

Oracle® Revenue Management and Billing Analytics

Version 2.2.0.0.0

Admin Guide

Revision 1.0

E69342-02

June, 2016

Oracle Revenue Management and Billing Analytics Admin Guide

E69342-02

Copyright Notice

Copyright © 2016, Oracle and/or its affiliates. All rights reserved.

Trademark Notice

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

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

License Restrictions Warranty/Consequential Damages Disclaimer

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

Warranty Disclaimer

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

Restricted Rights Notice

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

U.S. GOVERNMENT RIGHTS

Oracle programs, including any operating system, integrated software, any programs installed on the hardware, documentation, and/or technical data delivered to U.S. Government end users 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 any operating system, integrated software, any programs installed on the hardware, documentation, and/or technical data shall be subject to license terms and restrictions as mentioned in 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). No other rights are granted to the U.S. Government.

Hazardous Applications Notice

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

Third Party Content, Products, and Services Disclaimer

This software and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third party content, products, or services.

Preface

About This Document

This guide aims to act as a reference guide to an administrator user and helps him with day-to-day tasks, as well as provides some pointers on how to handle some commonly seen change requests. The document is organized in the form of a comprehensive questionnaire and covers most of the administrative tasks.

Intended Audience

This document is intended for the following audience:

- End-Users
- System Administrators
- Consulting Team
- Implementation Team

Organization of the Document

The information in this document is organized into following sections:

Section No.	Section Name	Description
Section 1	Introduction	Gives an overview of the product architecture, along with an outline of implementation and administration processes.
Section 2	Job Scheduling	Lists the common FAQs on job scheduling via ORMBA Admin UI.
Section 3	Dashboard User Extensions	Lists the main FAQs on dashboard extensions available to users.
Section 4	ETL User Extensions	Lists the main FAQs addressing common ETL user extensions and how to handle them.
Section 5	Handling Globalization	Gives an overview of the ORMBA features to support a global implementation of the product.
Section 6	Monitoring & Troubleshooting	Common maintenance tasks.
Section 7	Generic	Some general queries regarding the ORMBA processes.
Section 8	Performance & Scalability	Tips and tricks on improving system performance and scalability.
Section 9	Archival	Lists common FAQs about archiving data in various layers of ORMBA system.

Related Documents

You can refer to the following documents for more information:

Document	Description
<i>Oracle Revenue Management and Billing Analytics Functional Overview</i>	Lists the features and architecture of Oracle Revenue Management and Billing Analytics.
<i>Oracle Revenue Management and Billing Analytics Install Guide</i>	Lists the pre-requisites, supported platforms, and hardware and software requirements for installing the Oracle Revenue Management and Billing Analytics application. It also explains how to install the Oracle Revenue Management and Billing Analytics application.
<i>Oracle Revenue Management and Billing Analytics Security Guide</i>	Describes how to configure security for Oracle Revenue Management and Billing Extractors and Schema and Oracle Revenue Management and Billing Analytics.

Change Log

Revision	Last Update	Updated Section	Comments
1.0	June 2016		For ORMBA v 2.2.0.0.0

Contents

1.	Introduction	8
1.1	Administration Overview	8
2.	ETL Job Management.....	9
2.1	How do I schedule jobs?	9
2.2	How do I enable or disable a job?.....	10
2.3	Can I enable or disable multiple jobs in one go?	11
2.4	How do I set up job frequency?	11
2.5	How can I configure job dependency?.....	12
2.6	Can I prioritize jobs?	13
2.7	What happens when a job fails?	14
2.8	What happens when a job fails even after maximum number of retries?.....	14
2.9	What happens when a parent job fails?	15
2.10	What happens when schedules for a job overlap?.....	15
2.11	Which time zone will a job follow?	15
2.12	Can I work with an external scheduler?.....	15
3.	Dashboard User Extensions	16
3.1	How do I change labels in a report?	17
3.2	How do I add a new attribute into a report?.....	18
3.3	How do I hide or remove an attribute from a report?	19
3.4	How do I create a new report?	20
3.5	How do I bring a new user defined attribute into a subject area?	21
3.6	How do I add a filter to a dashboard?	22
3.7	How do I add a filter to an analysis?	23
3.8	How do I create a new dashboard?	24
3.9	How do I export data from a dashboard?.....	24
3.10	How do I share the results of an analysis?.....	25
3.11	How do I change theme and settings?.....	25
4.	ETL User Extensions	27
4.1	How can I extend the replication layer?	27
4.2	How can I extend the replication layer to add a new replication group?	27
4.3	What are the different ways to extend an ETL layer?	28
4.4	How do I populate one or more UDFs in a dimension table?.....	28
4.5	How do I populate one or more CDFs in a dimension table?	30
4.6	How do I populate UDDGENs or UDMs in a fact table?.....	31
4.7	How do I create a new Mview over an existing or a newly created fact?	31
4.8	Can I make an extension procedure specific to a particular RMB instance?	32

