, all data categorized as run time is exported (information contained in the DAC Execute view). Otherwise, this parameter is ignored.
systemFlag	If the value of this parameter is true, all data categorized as run time is exported (information contained in the DAC Setup view). Otherwise, this parameter is ignored.

### Create Schema

The CREATESCHEMA option creates the schema of a new DAC repository.

Syntax:

```
CREATESCHEMA <unicodeFlag> <workspace name>
```

where:

unicodeFlag	If the value of this parameter is true, the schema is created as unicode. If the value is false, it is not created as unicode.
workspace name	The name of the workspace in which the schema is created.

### Drop Schema

The DROPSHEMA option drops the schema of the DAC repository.

Syntax:

```
DROPSHEMA
```

### Analyze

The ANALYZE option analyzes the DAC repository tables.

Syntax:

```
ANALYZE
```

### Upgrade

The UPGRADE option upgrades the DAC repository.

Syntax:

```
UPGRADE
```

### Set Password

The SETPASSWORD option sets the passwords for the Informatica servers and physical data sources in the DAC repository.

Syntax:

```
SETPASSWORD <type> <Logical Name> <password>
```

where:

type                      Possible values are server or dbconn.

Logical Name            Logical name of the server or data source record in the DAC.

**NOTE:** If the logical name or password contains spaces, quotes are required.

## Replacing an Informatica Workflow with a Custom SQL File

You can improve the performance of loads by replacing Informatica workflows with custom SQL files.

### *To replace an Informatica workflow with a custom SQL file*

- 1 Create a SQL file to be used to load the table, and unit test it.



- 2 Create an XML or SQL file with one or more SQL statements in the format that the DAC can understand.

For more information about creating an XML or SQL file, see [“Using SQL Files as an Execution Type in the DAC” on page 77](#).

You can create one file for a full load and one for an incremental load, or you can use the same file for both full and incremental loads.

- 3 Save the file in the OracleBI\DAC\CustomSQLs directory.
- 4 In the Tasks tab of the DAC Design view, query for the task for which you want to replace the Informatica workflow.
- 5 Replace the workflow name in the Command for Incremental Load or Command for Full Load fields with the XML or SQL file.
- 6 Change the Execution Type to SQL.

## Determining the Informatica Server Maximum Sessions Parameter Setting

You set the Maximum Sessions parameter value when you register the Informatica Server in the DAC client. This parameter specifies the maximum number of workflows that can be executed in parallel on the Informatica server. If the number of sessions is zero or is not specified, the DAC server assigns the default value of 10.

You should consider the following factors when determining the Maximum Sessions parameter value:

- How powerful the machine is that hosts the Informatica Server.
- How many Informatica Server servers are available.

- The number of Runnable tasks in the queue. A Runnable task is a task for which the Depends On tasks have completed and is ready to be run but is waiting for an Informatica slot to be available. For information about the different task run statuses, see [“About the DAC Current Run Tab” on page 119](#).

For an optimal run, the runnable queue should be at zero or should reach zero in a short time. For example, [Figure 10](#) shows an ideal performance run when 15 sessions were run in parallel. There were many runnable tasks before the process began, but the queue soon reached zero.

You can generate a run analysis such as [Figure 10](#) from the right-click menu (Get Run Information > Get Graph) on the DAC Current Run and Run History tabs. If you find that the runnable curve does not reach zero in a short time, you should increase the Maximum Sessions parameter value to make more Informatica slots available.

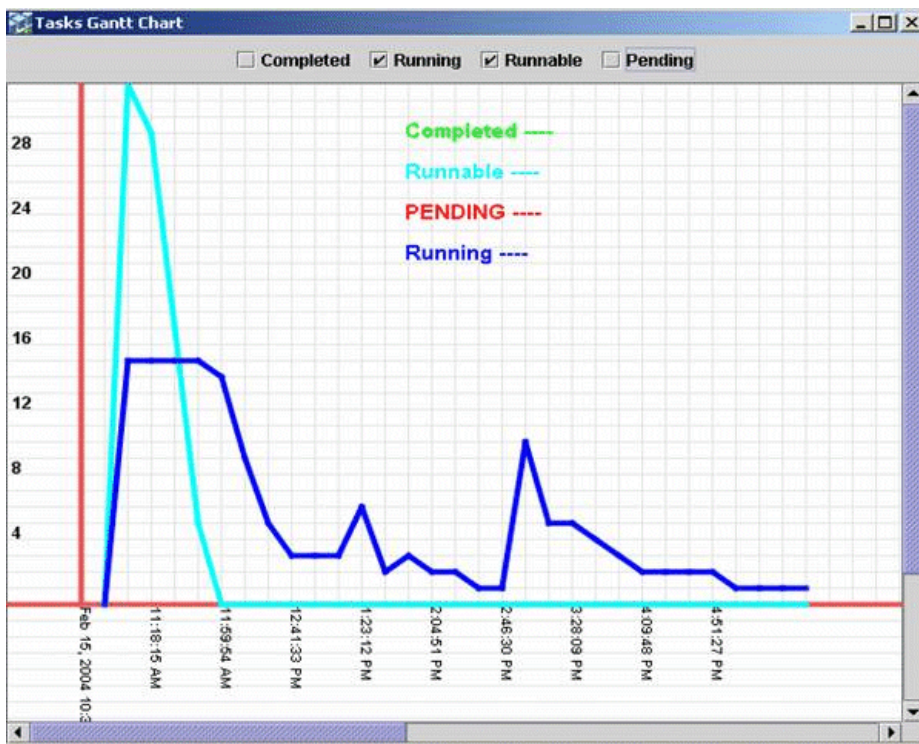


Figure 10. Sample Performance Run

## Determining the Number of Transactional and Data Warehouse Database Connections

This section describes how to determine the maximum number of database connections you need between the DAC server and the transactional database and the DAC server and the data warehouse database. You set the Max Num Connections parameter when you create the transactional and data warehouse database connections.

For the transactional database, the DAC server uses these connections to perform change capture. The number of connections you set for this connection pool determines how many change capture processes can run concurrently. If you have a powerful transactional database server and are going to run ETL processes during off-peak times, you can increase the Max Num Connections setting to 15 or 20 (10 is the default). If you have a less powerful transactional database server, you should not overload the operational system with ETL processes. Therefore, you should set the value below 10.

For the data warehouse database, the DAC server uses these connections to perform processes such as truncate tables, drop and create indices, and analyze tables. You should not set the Max Num Connections value higher than the Maximum Sessions parameter value (the maximum number of workflows that can be executed in parallel on the Informatica server) because these values have a one to one relationship.

## Running Two DAC Servers on the Same Machine

You can run two DAC servers on the same machine as long as they are listening on different ports and pointing to two different repositories.

### *To run two DAC servers on the same machine*

- 1 Install instances of the DAC server in two different directories.
- 2 For each instance, edit the config.bat file to set the DAC\_HOME variable appropriately.
- 3 Launch each of the DAC clients by navigating to the DAC directories and double-clicking the startclient.bat file.
- 4 For each instance, configure the DAC repository connection.
  - a Navigate to Tools > DAC Server Management > DAC Server Setup.  
An informational dialog box states this operation should be performed on the machine running the DAC server. It asks whether you want to continue.
  - b Click Yes.
  - c In the Repository Connection Information tab, enter the appropriate information for each instance. The Database Host should be the same for each instance, and the Database Port should be different.

- 5 For each instance, set up the DAC server system properties.
  - a Navigate to Setup > DAC System Properties.
  - b Set the DAC Server Host, OS, and Port properties.
- 6 Start each DAC server from its directory.

## Customizing Index and Analyze Table Syntaxes

The `customsql.xml` file, located in the `OracleBI\DAC\CustomSQLs` directory, contains the syntax for dropping and creating indexes and analyzing tables. You can edit the `customsql.xml` file to change the behavior of these operations.

### *To edit the Analyze Table syntax*

- 1 Open the `customsql.xml` file located in the `OracleBI\DAC\CustomSQLs` directory.
- 2 Locate the Analyze Table syntax for the appropriate database type.

For example, the syntax for an Oracle database is as follows:

```
<Sql Query name = "ORACLE_ANALYZE_TABLE" STORED_PROCEDURE = "TRUE">
DBMS_STATS.GATHER_TABLE_STATS(ownname => '@TABLEOWNER', tabname => '%1',
estimate_percent => 30, method_opt => 'FOR ALL COLUMNS SIZE AUTO', cascade => true )
</Sql Query>
```

- 3 Edit the syntax.

For example, to gather statistics for only the indexed columns, edit the syntax as follows:

```
<Sql Query name = "ORACLE_ANALYZE_TABLE" STORED_PROCEDURE = "TRUE">
DBMS_STATS.GATHER_TABLE_STATS(ownname => '@TABLEOWNER', tabname => '%1',
estimate_percent => 30, method_opt => 'FOR ALL INDEXED COLUMNS', cascade => true )
</Sql Query>
```

**NOTE:** The variables `@TABLEOWNER`, `%1`, `%2`, and so on, will be substituted appropriately by the DAC when the statement is executed.

### *To edit the Create Index syntax*

- 1 Open the `customsql.xml` file located in the `OracleBI\DAC\CustomSQLs` directory.
- 2 Locate the Create Index syntax for the appropriate database type, and edit the syntax.

## Using SQL Files as an Execution Type in the DAC

There are two types of custom SQL files that can be executed through the DAC: XML formatted .xml files, and plain text .sql files. For examples of XML and SQL files, go to the Oracle BI\DAC\CustomSQLs folder.

### XML Formatted Files

An XML file consists of a set of SQL statements for which the name, type, and Continue on Fail option are defined using XML attributes. The set of SQL statements is in a CDATA section which allows for special characters (like <, >, \) to be used without breaking the XML structure. A CDATA section looks like the following:

```
<![CDATA[This is a place for a SQL statement]]>
```

An example of an XML file follows:

```
<?xml version="1.0"?>
<CustomSQLs>
<sql name="Create Temp" type="SQL" continueOnFail="true">
<![CDATA[CREATE TABLE w_etl_temp (row_wid varchar(50))]]>
</sql >
<!-- This is how a comment section looks in XML -->
<!-- It will be ignored-->
<sql name="Update Temp" type="SQL">
<![CDATA[UPDATE w_etl_temp SET row_wid = 'qwe' WHERE row_wid LIKE '5*']]>
</sql >
<sql name="Drop Temp" type="SQL">
<![CDATA[DROP TABLE w_etl_temp]]>
</sql >
</CustomSQLs>
```

This example consists of three SQL statements: Create Temp, Update Temp, and Drop Temp. These names will be used in DAC run task detail description to give details of the errors encountered during execution (if any). In this particular example all three statements are of type SQL. Regular update SQL statements and the Stored Procedure type can also be used.