5.	Handling Globalization.....	33
5.1	How do I handle multiple divisions of business?	33
5.2	How do I work with multiple currencies?	33
5.3	How do I work with multiple languages?.....	33
5.4	How do I handle data without division?	34
5.5	How do I work with multiple ORMB instances?	34
6.	Monitoring & Troubleshooting.....	36
6.1	How do I monitor replication?	36
6.2	How do I stop replication?	36
6.3	Where can I view replication status?.....	37
6.4	How do I enable automated consistency checks?	38
6.5	Will I receive a notification if an ETL job fails?.....	39
6.6	Where can I view ETL errors?.....	39
6.7	How do I configure failure notification for replication?	41
6.8	How do I configure delay notification for replication?	42
6.9	How do I change GG parameters?	43
6.10	Where can I view the currently active job queue?	45
6.11	Where can I view the pending job queue?	45
7.	Generic.....	46
7.1	Is there a risk of processing partial transactions?	46
7.2	Why do we need Oracle GoldenGate when ODI can extract data?.....	46
7.3	How do we replicate historical data during initial load?	46
8.	Performance and Scalability	47
8.1	What are the recommendations for best performance?	47
8.2	What are the different levels of scalability?	48
8.3	How do I utilize Application level parallelism?	48
8.4	How do I utilize DB level parallelism?	49
8.5	How do I enable parallelism for a fact/dimension?.....	50
8.6	What are the recommended DB settings?.....	50
8.7	How do I perform DB tuning for performance?.....	52
8.8	What is the average lag while replicating data?	52
9.	Archival	53
9.1	How do I purge data in the replication schema?	53
9.2	How do I purge data from the staging schema?	54
9.3	How do I archive data from the warehouse?	55

1. Introduction

1.1 Administration Overview

An administrator has a pivotal role in the configuration, maintenance and troubleshooting of the product. Some of the administrative tasks are one-time configurations, whereas some others are routine checks. The administrative tasks are done mainly using the following tools:

- ORMBA Administration UI
- ORMBA Dashboards
- Oracle Data Integrator (ODI) Studio
- Oracle GoldenGate (GG) Monitor
- Business Intelligence (BI) Admin Tool

This manual is organized as a collation of most frequently asked questions about ORMBA installation, implementation and maintenance. The questions are grouped under various heads for easy reference. The main groups are:

- Job Scheduling
- Dashboard User Extensions
- ETL User Extensions
- Handling Globalization
- Monitoring & Troubleshooting
- Performance & Stability
- Archival

FAQs under each of these aspects are detailed below.

2. ETL Job Management

ORMBA allows you to schedule the data-loading job. You can perform data loading in two steps:

1. Initial Load: This job involves one-time loading of historical data from source system and can be initiated at any point of time.
2. Incremental Load: This job runs on a pre-defined schedule and normally does not require manual intervention.

As part of installation, the system is set on Initial Load mode where you can do the historical data loading in one go. Once you have done the Initial Loading, for further upload, you can use the Incremental Load mode, during which data loading occurs based on a defined schedule.