If the Continue on Fail option is set to True for a given statement, the execution will continue even if the task fails. If the Continue on Fail option is set to False for a given statement, or is omitted, the subsequent statements will not be executed and the Task Run status will be Failed.

## Plain Text SQL Files

Plain text SQL files consist of a set of SQL statements (no stored procedure calls). The SQL statements are separated by a semicolon (;), and comment tags are allowed (//, /\* comment \*/, --). If any of the SQL statements fail, the Task Run status will be Failed.

An example of a plain text SQL file follows:

```
CREATE TABLE w_etl_temp (name varchar(50))

;

UPDATE w_etl_temp

SET name = 'that's right' //this line demonstrates the use of ' in a text area

WHERE name LIKE 'gone fishing%';

/*

*some

*query

*statement

*/

SELECT * FROM w_etl_temp

;

DROP TABLE w_etl_temp

;

/*end of file*/
```

## Overview of Change Capture Process (Siebel Sources Only)

This section describes the change capture process used to extract data from the Siebel transactional database. It includes the following topics:

## Initial Data Capture

For each Siebel transactional source table (S\_) from which data is extracted, there is one S\_ETL\_I\_IMG\_ table and one S\_ETL\_R\_IMG\_ table.

The first time a staging table is extracted, rows that are in the specified period are selected and inserted into the appropriate S\_ETL\_R\_IMG\_ table. The specified period is set in the Prune Days parameter, in the Execution Plans tab. For more information about the Prune Days parameter, see [“About the DAC Execution Plans Tab” on page 115](#).

## Change Capture Mechanisms

There are two kinds of change capture mechanisms used:

- Change capture using tables

This is the most common method. It uses S\_ETL\_I\_IMG\_ and S\_ETL\_R\_IMG\_ table types and the LAST\_UPD column of the source tables.

- Change capture using the date column

In some cases, a predefined date column is used to enable change capture (without use of S\_ETL\_I\_IMG\_ and S\_ETL\_R\_IMG\_ tables).

### Change Capture Using Tables

When S\_ tables are extracted, the process looks for row ID information (the combination of ROW\_ID and MODIFICATION\_NUM from the S\_ table) that does not exist in the row image table (S\_ETL\_R\_IMG\_) and in columns where the value for LAST\_UPD is more recent than that of LAST\_REFRESH\_DATE minus the Prune Days parameter setting. This information is inserted into S\_ETL\_I\_IMG\_ table. The S\_ETL\_I\_IMG\_ table is joined with the base table during the SDE extraction process to extract only the change capture rows during refresh.

S\_ETL\_R\_IMG\_ tables store the ROW\_ID, MODIFICATION\_NUM, and LAST\_UPD for the rows in the S\_ table that have the LAST\_UPD in the defined Prune Days period. The LAST\_UPD column is used to delete the records from the S\_ETL\_R\_IMG\_ table. Records are deleted as soon as they go beyond the Prune Days period. This table is used to make sure that records that fall in the Prune Days period are not captured as updates unless they have actually been updated. This guarantees an efficient and faster change capture refresh. For information about tracking deleted records, see

Once the ETL process is completed, the data from the change capture image tables (S\_ETL\_I\_IMG\_) is pushed into the row image (S\_ETL\_R\_IMG\_) table. The S\_ETL\_R\_IMG\_ information is subsequently used in the next refresh.

Although the LAST\_UPD column in Siebel transactional tables is used for change capture, the timestamp reflects the time the data is committed in the database, rather than the actual transaction event time. This may happen because of remote synchronization, handheld synchronization, UTC conversion, and other processes in which there may be a significant lag between the time of a transaction and its commitment to the database. It is possible, therefore, for a data row to get committed to the transactional database with a LAST\_UPD date that is older than the date on which last refresh was executed. Consequently, if the extract is based purely on the date in LAST\_UPD, some rows might be missed during extraction.

The LAST\_UPD date column, however, still permits change capture process efficiency by limiting the number of rows that have to be compared. The rows from transactional tables are filtered based on the LAST\_UPD date being more recent than the LAST\_REFRESH\_DATE, minus the prune days. Then the ROW\_ID and MODIFICATION\_NUM combination is compared with the row image table to discover the changed records.

The Prune Days parameter ensures that the rows having LAST\_UPD values older than LAST\_REFRESH\_DATE are not missed. This is a parameter that customers can set based on experience with processes (such as remote synchronization) that may cause records to be missed.

### Primary and Auxiliary Tables

The DAC performs change capture for both primary and auxiliary tables. When more than one source table is involved, then both the auxiliary and primary table records need to be marked as changed. For auxiliary tables, you need to write auxiliary mappings to mark the primary tables as changed. The SQL queries that do this are part of the mapping SDEINC\_FindAux\_.

The extract logic sometimes requires that rows be considered as changed, even if the record has not been updated in the primary table (and therefore extracted during the SDE process). This situation occurs when child table rows have changed and the header/master rows need to be extracted so that data warehouse tables are loaded with a consistent set.

When the S\_CONTACT\_X row is changed, the corresponding S\_CONTACT also needs to be extracted. In this case, rows in S\_CONTACT are also marked as changed by inserting rows in the change capture row image table.

When the S\_ORDERITEM row is changed, the corresponding S\_DOC\_ORDER also needs to be extracted. In this case, rows in S\_DOC\_ORDER are also marked as changed (by inserting rows in the change capture row image table).

These auxiliary changes capture processes are heavily dependent on the data warehouse data model and are required to support the ETL logic.

### Example: Building S\_ETL\_I\_IMG\_ Table for Loading Account Dimension

This section gives an extended example of the process of change capture using tables.



1 Load image tables for all relevant source tables.

The content of this entity comes from the S\_ORG\_EXT and S\_ORG\_EXT\_X tables. Whenever any of the rows change in either of these tables, the record is marked as changed.

The image table for S\_ORG\_EXT is S\_ETL\_I\_IMG\_26. The image table prefix can be found using the DAC to view any source table. This table is truncated before loading with fresh data during every refresh.

During the ETL, process rows are inserted into S\_ETL\_I\_IMG\_26 by selecting ROW\_ID information from S\_ORG\_EXT, for rows (combined ROW\_ID and MODIFICATION\_NUM) that do not exist in the S\_ETL\_R\_IMG\_26 and for which LAST\_UPD is more recent than LAST\_REFRESH\_DATE minus the Prune Days setting. This is done during the ETL execution by the DAC's internal image building tasks.

Similarly, the image table S\_ETL\_I\_IMG\_27 for S\_ORG\_EXT\_X is loaded.

2 Load the image table for auxiliary table-based changes.

In addition to the basic change capture, extra processing might be required due to special ETL requirements. In this example, it happens that S\_ORG\_EXT needs to be extracted for processing even if only S\_ORG\_EXT\_X changes. This is because both the tables are joined to form W\_ORG\_D, and the extract process of W\_ORG\_D (a SDE mapping) looks for a changed ROW\_ID in the change capture row image table for the primary table S\_ORG\_EXT only. Therefore, the extract happens only when the ROW\_ID for S\_ORG\_EXT exists in the row image table.

In this case, the SDEINC\_FindAux\_ mapping is needed to insert corresponding rows of S\_ORG\_EXT.ROW\_ID in the change capture row image table whenever S\_ORG\_EXT\_X changes. The following logical statement shows the method:

Identify the records that have changed in the S\_ORG\_EXT\_X (rows in S\_ETL\_I\_IMG\_27) table and then find the corresponding rows in S\_ORG\_EXT. Insert the ROW\_ID and MODIFICATION\_NUM of those corresponding rows from S\_ORG\_EXT into S\_ETL\_I\_IMG\_26 table.

Using Informatica, the auxiliary mapping SDEINC\_FindAux\_ has to be written for each primary table that requires it, depending on data warehouse extract logic. Using the DAC, this auxiliary task has to be linked as a parent to the extract mapping for the base table (S\_ORG\_EXT in this case).

This is the SQL override for the SDEINC\_FindAux Informatica mapping:

```
SELECT
    S_ORG_EXT.ROW_ID,
    S_ORG_EXT.MODIFICATION_NUM,
    S_ORG_EXT.LAST_UPD
FROM
    S_ORG_EXT,
    S_ORG_EXT_X,
```

```
S_ETL_I_IMG_27 IMG
WHERE
(
  IMG.ROW_ID = S_ORG_EXT_X.ROW_ID
  AND
  S_ORG_EXT_X.PAR_ROW_ID = S_ORG_EXT.ROW_ID
)
AND NOT EXISTS
( SELECT 'X'
  FROM
  S_ETL_I_IMG_26 IMG1
  WHERE
  IMG1.ROW_ID = S_ORG_EXT.ROW_ID
)
```

**3** Extract source table information using change capture image information.

After the records that are new or modified are identified, those rows are loaded into the staging tables. The Informatica mappings that load the staging tables use the ROW\_ID information captured in the image tables.

This example shows the loading of staging table W\_ORG\_DS. The main logic of populating this table lies in the SQL override of the mapping SDE\_OrganizationDimension.

The DAC creates views on tables that are being extracted. The views are different, depending on whether a table is extracted the first time or is a change capture extract.

- If extracting for the first time, the view is created as `SELECT * FROM S_ORG_EXT`.
- If it is a change capture extract, the view is created as `SELECT * FROM S_ORG_EXT, S_ETL_I_IMG_26 IMG WHERE S_ORG_EXT.ROW_ID = IMG.ROW_ID`.

The SQL override in the mapping uses the view to extract the data.

```
SELECT
  S_ORG_EXT.ROW_ID,
  S_ORG_EXT.NAME, ....
...
...
FROM
```

```
V_ORG_EXT,  
S_ORG_EXT_X,  
.... .  
WHERE  
{  
  V_ORG_EXT S_ORG_EXT  
  LEFT OUTER JOIN S_ORG_EXT_X ON  
  S_ORG_EXT.ROW_ID = S_ORG_EXT_X.PAR_ROW_ID  
  .... .  
}  
AND  
  S_ORG_EXT.ROW_ID <> 'INT_COMPANY_ID'
```

## Change Capture Using the Date Column

Forecasts are extracted without using the image table. The value S\_FCSTSER\_DATE is tracked using the date column ARCHIVE\_TS. The administrator sets the ARCHIVE\_TS for forecasts that are submitted, frozen, and ready to be loaded into the Oracle Business Analytics Warehouse. S\_ETL\_RUN stores the previous ETL date when forecasts were extracted and the current ETL date when the forecasts are being extracted. All forecasts with ARCHIVE\_TS values greater than that of the previous ETL date and ARCHIVE\_TS (less than the current ETL date) are extracted in the current ETL. Both ETL date and ARCHIVE\_TS (less than the current ETL date) are stored in S\_ETL\_CURR\_RUN.

**NOTE:** Forecasts in the Oracle Business Analytics Warehouse are never updated. Once loaded, they are frozen.

```
SELECT  
.... .  
FROM  
  S_FCSTSER_DATE,  
  S_FCSTSER,  
  S_ETL_CURR_RUN,  
  ..... .  
WHERE  
  S_FCSTSER_DATE.FCSTSER_ID = S_FCSTSER.ROW_ID  
  AND S_FCSTSER_DATE.ARCHIVE_TS > S_ETL_CURR_RUN.PREV_LOAD_DT
```

```
AND S_FCSTSER_DATE.ARCHIVE_TS <= S_ETL_CURR_RUN.LOAD_DT
```

## Using the Change Capture Filter

The change capture filter allows you to selectively load data from the Siebel transactional database into the data warehouse. You can set a condition in the ChangeCaptureFilter.xml file to filter data from specific tables. This file is located in the OracleBI\DAC\CustomSQLs directory. It provides an XML sample that you can copy and alter to fit your needs. Detailed instructions for using this feature are included at the top of the file.

## Tracking Deleted Records

The Oracle Business Analytics Warehouse change capture process uses delete triggers to identify records for deletion on the Siebel transactional database. The deleted records are stored in S\_ETL\_D\_IMG tables. During the change capture process, the DAC server moves the data from the S\_ETL\_D\_IMG tables to the S\_ETL\_I\_IMG tables, where D appears in the OPERATION column to show the records were deleted. During the change capture sync process, the records in the S\_ETL\_D\_IMG tables that were moved to the S\_ETL\_I\_IMG tables are flushed. In the DAC, you can view the SQL that runs during the change capture and change capture sync processes by using the Output to File right-click command in the Tasks tab of the Design view.

The preconfigured ETL process captures deleted records for the target tables W\_ORG\_D and W\_PERSON\_D, the source tables for which are S\_ORG\_EXT, S\_CONTACT, and S\_PRSP\_CONTACT. These source tables need to have delete triggers created in the Siebel transactional database in order for deleted records to be tracked.

For vertical applications, the preconfigured ETL process captures deleted records for W\_FUND\_F and W\_ALIGNMT\_DH. You need to create delete triggers in the transactional database for the following additional tables: S\_MDF\_TXN, S\_ASGN\_GRP\_POSTN, S\_ASGN\_RULE\_ITEM.

In the Oracle Business Analytics Warehouse, preconfigured visibility tables are inactivated. If you activate visibility tables, you should also create delete triggers on the optional tables.

The preconfigured SIA Account and Contact visibility tables are activated by default for vertical applications. If your organization is not going to use any of the visibility tables, you need to inactivate them in the DAC.

On the target tables for which deleted records are tracked, a D appears in the INACTIVE\_FLG column to show the records as deleted when the source records are deleted. This method of flagging a record as deleted is known as a *soft delete*, as compared to a *hard delete* when the record is physically deleted. When deleted records are tracked on visibility-related data warehouse tables, the records are physically deleted. The general rule is that soft deletes should be used for tables that are referenced by other tables. If a table is not referenced by any other table, then you can use hard deletes.

Aggregate tables are rebuilt during each ETL process. Therefore, records can be physically deleted from the base tables without consequence. If you want to use the soft delete method, you should consider changing the aggregate building mappings so that the deleted records are omitted.

**NOTE:** The Oracle BI Server does not recognize soft deletes. Therefore, you have to modify the .rpd file so that it does not pick up soft-deleted records for reporting.

***To create delete triggers for preconfigured ETL change capture***

- 1 From the DAC menu bar, click Tools > ETL Management > Configure.
- 2 In the Sources dialog box, select the database platform for the target and transactional databases, and click OK.
- 3 In the Data Warehouse Configuration Wizard, select the Create Delete Triggers in Transaction Database check box, and click Next.

The Delete Triggers tab is active.

- 4 Select one of the following:

Option	Description
Create Triggers	Executes the trigger statements directly.
Write Script to File	Writes the trigger statements to a file, which can be executed by a database administrator.

- 5 Select the database type as defined in the DAC.
- 6 For DB2 zSeries databases, enter the base table owner.
- 7 (Optional) Select the Include Optional Triggers check box to create triggers for the optional tables.
- 8 Click Start.

***To create delete triggers for new source tables***

- 1 In the DAC, navigate to Design > Tables.
- 2 Select the table for which you want to track deleted records.  
Make sure the table has an image suffix.
- 3 Right-click the table and select Change Capture Scripts > Generate Image and Trigger Scripts.
- 4 In the Triggers and Image Tables dialog box, select the database type of the source database.
- 5 Make sure the Generate Image Table Scripts and Generate Trigger Script(s) options are selected.
- 6 Execute the script on the database.

### *To track deleted records*

- 1 Make sure the delete triggers are enabled for the appropriate tables.
- 2 Write custom Informatica workflows with a clause WHERE operation = 'D' to the appropriate I\_IMG table to take them across to the dimension and fact tables.
- 3 In the DAC, register the workflows as tasks.
- 4 Define the appropriate dependencies.

For an example of such a workflow, see the preconfigured task SDE\_OrganizationDimension\_LoadDeletedRows.

## Handling ETL Failures with the DAC

This section includes the following topics:

- [When the Execution of an Execution Plan Fails on page 86](#)
- [Discarding the Current Run Execution Plan on page 87](#)
- [Failure of Aggregator Transformation Tasks with Sorted Input on page 87](#)

### When the Execution of an Execution Plan Fails

When an execution plan is executed, if a task fails, the status of the tasks that are dependent on the failed task is changed to Stopped. While tasks are still running, the execution plan's status is Running. When all the tasks have been run, and if one or more tasks have failed, the execution plan's status is changed to Failed.

You can check the tasks that have failed in the Current Run tab of the Execute view, fix the problems, and then requeue the failed tasks by changing the status to Queued. You can then restart the ETL. All the tasks will then be rerun. You can also manually run a task, change its status to Completed, and then restart the ETL. Tasks with a Completed status are skipped.

**CAUTION:** The DAC server does not validate tasks that have been run manually.

To restart a Failed ETL, click Run Now from the Current Run tab of the Execute view.

### In Case of Abnormal Termination of the DAC Server

If the DAC server fails during the execution of the ETL, the status of the ETL execution will remain as Running. When the DAC server is started again, it will automatically run the ETL if the Auto Restart ETL DAC system property is set to True. If the same system property is set to False, when the server restarts, it will set the correct status as Failed. In order to execute the ETL from the point of failure, submit the request to the server again.

The DAC server will automatically terminate if it loses connection to the DAC repository.

## Discarding the Current Run Execution Plan

You can discard an execution plan that failed by navigating to the Current Run tab, right-clicking on the execution plan and changing its status to Mark as Completed. This will force the run status to be updated as Completed. When you submit a request for another run, the DAC server creates another instance of it.

**CAUTION:** Perform this procedure in a development or testing environment only, since it might leave the data in an inconsistent state, causing you to have to reload all of the data.

## Failure of Aggregator Transformation Tasks with Sorted Input

Tasks that use Informatica Aggregator transformation can fail when the Sorted Input option is active. The tasks SDE\_DTLFORECASTFACT and SDE\_COSTLIST are examples of tasks that can fail in such a situation.

To prevent such tasks from failing, in Informatica Designer, navigate to Mapping Designer, open the corresponding mapping, and in the Aggregator transformation, remove the check from the Sorted Input check box.





# 7

## DAC Functional Reference

This chapter describes the functionality available in the Data Warehouse Administration Console (DAC). It contains the following topics:

- [Common Elements of DAC Interface Tabs on page 89](#)
- [About the DAC Design View on page 89](#)
- [About the DAC Setup View on page 106](#)
- [About the DAC Execute View on page 114](#)

### Common Elements of DAC Interface Tabs

Some of the DAC interface tabs have common elements, such as columns or child tabs. [Table 13](#) provides a description of these common elements.

Table 13. Common Elements of DAC Interface Tabs

Column or Child Tab	Description
Name	A column that specifies the name of the repository, warehouse, or transactional database object.
Inactive	A column that indicates whether an object is inactive. Inactive objects do not participate in the ETL process.
Owner	A column that specifies the source system container in which the object was created.
Edit	A child tab that allows you to edit an object that is selected in the top pane window.
Description	A child tab that displays and allows you to edit a description of the object selected in the top pane.

### About the DAC Design View

The DAC Design view provides access to functionality related to creating and managing subject areas.

This section includes the following topics:

- [About the DAC Subject Areas Tab on page 90](#)
- [About the DAC Tables Tab on page 92](#)
- [About the DAC Indices Tab on page 94](#)

- [About Advanced Custom Index Management on page 96](#)
- [About the DAC Tasks Tab on page 97](#)
- [About the DAC Tasks Tab on page 97](#)
- [About the DAC Configuration Tags Tab on page 105](#)
- [About the DAC Source System Parameters Tab on page 106](#)
- [About the DAC Source System Folders Tab on page 106](#)
- [About the DAC System Properties Tab on page 107](#)

## About the DAC Subject Areas Tab

The Subject Areas tab lists all the subject areas associated with the selected source system container. It allows you to view and edit existing subjects areas and to create new ones. [Table 14](#) provides a description of the properties in the top pane of the tab.

Table 14. DAC Subject Areas Tab Top Pane Properties

Column	Description
Configuration Tag tasks only	Indicates whether configuration tag tasks are the only tasks associated with this subject area that will participate in the ETL process. If this check box is selected, only the tasks associated with the configuration tag will be chosen by the DAC when the subject area is assembled.  For more information about Configuration Tag tasks, see <a href="#">“About the DAC Configuration Tags Tab” on page 105</a> .
Last Designed	Date and time the subject area was last assembled.

Table 15 provides a description of the child tabs in the bottom pane of the Subject Areas tab.

Table 15. DAC Subject Areas Child Tabs

Child Tab	Description
Tasks	<p>Displays the tasks associated with the selected subject area, and allows you to add tasks to a subject area, inactivate tasks, and remove tasks from a subject area.</p> <p><b>NOTE:</b> When you <i>inactivate</i> a task, it remains inactive even if you reassemble the subject area. When you <i>remove</i> a task from a subject area, it will be added back to the subject area upon reassembly.</p> <p>It includes the following properties:</p> <ul style="list-style-type: none"> <li>■ <b>Parent Group.</b> If the task belongs to a task group, this column displays the task group name.</li> <li>■ <b>Phase.</b> Task phase of the ETL process.</li> <li>■ <b>Autogenerated.</b> Indicates whether the task was automatically generated by the DAC's task generation process.</li> <li>■ <b>Is Group.</b> Indicates whether the task is a task group.</li> </ul>
Tables	<p>Displays the tables that are associated with the selected subject area. It allows you to add tables to subject areas or to remove them.</p>
Configuration Tags	<p>Displays the configuration tags that are associated with this subject area.</p> <p>It includes the following properties:</p> <ul style="list-style-type: none"> <li>■ <b>Include Tasks.</b> This field is read-only. It indicates whether the configuration tag tasks will be executed with the selected subject area.</li> <li>■ <b>Context Disabled.</b> If this check box is selected, the configuration tag is globally disabled (set as Inactive in the Configuration Tags parent tab).</li> </ul>

## About the DAC Tables Tab

The Tables tab lists all the tables associated with the selected source system container. It allows you to view and edit existing tables and to create new ones. [Table 16](#) provides a description of the table properties in the top pane of the tab.

Table 16. DAC Tables Tab Top Pane Properties

Column	Description
Table Type	Indicates the type of table.
Warehouse	Indicates whether the table is a warehouse table.  <b>NOTE:</b> If the warehouse flag is not selected, the schema creation process will not include this table.
Image Suffix	Suffix for image tables. Applicable to Siebel source tables only. For more information about image tables, see <a href="#">“Generate Image and Trigger Scripts Command” on page 93</a> .
Is MultiSet	Indicates whether the table is a MultiSet table. Applicable only to Teradata databases.
Has Unique Primary Index	Indicates whether the table has a Unique Primary Index. Applicable only to Teradata databases.

[Table 17](#) provides a description of the child tabs in the bottom pane of the Tables tab.

Table 17. DAC Tables Child Tabs

Tab	Description
Related Tables	Lists tables that are related to the selected table. Related tables participate in the ETL process in addition to the tables that are associated with this table.
Columns	Lists the columns associated with the selected table and allows you to add columns to the selected table. Also allows you to enter foreign key table and foreign key column relationships.
Multi-column Statistics	Applicable to Teradata databases only.
Indices (RO)	Displays a read-only list of indices that belong to the selected table.
Source for Tasks (RO)	Displays a read-only list of tasks that use the selected table as a source.
Target for Tasks (RO)	Displays a read-only list of tasks that use the selected table as a target.
Conditional for Tasks (RO)	Displays a read-only list of tasks that are optional tasks for the selected table.

## Actions Available

The following actions are available in the top pane toolbar and in the right-click menu when the Tables tab is active.

### Generate Index Scripts Command

This command generates drop index, create index, and analyze table scripts for all the tables that participate in the ETL process. You can query for one or more tables, and the command will generate all the scripts for all the tables. The results are stored in the log\scripts directory. Files are generated that provide the following information: drop indices of type ETL, create indices of type ETL, drop indices of type Query, create indices of type Query, and analyze tables.

### Generate Image and Trigger Scripts Command

This command generates change capture scripts for tables with defined image suffixes (Siebel sources only). The scripts include creation of the image tables to assist in the change capture process and their indices. You can also create the triggers for deletes. Triggers for deletes should be used only for new source tables for which you want to use change capture. The scripts will be generated for the selected table or all the tables shown in the list applet. You can narrow down the tables of interest by querying. The results are shown in a text box, which can be supplied to the DBA for executing.

### Generate View Scripts

This command generates change capture view scripts for full or incremental loads for any table that participates in the change capture process. This command can be useful when the change capture fails and you want to run a task manually. It can also be used for unit testing.

### Generate Change Capture SQL

This command generates change capture SQL scripts for full or incremental mode for tables that participate in the change capture process. This command can be useful when the change capture fails and you want to run a task manually. It can also be used for unit testing.

### Import Database Tables Command

This command allows you to import table definitions from a selected database. This command does not allow you to import table column definitions. You need to use the Import Database Columns command to do so.

### Import Database Columns Command

This command allows you to import table column definitions.

### Import Indices Command

This command allows you to import index definitions from the database for one or more tables as listed in the result of the query.

**NOTE:** In MSSQL Server databases, when you import indices with columns whose sort order type is "Desc," the sort order type appears in the DAC as "Asc." You have to set the sort order manually to "Desc."

## About the DAC Indices Tab

The Indices tab displays a list of all the indices associated with the selected source system container. It is recommended that you do not register any indices for source tables. During the ETL process, when a table is going to be truncated, all the indices as defined in the repository will be dropped before the data is loaded and will be created after the data is loaded automatically. While this improves the ETL performance, the preconfigured workflows have the bulk load option turned on. The bulk load will fail if there are indices on the table. Therefore, it is important to keep the index definitions in sync with the database. For example, if you create an index on the database, and it is not registered in the repository, the index will not be dropped and the load will fail.

For Teradata databases, only secondary indices should be registered in the DAC. You should not register primary indices or the more complex indices, such as single- and multi-table indices, because they cannot be dropped and recreated. You can use SQL commands to drop and create such tasks in the DAC.

Table 18 provides a description of the index properties displayed in the top pane of the Indices tab.

Table 18. DAC Indices Tab Top Pane Properties

Column	Description
Table Name	Table for which an index is created.
Index Usage	Usage of index: ETL or Query. An ETL index is typically used during the ETL process. A Query index is an index used only during the reporting process. It is recommended that you have a clear understanding of when and where the index will be used at the time of registering the index.
Databases	Allows you to associate an index with a database type.
# Unique Columns	For unique indices, the number of columns that will be unique.
Is Unique	Indicates whether the index is unique.
Is Clustered	Indicates whether the index is clustered. There can be only one clustered index per table.
Is Bitmap	Indicates whether the index is of the bitmap type.

Table 18. DAC Indices Tab Top Pane Properties

Column	Description
Allow Reverse Scan	Applicable only for DB2-UDB databases. The index will be created with the Allow Reverse Scan option.
Always Drop & Create	Indicates whether the index will be dropped and created regardless of whether the table is being loaded using a full load or incremental load.

Table 19 shows which index properties are available for the different database types that are supported.

Table 19. Index Properties and Database Types Matrix

Column/ Database	Is Unique	Is Clustered	Is Bitmap	Allow Reverse Scan	# Unique Columns
Oracle	X		X		
DB2	X	X		X	X
DB2-390	X				
MSSQL	X	X			

Table 20 provides a description of the child tabs in the bottom pane of the Indices tab.

Table 20. Indices Child Tabs

Tab	Description
Columns	<p>Displays the list of columns the index is made of.</p> <p>It includes the following properties:</p> <ul style="list-style-type: none"> <li>■ <b>Position.</b> Position of the column in the index.</li> <li>■ <b>Sort Order.</b> Indicates whether the sort order is ascending or descending.</li> </ul>
Databases	Lists the database types that apply to the selected index. If no database type is indicated, the index will not be created.

## Actions Available

The following actions are available in the top pane toolbar and in the right-click menu when the Indices tab is active.

### Database Filter

The database filter appears in the Indices tab to the right of the top pane toolbar, as shown in Figure 11. Click on the words “Database Type” to open the Index Filtering dialog box. It allows you to filter the set of displayed indices based on the database type. To display all indices regardless of database type, select the option All.

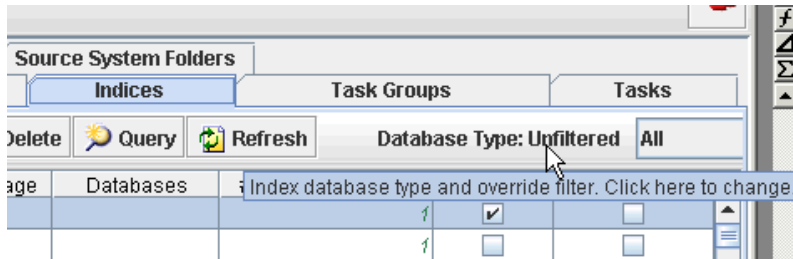


Figure 11. Indices Tab Database Filter

## About Advanced Custom Index Management

The DAC allows you to drop and recreate indices during the load process, which reduces the overall load time during a full load. The DAC drops and recreates indices based on the index definitions stored in the DAC metadata.

**NOTE:** More advanced index management needs to be handled outside of the DAC, such as creating partitioned indices on Oracle databases, which requires specific syntax, and creating single-table and multi-table join indices on Teradata databases. In such cases you can use DAC SQL tasks placed appropriately in the task dependencies, or you can use a pre-session or post-session script in Informatica. You must drop the partitioned index before a task runs to load data using the bulk loader, because the bulk loader may fail if there are indices on the target database.

If you need to modify the preconfigured indices with extra options or syntax, you must inactivate them in the DAC metadata so that the DAC server does not try to drop and recreate them. You can then manage these indices in the same manner as the advanced index management described above.



## About the DAC Tasks Tab

The Tasks tab lists all the tasks associated with the selected source system container. [Table 21](#) provides a description of the task properties displayed in the list of the top pane of the Tasks tab.

Table 21. Tasks Tab Top Pane Properties

Column	Description
Parent Group	If the task is a member of a group, this field lists the task group name.
Command for Incremental Load	<p>A table can be loaded in Full Mode or Incremental Mode. Full Mode refers to data loaded for the first time or data that is truncated and then loaded. Incremental Mode refers to new or changed data being added to the existing data.</p> <p>The DAC maintains a last refresh timestamp whenever a table is changed during the ETL process. (You can view this timestamp under Setup &gt; Physical Data Sources &gt; Refresh Dates.) If a table has a timestamp, the command appearing in this column is executed. If a table does not have a timestamp, the command for a full load is executed. If the execution type is Informatica, the workflow name is used as the command.</p>
Command for Full Load	If a table has no last refresh timestamp, this command is executed.
Folder Name	<p>Only for execution type of Informatica. The folder in which the workflow resides.</p> <p><b>NOTE:</b> The name cannot contain spaces.</p>
Primary Source	Logical database connection for the primary source database.
Primary Target	Logical database connection for the primary target database.
Task Phase	Task phase of the ETL process. This information is primarily used for dependency generation. Certain phases, such as Change Capture and Query Index Creation, are not available for you to assign to a task. The DAC server uses the task phase to prioritize tasks and to generate a summary of the time taken for each of the phases.

Table 21. Tasks Tab Top Pane Properties

Column	Description
Execution Type	<p>Tasks are executed based on their execution type. The following types are supported:</p> <ul style="list-style-type: none"> <li>■ <b>Informatica.</b> Task is invoked on an Informatica Server using pmcmd.</li> <li>■ <b>External Program.</b> Task is an operable program on the operating system where the DAC server is running. This program can be a batch file, shell script, or any other program that can be run like a bulk loader.</li> <li>■ <b>SQL File.</b> Task is a SQL script in .xml or .sql format.</li> <li>■ <b>Stored Procedures.</b> Task is a stored procedure that has been defined on the databases.</li> </ul> <p>In addition, there are several internal execution types that you will not be able to select when creating new tasks. These tasks are categorized as either internal change capture tasks or internal data warehouse tasks; all of these tasks are color-coded in pink in the Tasks tab.</p> <ul style="list-style-type: none"> <li>■ <b>IMG_BUILD</b> - internal change capture. If you are using multiple Siebel transactional sources, you cannot change the behavior of the change capture process. This task requires change capture tables to be created on the other sources also. When adding additional Siebel sources, go to Design &gt; Tables, right-click and select Change Capture Tasks. This action generates change capture tasks. Use this same action to disable or delete change capture tasks.</li> <li>■ <b>IMG_SYNC</b> - internal change capture. If you are using multiple Siebel transactional sources, you can create this task for the additional tasks for doing similar change capture sync processes. You cannot change the behavior of the change capture sync process. This task requires change capture tables to be created on the other sources also. This task should be used with discretion for Siebel sources only.</li> <li>■ <b>QUERY_INDEX</b> - internal data warehouse. This task allows you to alter when the Query indices are created. The DAC server drops all indices before loading when the CreateQueryIndexesAtTheEnd setting is set to True. When this setting is set to False, all the indices, regardless of the index type, get created as part of the task that does the loading.</li> <li>■ <b>UPDATE_ETL_PARAM</b> - internal data warehouse. This task is used only to update W_PARAM_G from the DAC server. This task populates the system properties to the W_PARAM_G table in the data warehouse by querying values defined in the DAC repository. Because only one data warehouse per DAC repository is supported, this execution type should not be chosen for any task.</li> </ul>
Priority	<p>Indicates the order in which the task runs. If two or more tasks have the same priority, the order occurs randomly.</p>

Table 21. Tasks Tab Top Pane Properties

Column	Description
Pre-SQL for Full Load	The SQL script (derived from a SQL or XML file) that is executed before the specified task when the task is participating in a full load.
Pre-SQL for Incremental Load	The SQL script (derived from a SQL or XML file) that is executed before the specified task when the task is participating in an incremental load.
Post-SQL for Full Load	The SQL script (derived from a SQL or XML file) that is executed after the specified task when the specified task is participating in a full load.
Post-SQL for Incremental Load	The SQL script (derived from a SQL or XML file) that is executed after the specified task when the specified task is participating in an incremental load.
Build Image	Applicable for Siebel transactional sources only. Indicates the change capture for the primary/auxiliary source tables will be executed.
Analyze Tables	The DAC automatically analyzes tables when tasks truncate tables. By selecting this check box, however, you can force the DAC to analyze tables even when they are not truncated.
Continue on Error	When this check box is selected, if the command fails, the dependent tasks are not stopped. However, if any autogenerated tasks fail, the dependent tasks are stopped.

Table 22 provides a description of the child tabs in the bottom pane of the tab.

Table 22. Tasks Child Tabs

Tab	Description
Source Tables	<p>Displays the list of tables from which the selected task extracts data.</p> <p>It includes the following properties:</p> <ul style="list-style-type: none"> <li>■ <b>Type.</b> Table type. Possible values are:                             <ul style="list-style-type: none"> <li>■ <b>Primary.</b> Indicates whether the table is a primary source of data.</li> <li>■ <b>Auxiliary.</b> Indicates whether the table is a secondary source of data.</li> <li>■ <b>Lookup.</b> Indicates whether the table is a lookup table.</li> </ul> </li> <li>■ <b>Data Source.</b> When a data source is not specified, the default is the task's primary source.</li> </ul> <p><b>NOTE:</b> If a table is marked as Primary or Auxiliary and the Build Image property of the task is selected, the change capture process is invoked. There are special tasks that force the base table data to be extracted when data in auxiliary tables change.</p> <p>A table can be neither Primary nor Auxiliary but still be used for getting some attributes to populate a dimension or fact table. The changes in these kinds of source tables are not reflected in the dimension or fact table once the data is populated.</p>
Target Tables	<p>Displays the list of tables into which the selected task loads data.</p> <p>It includes the following properties:</p> <ul style="list-style-type: none"> <li>■ <b>Data Source.</b> If no data source is specified, this value defaults to the task's primary target.</li> <li>■ <b>Truncate Always.</b> Indicates the target tables will be truncated regardless of whether a full or incremental load is occurring. Any indices registered for this table are dropped before the command is executed and are recreated after the command completes successfully. When indices are dropped and created, the table is analyzed so that the index statistics are up-to-date.</li> <li>■ <b>Truncate for Full Load.</b> Indicates the target tables will be truncated only when a full load is occurring. Any indices registered for this table are dropped before the command is executed and are recreated after the command completes successfully. When indices are dropped and created, the table is analyzed so that the index statistics are up-to-date. When the Truncate Always option is selected, this option is unnecessary.</li> </ul>

Table 22. Tasks Child Tabs

Tab	Description
Conditional Tables	<p>Displays the tables that, if included in an execution plan, cause the optional task selected in the top pane to be executed.</p> <p>For example, the Order Item fact table is a conditional table associated with the optional task called UpdateRecencyCat in Person Dimension. The UpdateRecencyCat in Person Dimension task is executed only when the Order Item fact table is included in an execution plan.</p>
Phase Dependency	<p>The DAC server uses the ETL phase property to prioritize tasks. By changing the phase property of a task, you change the task's execution order.</p> <p>For instructions on setting a phase dependency, see <a href="#">"Setting a Task Phase Dependency" on page 48</a>.</p> <p>This child tab includes the following properties.</p> <ul style="list-style-type: none"> <li>■ <b>Action.</b> The action to be taken in relation to the phase dependency. Possible values are: <ul style="list-style-type: none"> <li>■ <b>Wait.</b> Indicates the selected task will wait to be executed until the tasks of a specified phase have been executed.</li> <li>■ <b>Block.</b> Indicates the selected task will block all tasks of the specified phase from being executed until is has been executed.</li> </ul> </li> <li>■ <b>Grain.</b> Possible values are: <ul style="list-style-type: none"> <li>■ <b>All.</b> Indicates the action will affect <i>all</i> tasks.</li> <li>■ <b>Related.</b> Indicates the action will affect only <i>related</i> tasks. You can view a task's related tasks by navigating to Execution Plans &gt; All Dependencies and viewing the specified task's predecessor tasks.</li> </ul> </li> <li>■ <b>Phase.</b> The ETL phase that will apply to the Action and Grain properties.</li> </ul>

Table 22. Tasks Child Tabs

Tab	Description
Configuration Tags	<p>Displays the configuration tags to which the selected task belongs. It also allows you to associate the selected task with a configuration tag.</p> <p>It includes the following properties:</p> <ul style="list-style-type: none"> <li>■ <b>Include Tasks.</b> This field is read-only. It indicates whether the configuration tag tasks will be executed with the selected subject area.</li> <li>■ <b>Context Disabled.</b> If this check box is selected, the configuration tag is globally disabled.</li> </ul>
Parameters	<p>Displays the parameters associated with the selected task. It allows you to configure task level parameters. This parameter takes precedence over source system parameters when the name is the same.</p> <ul style="list-style-type: none"> <li>■ <b>Data Type.</b> Possible values are Text, Timestamp, and SQL.</li> <li>■ <b>Value.</b> The value of the parameter.</li> </ul>

## About the DAC Task Groups Tab

The Task Groups tab lists all the task groups associated with the selected source system container. A task can belong to only one group. [Table 23](#) provides a description of the task group properties displayed in the top pane of the tab.

Table 23. Task Groups Tab Top Pane Properties

Column	Description
Restart All on Failure	Indicates the tasks in this task group will be restarted if one or more tasks fails during an ETL process.
Execute Serially	Indicates the tasks in this task group will be executed sequentially. This property overrides the execution order.

Table 23. Task Groups Tab Top Pane Properties

Column	Description
Truncate Always	<p>Indicates the target tables are truncated regardless of whether a full or incremental load is occurring. Any indices registered for the target table are dropped before the command is executed and are recreated after the command completes successfully. When indices are dropped and created, the table is analyzed so that the index statistics are up-to-date.</p> <p><b>NOTE:</b> Make sure if you select this option that all the tasks write to the same data source.</p>
Truncate for Full Load	<p>Indicates the target tables will be truncated only when a full load is occurring. Any indices registered for the target table are dropped before the command is executed and are recreated after the command completes successfully. When indices are dropped and created, the table is analyzed so that the index statistics are up-to-date.</p>

Table 24 provides a description of the child tabs in the bottom pane of the Task Groups tab.

Table 24. Task Groups Child Tabs

Column	Description
Child Tasks	<p>Lists the tasks that belong to the selected task group.</p> <p>It includes the following properties:</p> <ul style="list-style-type: none"> <li>■ <b>Execution Order.</b> Order among the tasks in the task group in which this task will be executed. If two or more tasks have the same execution order and the Execute Serially flag is not checked, the DAC will run the tasks in parallel.</li> </ul>
Source Tables	<p>Read only. Displays the tables used for getting data by the selected task group.</p> <p>It includes the following properties:</p> <ul style="list-style-type: none"> <li>■ <b>Task.</b> Task that extracts data from the table.</li> <li>■ <b>Primary.</b> Indicates whether the table is a primary source of data.</li> <li>■ <b>Auxiliary.</b> Indicates whether the table is a secondary source of data.</li> <li>■ <b>Type.</b> Source table type.</li> </ul> <p><b>NOTE:</b> If a table is marked as Primary or Auxiliary and the Build Image property of the task is selected, the change capture process is invoked. There are special tasks that force the base table data to be extracted when data in auxiliary tables changes.</p> <p>A table can be neither Primary nor Auxiliary but still be used for getting some attributes to populate a dimension or fact table. The changes in these kinds of source tables are not reflected in the dimension or fact table once the data is populated.</p>
Target Tables	<p>Read only. Displays the tables into which the task group (selected in the top pane) loads data.</p> <p>It includes the following properties:</p> <ul style="list-style-type: none"> <li>■ <b>Table.</b> Name of the target table.</li> <li>■ <b>Task.</b> Task that loads data into the target table.</li> <li>■ <b>Primary.</b> Indicates whether the table is a primary target table. For documentation purposes only.</li> <li>■ <b>Auxiliary.</b> Indicate whether the table is a secondary target table.</li> <li>■ <b>Type.</b> Target table type.</li> </ul>



## About the DAC Configuration Tags Tab

A configuration tag is an object that controls the inclusion of tasks in subject areas. When a task is tagged, it is not eligible to be included in the collection of tasks for any subject area, unless the tag is part of the subject area definition "Include Task" property.

A configuration tag can function in one of the following ways:

- **Remove tasks from all subject areas**
- **Reassign autogenerated tasks to a specific subject area**
- **Add non-autogenerated tasks to a subject area**

For instructions on creating configuration tags, see ["Working with Configuration Tags" on page 51](#).

[Table 25](#) provides a description of the columns in the top pane of the Configuration Tags tab.

Table 25. Configuration Tags Tab Top Pane Properties

Column	Description
Include Tasks	<p>If this check box is selected, the tasks that are assigned to a configuration tag will participate in the ETL process for the subject area to which this configuration tag is assigned.</p> <p>For example, suppose Configuration Tag 1 is made up of Task 1 and Task 2, and Configuration Tag 1 is assigned to Subject Area 1. Task 1 and Task 2 will be executed when the execution plan for Subject Area 1 is executed, whether or not Task 1 and Task 2 relate to the tables that make up the subject area.</p>

[Table 26](#) provides a description of the child tabs in the bottom pane of the Configuration Tags tab.

Table 26. Configuration Tags Child Tabs

Child Tab	Description
Subject Areas	<p>This child tab allows you to assign a configuration tag to a subject area.</p> <ul style="list-style-type: none"> <li>■ <b>Configuration Tag tasks only.</b> This field is read-only. It indicates whether configuration tag tasks are the only tasks associated with this subject area that will participate in the ETL process. If this check box is selected, the tasks shown in the Tasks child tab will not participate in the ETL process; only the tasks associated with the configuration tag will participate. You set this flag in the Subject Areas parent tab.</li> </ul>
Tasks	This child tab allows you to assign tasks to a configuration tag.

## About the DAC Source System Parameters Tab

The Source Systems Parameters tab lists all the source system parameters associated with the selected source system container. It allows you to view and edit existing parameters and to configure new ones. [Table 27](#) provides a description of the properties in the top pane of the tab.

Table 27. Source Systems Parameters Tab Top Pane Properties

Column	Description
Data Type	The parameter data type. Possible values are: <ul style="list-style-type: none"> <li>■ Text</li> <li>■ Timestamp</li> <li>■ SQL</li> </ul>
Value	The parameter value.

## About the DAC Source System Folders Tab

The Source System Folders tab lists the Informatica folders associated with the selected source system container. It allows you to view existing folders and to create new ones. [Table 28](#) provides a description of the tab properties.

Table 28. Source Systems Parameters Top Pane Properties

Column	Description
Logical Folder	Name of the logical Informatica folder. This name is used in the task definition (in the Tasks tab) so that task definitions do not have to be cloned.
Physical Folder	Name of the physical Informatica folder. The physical Informatica folder corresponds to the actual folder in the Informatica repository. This name is used in the Ordered Tasks child tab of the Execution Plans tab.

## About the DAC Setup View

The Setup View provides access to functionality related to setting up DAC system properties, Informatica servers, database connections, and email notification.

This section includes the following topics:

- [About the DAC System Properties Tab on page 107](#)
- [About the Informatica Servers Tab on page 111](#)
- [About the Physical Data Sources Tab on page 112](#)
- [About the Email Recipients Tab on page 114](#)

## About the DAC System Properties Tab

The DAC System Properties tab allows you to configure various properties that determine the behavior of the DAC server. [Table 29](#) provides a description of the DAC system properties.

Table 29. DAC System Properties Tab Properties

Property	Description
Analyze Frequency (in days)	For DAC metadata tables, the frequency (in days) the DAC client automatically updates the table and index statistics for the DAC repository. The value must be numerical.
Auto Restart ETL	<p>Possible values are True and False.</p> <p>When set to True: An ETL that is running when the DAC server abnormally terminates will continue running when the DAC server is restarted.</p> <p>When set to False: An ETL that is running when the DAC server abnormally terminates will not automatically restart when the DAC server restarts. The ETL status will be updated to Failed. An administrator will have to manually restart the ETL.</p>
CreateQueryIndexesAtTheEnd	<p>Possible values are True and False.</p> <p>During the ETL process, the DAC server automatically drops and creates indices. When set to True, this property groups all indices of the Query type and creates them after the ETL is complete.</p> <p><b>NOTE:</b> The DropAndCreateIndexes property takes precedence over this property. Therefore, if the DropAndCreateIndexes property is set to False, you cannot set the property CreateQueryIndexesAtTheEnd to True to have indices of the Query type created at the end of the ETL process.</p> <p>Also, be aware that when this property is set to True, tables will be analyzed twice. If any indices are marked as Query type indices, and are used by ETL processes, it can adversely affect the performance of the ETL process.</p>
DAC Alternate Server	Host name of the machine where the alternate DAC server resides. The alternate DAC server is used for failover purposes. The DAC client cannot talk to the alternate server unless the main DAC server is not running.

Table 29. DAC System Properties Tab Properties

Property	Description
DAC Server Host	<p>Host name of the machine where the DAC server resides. You cannot use an IP address for this property.</p> <p><b>NOTE:</b> The DAC server and a given DAC repository have a one-to-one mapping. That is, you can only run one DAC server against any given DAC repository. Thus, in the repository you must specify the network host name of the machine where the DAC sever is to be run.</p> <p><b>NOTE:</b> This property also takes the value <i>localhost</i>. However, this value is provided for development and testing purposes and should not be used in a production environment.</p>
DAC Server OS	<p>Operating system of the machine where the DAC server resides. Possible values are Windows, Solaris, HP, or AIX.</p> <p><b>NOTE:</b> If you move the DAC server from another operating system to AIX, you need to do the following: change the DAC server host to the appropriate value; restart the DAC client; reenter all the password fields for the Informatica servers and database connections; and reconfigure the DAC server on the AIX machine by running <code>serverSetupPrompt.sh</code>.</p>
DAC Server Port	<p>Network port to which the DAC server binds in order to listen to client requests. The default value is 3141. If this port has been assigned to another process, you can enter any numerical port value greater than 1024.</p>
Drop and Create Change Capture Views Always	<p>Possible values are True and False.</p> <p>When set to True (the default value), the DAC server drops and creates change capture views every time it performs a change capture process, including for both full and incremental loads.</p> <p>Setting this property to True can create system catalog lock up for DB2-UDB and DB2-390 databases. Therefore, by setting the property to False, the DAC server will drop and create views selectively, using the following rules:</p> <ul style="list-style-type: none"> <li>■ In full mode: <ul style="list-style-type: none"> <li>■ During the change capture phase, views will be dropped and created as full views.</li> <li>■ During the change capture sync process, incremental views will be generated.</li> </ul> </li> <li>■ In incremental mode: <ul style="list-style-type: none"> <li>■ If the view exists, it will not be dropped and created.</li> <li>■ If the view does not exist, the incremental view will be created.</li> </ul> </li> </ul>

Table 29. DAC System Properties Tab Properties

Property	Description
DropAndCreateIndexes	<p>Possible values are True and False.</p> <p>Indicates whether, during the ETL process, the DAC server automatically drops and creates indices.</p> <p><b>NOTE:</b> This property takes precedence over the CreateQueryIndexesAtTheEnd. Therefore, if the DropAndCreateIndexes property is set to False, you cannot set the property CreateQueryIndexesAtTheEnd to True to have indices of the Query type created at the end of the ETL process.</p>
Dryrun	<p>Possible values are True and False.</p> <p>Indicates whether tasks are executed without invoking Informatica workflows. The following processes are executed: change capture, truncation of tables, drop and creation of indices, and analyze statements.</p> <p><b>NOTE:</b> This option should be used for debugging purposes only and not used in a production environment.</p>
Generic Task Concurrency Limit	<p>Determines how many tasks with execution types other than Informatica can be run concurrently. The value must be numerical.</p> <p>To set this value, you should consider what the external tasks do. For example, if the tasks open connections to a database, you should consider how this would affect the preconfigured tasks.</p>
HeartBeatInterval	<p>Frequency (in seconds) the DAC server checks on the health of the database connections. The value must be numerical. For example, a value of 300 (the default value) indicates the system will perform subsystem diagnostics and recovery procedures every 300 seconds.</p>
InformaticaFileParameter Location	<p>Directory where the Informatica parameter file is stored.</p>
Main Informatica Repository	<p>This property is obsolete from Release 7.9 onwards.</p>

Table 29. DAC System Properties Tab Properties

Property	Description
Output Redirect	<p>Indicates whether logging information and standard output and errors are redirected to files in the log directory (when property is set to True). The file containing standard output starts with out_ and ends with the .log extension. The standard error messages are in the file starting with err_ and ending with the .log extension.</p> <p>If this property is set to False, the logging information is directed to the machine's standard output and error files, which typically defaults to the console from which the DAC server was launched if the server was launched in a visible console mode. If the server is launched as a Windows service, the logging information is directed to the service log. If the server is launched with the command shell not visible, all logging information is deleted.</p>
Repository Name	Unique name for the DAC repository.
Scheduler.Poll.Interval	Frequency (in seconds) the DAC server polls for changes in the schedule configuration.
Script After Every ETL	<p>The name of the script or executable to be run after every execution plan.</p> <p>For more information, see the description of the property Script Before Every ETL.</p>
Script Before Every ETL	<p>The name of the script or executable to be run before every execution plan.</p> <p>For example, before running an execution plan, you might want to run a process or perform certain tasks. These can be contained in a script or executable. This file should be placed in the scripts subdirectory of the DAC server.</p> <p>The execution plan runs only after the external process has finished. Therefore, it is important that the script or executable does not fail.</p>
Server Log Level	Output logging level. Possible values are Finest, Finer, Fine, Config, Info, Warning, and Severe. The Severe value produces minimal log details, and Finest produces the most extensive amount of reporting.
SQL Trace	<p>Possible values are True and False.</p> <p>Indicates whether the SQL statements to the DAC repository and database connections are added to the log file. Possible values are True and False. The True value sends a hint to the database connectivity layer of the DAC server to enable SQL tracing; thus, every SQL statement that is run by the DAC server is spooled to the appropriate output log file.</p> <p>It is recommended that you set this property to False.</p>

Table 29. DAC System Properties Tab Properties

Property	Description
Test Run	<p>Possible values are True and False.</p> <p>When set to True, the execution plan will not stop on errors.</p>
Verify and Create Non-Existing Indices	<p>Possible values are True and False.</p> <p>Indicates whether indices defined in the DAC repository will be automatically created in the data warehouse database during an incremental load.</p> <p><b>NOTE:</b> When this system property is set to True, the DAC server verifies whether indices defined in the DAC repository are also defined in the data warehouse database. This verification process can delay the execution of an execution plan.</p>

## About the Informatica Servers Tab

The Informatica Servers tab allows you to register one or more Informatica servers and one Informatica Repository server and to specify how many workflows can be executed in parallel on each server. The DAC server automatically load balances across the servers. [Table 30](#) describes the properties of the tab.

Table 30. Informatica Server Tab Properties

Property	Description
Name	Name of Informatica Server or Informatica Repository Server.
Type	<p>Type of server:</p> <ul style="list-style-type: none"> <li>■ <b>Informatica.</b> The Informatica Server.</li> <li>■ <b>Repository.</b> The Informatica Repository Server.</li> </ul>
Server Hostname	The host machine name where the Informatica Server or Informatica Repository Server is installed.
Server Port	Port number used by the Informatica Server or Informatica Repository Server to listen to requests.
Login	Informatica Repository user login.
Password	Informatica Repository password.
Maximum Sessions	The maximum number of workflows that can be executed in parallel on the Informatica Server.
Repository Name	<p>Informatica Repository name.</p> <p><b>NOTE:</b> You deploy only one Informatica Repository Server, but you can deploy multiple Informatica Servers.</p>

## Actions Available

The following actions are available in the top pane toolbar and in the right-click menu when the Informatica Servers tab is active.

### Test Connection

This command allows you to test whether the Informatica Server or Informatica Repository Server is running. The DAC client will attempt to delegate the test to the DAC server. If the DAC server is not running, you will be asked whether you want to test the connection on the local client machine. The results of a test on the local client machine may not be true for the DAC server environment.

## About the Physical Data Sources Tab

The Physical Data Sources tab provides access to the connection properties for the physical data sources. In this tab, you can view and edit existing physical data source connections and create new ones. [Table 31](#) provides a description of the physical data source properties.

Table 31. Physical Data Sources Top Pane Properties

Property	Description
Name	Logical name for the physical data source.
Type	Physical data source type. Possible values are: <ul style="list-style-type: none"> <li>■ Source</li> <li>■ Warehouse</li> <li>■ Informatica Repository</li> <li>■ DAC Repository</li> <li>■ Other</li> </ul>
Connection Type	Type of database connection. Possible values are the following: <ul style="list-style-type: none"> <li>■ <b>Oracle (OCI8)</b>. Connect to Oracle using the tnsnames entry.</li> <li>■ <b>Oracle (Thin)</b>. Connect to Oracle using thin driver.</li> <li>■ <b>DB2</b>. DB2 UDB database.</li> <li>■ <b>DB2-390</b>. DB2 390 database.</li> <li>■ <b>MSSQL</b>. Microsoft SQL Server.</li> <li>■ <b>Teradata</b>. Teradata database.</li> <li>■ <b>Flat File</b></li> </ul>



Table 31. Physical Data Sources Top Pane Properties

Property	Description
Connection String	<p>If you are using...</p> <ul style="list-style-type: none"> <li>■ Oracle (OCI8), use the tnsnames entry.</li> <li>■ Oracle (Thin), use instance name.</li> <li>■ SQL server, use the database name.</li> <li>■ DB2-UDB/DB2-390, use the connect string as defined in the DB2 configuration.</li> <li>■ Teradata, use the database name.</li> </ul>
Table Owner	Name of table owner.
Max Num Connections	<p>Maximum number of database connections this connection pool can contain.</p> <p>For information about determining how to set this value, see <a href="#">“Determining the Number of Transactional and Data Warehouse Database Connections” on page 75</a></p>
DBHost	Host machine where the database resides. This field is mandatory if you are using Oracle (Thin), MSSQL, or Teradata, but is not required if you are using Oracle (OCI8), DB2, or DB2-390.
Port	Port where the database receives requests. Required for Oracle (Thin) and MSSQL databases. Not required for Oracle (OCI8), DB2, or DB2-390, or Teradata databases.
Priority	User-defined priority of the data source.
Data Source Number	User-defined number of the data source.
Default Index Space	Applicable to Oracle databases only. The default index space for the physical data source. When indices are dropped and created, they are created in this index space.

Table 32 provides a description of the child tabs in the bottom pane of the Physical Data Sources tab.

Table 32. Physical Data Sources Child Tab

Refresh Dates	<p>During an ETL process, this date is captured for all target tables and source tables of the type primary and auxiliary. The DAC uses this date in the change capture process, during parameter generation, when choosing between full and incremental loads, and when deciding whether to truncate a table. (Does not apply to micro ETL processes.)</p> <p><b>NOTE:</b> Refresh dates for micro ETL processes are captured in the Refresh Dates child tab of the Execution Plans tab.</p>
---------------	---

## Actions Available

The following actions are available in the top pane toolbar and in the right-click menu when the Physical Data Sources tab is active.

### Test Connection

This command allows you to test the connection to the database. The DAC client will attempt to delegate the test to the DAC server. If the DAC server is not running, you will be asked whether you want to test the connection on the local client machine. The results of a test on the local client machine may not be true for the DAC server environment.

## About the Email Recipients Tab

This tab allows you to set up a list of email addresses that will be notified about the status of the ETL process. [Table 33](#) describes the properties in this tab.

Table 33. Email Recipients Top Pane Properties

Column	Description
Name	Logical name of the user to be notified.
Email Address	Email address where the notification is sent.
Notification Level	<p>The notification levels are as follows:</p> <ul style="list-style-type: none"> <li>■ 10 -- Notifies recipient of success or failure of each task.</li> <li>■ 5 -- Notifies recipient of success or failure of the entire ETL process.</li> <li>■ 1 -- Notifies recipient that ETL completed successfully.</li> </ul>

## About the DAC Execute View

The Execute View provides access to functionality that allows you to run, schedule, and monitor execution plans.

This section includes the following topics:

- [About the DAC Execution Plans Tab on page 115](#)
- [About the DAC Current Run Tab on page 119](#)
- [About the DAC Run History Tab on page 122](#)
- [About the DAC Scheduler Tab on page 122](#)

## About the DAC Execution Plans Tab

The Execution Plans tab allows you to view and edit existing execution plans and to create new ones. [Table 34](#) provides a description of the properties in the top pane of the tab.

Table 34. Execution Plans Tab Top Pane Properties

Column	Description
Full Load Always	Indicates the specified ETL process will always execute a full load.
Keep Separate Refresh Dates	Used for micro ETL processes. Indicates refresh dates are kept separate for each run of the execution plan.
Prune Days	<p>When the source system is Oracle's Siebel CRM applications, the LAST_UPD column in the transactional database tables is used for incremental change capture. This timestamp reflects the actual event time. It is therefore possible for a data row to be committed to the transactional database with a LAST_UPD date that is older than the date on which the last refresh was executed. This will cause the data row to be missed in the subsequent extract (if based purely on LAST_UPD date).</p> <p>However, the LAST_UPD date column still provides an opportunity to improve the change capture process by overlapping the extraction date window by the number of days set in this parameter. The records extracted in the overlapped window are filtered by comparing this information with information in the Image table.</p> <p>The Prune Days setting ensures that the rows that had values in LAST_UPD older than values in LAST_REFRESH_DATE are not missed. This is a parameter that can be set based on experience with processes, such as remote sync, that potentially can cause records to be missed. This parameter cannot be less than 1.</p> <p>For example: Assume the table W_PERSON_D was refreshed on January 15th by querying the table S_CONTACT. And, the Prune Days setting was set to 5. The next time S_CONTACT is queried to load W_PERSON_D, the records that have a LAST_UPD value since January 10 are compared with the ROW_ID of the Image table to cover for any missing records between January 15 and January 10 (the overlap period).</p> <p>For source systems other than Siebel, the Prune Days setting is used in the same way except that the DAC subtracts the number of prune days from the LAST_REFRESH_DATE of a given source and supplies this as the value for the \$\$LAST_EXTRACT_DATE parameter.</p>
Last Designed	Date this execution plan was last designed.
Analyze	Indicates the tables associated with this execution plan will be analyzed.

Table 34. Execution Plans Tab Top Pane Properties

Column	Description
Analyze Truncated Tables Only	Indicates only truncated tables will be analyzed.
Drop/Create Indices	Indicates indices of the tables associated with this execution plan will be dropped and created.

Table 35 provides a description of the child tabs in the bottom pane of the Execution Plans tab.

Table 35. DAC Execution Plans Child Tabs

Tab	Description
Subject Areas	Displays the subject areas associated with the selected execution plan. You can also add subject areas to the selected execution plan.
Parameters	<p>Displays the selected execution plan's parameters for database connections and Informatica folders. It includes the following properties:</p> <ul style="list-style-type: none"> <li>■ <b>Type.</b> Possible values are:                             <ul style="list-style-type: none"> <li>■ <b>Folder.</b> Indicates an Informatica folder.</li> <li>■ <b>Datasource.</b> Indicates a database connection parameter.</li> </ul> </li> <li>■ <b>Name.</b> Logical name of the folder or database connection.</li> <li>■ <b>Value.</b> Physical name of the folder or database connection.</li> <li>■ <b>Source System.</b> Name of the source system associated with the parameter.</li> </ul>
Preceding Tasks	<p>Displays the tasks that must be completed before an ETL process is executed. Also allows you to add preceding tasks.</p> <p>It includes the following properties:</p> <ul style="list-style-type: none"> <li>■ <b>Priority.</b> Indicates the order in which the task runs. If two or more tasks have the same priority, the DAC will run them in parallel.</li> <li>■ <b>Inactive.</b> Indicates the task is inactive.</li> <li>■ <b>Command.</b> Command associated with the task.</li> <li>■ <b>Source System.</b> Source system container from which the task extracts data.</li> </ul>
Following Tasks	<p>Displays the tasks that must be completed after an ETL is executed. Also allows you to add tasks.</p> <p>It includes the same properties as the Preceding Tasks child tab.</p>

Table 35. DAC Execution Plans Child Tabs

Tab	Description
Ordered Tasks	<p>Displays tasks associated with the selected execution plan and the order in which they can be executed. It includes the following properties:</p> <ul style="list-style-type: none"> <li>■ <b>Depth.</b> The level of the task’s dependency. Tasks that have no dependencies are depth 0. Tasks that depend on other tasks of depth 0 are depth 1, and so on.</li> <li>■ <b>Primary Source.</b> Primary source table from which the task extracts data.</li> <li>■ <b>Primary Target.</b> Primary target table into which data is loaded.</li> <li>■ <b>Folder Name.</b> Name of the Informatica folder in which the task resides.</li> <li>■ <b>Task Phase.</b> Task phase of the ETL process. The DAC server uses the task phase to prioritize tasks and to generate a summary of the time taken for each of the phases.</li> <li>■ <b>Command.</b> Command associated with the task</li> <li>■ <b>Source System.</b> Source system container from which the task extracts data.</li> </ul> <p>The Details button opens a dialog box that contains the following information about the selected task:</p> <ul style="list-style-type: none"> <li>■ All Predecessors</li> <li>■ All Successors</li> <li>■ Immediate Predecessors</li> <li>■ Immediate Successors</li> <li>■ Source Tables</li> <li>■ Target Tables</li> <li>■ Conditional Tables</li> </ul>
Immediate Dependencies	<p>Displays the immediate dependent relationship between tasks that are generated during the automatic task generation process.</p> <p>It includes the following properties:</p> <ul style="list-style-type: none"> <li>■ <b>Task (Calculated).</b> Shows the named task’s source and target.</li> <li>■ <b>Predecessor Name.</b> Predecessor task for the named task.</li> <li>■ <b>Predecessor (Calculated).</b> Shows the source and target of the predecessor task.</li> </ul>

Table 35. DAC Execution Plans Child Tabs

Tab	Description
All Dependencies	Displays the dependent relationship for all tasks. The columns are the same as the Immediate Dependencies child tab.
Refresh Dates	<p>Applies to micro ETL execution plans (indicated by selecting the Keep Separate Refresh Dates check box in the Execution Plans tab).</p> <ul style="list-style-type: none"> <li>■ <b>Connection.</b> Logical name for database connection.</li> <li>■ <b>Refresh Date.</b> Last refresh time of the execution plan. This applies only when separate refresh dates are maintained. Used for micro ETL processing.</li> </ul>

## Actions Available

The following actions are available in the top pane toolbar and in the right-click menu when the Executions Plan tab is active.

### Run Now Command

This command submits a request to the DAC server to execute the execution plan. Inactive execution plans are not included in the process.

When an execution plan is run, the following logic is used to identify the tasks that need to be run and their order of execution.

- 1 Check on the previous run status.
  - a If the previous run executed a different execution plan from the one currently submitted (either manually or through the scheduler), abort the run.
  - b If the previous run executed is the same as the current execution plan and did not complete, continue from where it left off.
  - c If the previous run has completed, create a new run.
- 2 Check resources.
  - a If the DAC server cannot establish connections to the databases defined for that execution plan, abort the run.
  - b If there are no Informatica servers registered, abort the run.
  - c If the Informatica servers are registered and the DAC server is unable to ping them, abort the run.
- 3 Select initial tables.
  - a Find all the subject areas.
- 4 Select initial tasks: Find all the tasks that load into the above selected tables. (Tasks whose target table is one of the selected tables and that are non-optional.)
- 5 Recursively select all tasks. Depending on the non-optional dependency rules, figure out the prerequisite tasks recursively.

- 6 Select optional tasks (with the tasks already chosen).
- 7 Select dependencies: Load dependencies and compute for each task the number of all Depends On tasks. This number will be decremented for tasks higher up in the hierarchy as each of the tasks complete. When the number of Depends On tasks reaches zero, the task itself will become eligible to be run.
- 8 Identify the task details. Iterate through all the tasks selected, compute the task dependencies.
- 9 Identify the list of source tables for change capture and create tasks for change capture and sync for Siebel sources. Iterate through all the tasks and find the primary or auxiliary tables and group them by source database connection.
- 10 Execute pre-ETL tasks. If there are errors, stop the run.
- 11 Execute the change capture process. If there are errors, stop the process.
- 12 Execute the dependency graph. If there are errors, stop all the tasks that depend on the failed tasks. Continue to execute all others, until there are no more to run, and then stop.
- 13 Execute the change capture sync process. If there are errors, stop the process.
- 14 Execute the post-ETL process. If there are errors, stop the process.
- 15 At the end of the ETL process, the DAC server updates refresh timestamps for all source tables (primary or auxiliary) and target tables. The refresh timestamps are the database timestamps. The ETL history timestamps (Start and End timestamps) are the DAC server timestamps.

## About the DAC Current Run Tab

This tab displays a list of queued, running, and failed current ETL processes in the top pane. This list includes comprehensive information about each process.

Table 36 provides a description of the information provided in the top pane.

**NOTE:** Once an ETL process completes, it is accessible from the Run History tab.

Table 36. Current Run Tab Top Pane Properties

Column	Description
Execution Plan Name	The execution plan whose runtime instance is this record. This field is read only.
Process ID	ID for the process. This value is an integer that is incremented by 1 for every run. This value is stored as ETL_PROC_WID in all the data warehouse tables. This field is read only.

Table 36. Current Run Tab Top Pane Properties

Column	Description
Run Status	<p>Status of the run. The possible values are the following:</p> <ul style="list-style-type: none"> <li>■ <b>Queued.</b> Tasks for which the Depends On tasks are not yet completed. Displayed in yellow in the Current Run list.</li> <li>■ <b>Runnable.</b> Tasks for which the Depends On tasks have completed and are ready to be run but are waiting for an Informatica slot to be available.</li> <li>■ <b>Running.</b> Tasks for which the Depends On tasks have been completed, have gotten an Informatica slot, and are being executed. Displayed in blue.</li> <li>■ <b>Paused.</b> Task group members that are waiting for the other tasks in the group to be executed.</li> <li>■ <b>Failed.</b> Tasks that were executed but encountered a problem. Displayed in red.</li> <li>■ <b>Stopped.</b> Tasks for which one or more Depends On tasks have failed.</li> <li>■ <b>Completed.</b> All tasks have completed without errors. Displayed in green.</li> </ul>
Total Number of Tasks	Total number of tasks for this run. This field is read only.
Number of Failed Tasks	Sum total of tasks that have failed and that have stopped. This field is read only.
Number of Successful Tasks	Number of tasks whose status is completed. This field is read only.
Number of Tasks Still in Queue	Number of tasks whose prerequisite tasks have not completed, and the number of tasks whose prerequisite tasks are completed and are waiting for resources. This field is read only.
Start Timestamp	Start time of the ETL. Reflects the start time of every ETL attempt. For example, if the ETL fails and is run again, it gets a new start timestamp. The history of attempted runs is maintained in the audit trail for the run. This field is read only.
End Timestamp	End time of the ETL. Reflects the end time of every ETL attempt. For example, if the ETL fails and is run again, it gets a new end timestamp when the ETL is complete. The history of attempted runs is maintained in the audit trail for the run. This field is read only.
Duration	A calculated field that shows the difference between start and end time stamps.
Status Description	Displays messages generated during run time. You can add notes to this field for Completed runs.
Schedule Name	Name of a scheduled ETL process.



Table 37 provides a description of the tabs in the bottom pane.

Table 37. Current Run Child Tabs

Tab	Description
Tasks	<p>Displays runtime instances of the tasks. As the execution proceeds, the tasks are executed based on the dependency rules and some prioritization.</p> <p>As tasks complete, the tasks that depend on the completed tasks are notified and once their dependencies are completed, they become eligible to run. If a task fails, the administrator can address the failure and then requeue the task or mark it as completed. The DAC server polls for any changes in the failed task’s detail status. If a failed task detail is queued, the task itself gets back into the ready to run queue and all its dependent tasks get into the queued status.</p> <p>The rules of prioritization are as follows:</p> <ul style="list-style-type: none"> <li>■ Tasks with no dependencies are executed first.</li> <li>■ If a task has failed and has been requeued, it gets the maximum priority.</li> <li>■ Tasks with greater phase priorities are executed next. When several tasks of the same phase are eligible to run, the tasks with greater task priorities are executed next.</li> </ul>
Summary (Read Only)	Provides a summary (based on dynamic SQL) of the selected ETL. A start and end time for each phase is included.
Task Details	Provides a execution details about each task associated with the selected ETL.
Audit Trail (Read Only)	Provides the history of the selected ETL.

### Actions Available

The following actions are available in the top pane toolbar and in the right-click menu when the Current Run tab is active.

#### Start

This command restarts the selected ETL, after the ETL has failed, stopped, or been aborted.

#### Abort

This command causes an ETL in progress to abort. All currently running tasks will be aborted. The status of queued tasks and the ETL itself will change to Stopped.

### Auto Refresh

This command allows you to turn on and off the automatic screen refresh functionality and set the refresh interval.

### Get Log File

This command fetches the log file for this run from the DAC server and saves it in the ServerLog folder.

### Mark as Completed

This command changes the status of a stopped or failed ETL to Completed. In the audit trail for this ETL, the status is Marked as Completed. Use this command with caution; it can cause the data warehouse to be inconsistent.

### Analyze Run

This command saves a description of the run as an HTML file in the Log/Statistics folder.

## About the DAC Run History Tab

The Run History tab displays information about completed ETL processes. The information displayed in the top and bottom panes is the same as in the Current Run tab. For a description of the information displayed in the Run History tab and the actions available, see [“About the DAC Current Run Tab” on page 119](#).

## About the DAC Scheduler Tab

The Scheduler tab allows you to schedule ETL processes to be run either once at a later time or periodically. When you schedule an ETL or make changes to a schedule, the DAC server picks up the information from the DAC client. The DAC server polls the DAC repository for changes periodically at a frequency set in the DAC system properties.

The top pane of the Scheduler tab lists ETL processes that have been scheduled. The bottom pane allows you to schedule an ETL process. [Table 38](#) provides a description of the information displayed in the top pane.

Table 38. DAC Scheduler Tab Top Pane Properties

Column	Description
Execution Plan	Execution plan scheduled.
Last Schedule Status	The last run status of the scheduled ETL process. Possible values are Running, Completed or Stopped.
Next Trigger	Time the scheduled ETL process will next run.

Table 38. DAC Scheduler Tab Top Pane Properties

Column	Description
Status Description	Description of the last ETL run. Possible values are Running, Completed, or the reason the process stopped.
Recurrence	Indicates how often the schedule will be run.
Inactive	Indicates whether the schedule is inactive.



# Index

## C

### change capture

- about 78
- filter 84

### configuration tags

- working with 51

## D

### DAC repository

- command line options 70

### DAC server

- command line access 66
- monitor icons 33
- running automatically 65

### data flow

- Online Analytical Processing (OLAP) database, about and diagram 11

### data warehouse

- architecture 10
- overview 9

### Data Warehouse Administration Console (DAC)

- about 12
- Current Run tab 119
- DAC features 14
- DAC window 22
- Design view 89
- editable lists 34
- Email Recipients tab 114
- Execute view 114
- Execution Plans tab 115
- exporting metadata 64
- importing metadata 63
- Indices tab 94
- Informatica Servers tab 111
- menu bar 23
- navigation tree 34
- object ownership 38
- Physical Data Sources tab 112
- process life cycle 17
- repository objects, about 16
- Run History tab 122
- starting DAC client 19
- System Properties tab 107
- Tables tab 92
- Tasks tab 97
- top pane toolbar 27
- user interface 22

### deleted records

tracking 84

## E

### execution plan

micro ETL 58  
monitoring processes 61  
scheduling 60

## F

### flat views

querying 37

## I

### Informatica

mappings, creating 46  
replacing workflow with SQL file 72  
server sessions 73

### Informatica repository

importing objects 46

## O

### Online Analytical Processing database

Data Warehouse, data flow into 11

### Oracle Business Analytics Data Warehouse

overview 9

### Oracle Business Analytics Warehouse

adding columns 44  
adding indices 46  
adding new table 44  
architecture 10  
architecture components 11  
customizing 42  
overview 9

## Q

### query functionality

flat views querying 37  
query commands 35  
query operators 35  
query procedures 36

## R

### refresh dates

about 60

### right-click menu

common commands 29  
Design view commands 30  
Execute view commands 32  
Setup view commands 31

**S****source system container**

- about 15
- copying 41
- creating 41

**source system parameters**

- setting 49

**SQL files**

- using as execution type 77

**subject area**

- creating 56
- designing 54

**T****task group**

- creating 49

**task level parameters**

- setting 50

**task phase dependency**

- setting 48

**Tasks**

- creating tasks for workflows 47