Some frequently asked questions regarding job scheduling are listed below:

2.1 How do I schedule jobs?

You can schedule jobs to start at specified date and time. You can also define how frequent you want the job to execute.

Where: ORMBA Administration UI

To schedule a job, follow the procedure below:

1. Navigate to the Target Entity Definition page under the Administration menu. You will see a list of entities that are currently available in the system.
2. Select an entity; it can be a fact or dimension or a materialized view.
3. Under the Scheduler Configuration group box, check the Start Date field.
4. You can select the date on which jobs using this entity should start.
5. Depending on the values configured in Slice Duration Frequency and Slice Duration fields, the jobs are scheduled for execution. To know more about this configuration, see section [2.4](#).
6. Click Save.

Note: Despite this configuration, the actual execution of the jobs start on the 'Start Date' configured in ODI Scheduler, provided the job is enabled.

Scheduler Configuration

Start Date (yyyy-MM-dd HH:mm:ss)	2015-07-24 00:00:00		Maximum Retry *	<input type="text"/>
Schedule Type *	Incremental Load		Retry Interval (Minutes) *	<input type="text"/>
Slice Duration Frequency *	Day(s)		Analyze Table Flag *	No
Slice Duration *	<input type="text" value="1"/>		Daily Sync Flag *	No
Maximum Parallel Count *	<input type="text" value="2"/>			

Figure 1: Scheduling a Job

2.2 How do I enable or disable a job?


Where: ORMBA Administration UI

To enable a job, follow the procedure below:

1. Navigate to the Job Configuration page under the Administration menu. You will see the list of jobs pre-configured in the system.
2. Select a job.
3. Click Edit button. This opens the job and you can view the details.
4. Select **Yes** from the drop-down list in the Enabled field to enable a deactivated job. On the other hand, select No if you want to disable a job.
5. If you enable a disabled job, you can select a 'Slice Start Time' to indicate the time at which the job should next run. You can select any time after the last successful execution time of the job.

Note: ETL jobs works on data slices. While considering data for replication, the ETL job selects data with time stamp greater than or equal to 'Slice Start Time'. You can edit the 'Slice Start Time' of a job while enabling it, provided it has never been executed. Once executed, the field remains un-editable.

6. Click Save.



The screenshot displays the 'Job Configuration' form. The fields and their values are as follows:

Field	Value
Description	Fiscal Calendar Dimension
Instance Number	1
Target Entity	BD_FISCAL_CAL
Package Name	
Slice Start Time	2000-01-01 00:00:00
Enabled	Yes
User Extension Procedure	
User Extension Procedure (Post Job)	
Execution Sequence	

The 'Enabled' field is highlighted with a red box, and a red asterisk is visible next to it, indicating it is a required field. The dropdown menu for 'Enabled' is currently set to 'Yes'.

Figure 2: Enabling a Job

2.3 Can I enable or disable multiple jobs in one go?

Yes, you can enable or disable multiple jobs in one go, using the Bulk Edit facility in the Job Configuration page.

Where: ORMBA Administration UI

To enable or disable multiple jobs in a go, follow the procedure below:

1. Navigate to the Job Configuration page under the Administration menu. You will see the list of pre-configured jobs in the system.
2. Click Bulk Edit button.
3. Select the check box against the jobs you want to enable or disable.
4. Click the Enable or Disable button, depending on whether you want to enable or disable the selected jobs.
5. A confirmation message appears. Click Yes.

Job Configuration Bulk Update

View ▼	<input checked="" type="checkbox"/> Enable	<input checked="" type="checkbox"/> Disable		Detail
<input type="checkbox"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>
<input type="checkbox"/>	Description	Instance Number	Target Entity	
<input checked="" type="checkbox"/>	Account Dimension	1	BD_ACCT	
<input checked="" type="checkbox"/>	Adjustment Type Dimension	1	BD_ADJ_TYPE	
<input type="checkbox"/>	Fiscal Calendar Dimension	1	BD_FISCAL_CAL	
<input type="checkbox"/>	Service Agreement Type Dimension	1	BD_SA_TYPE	

Figure 3: Enabling or Disabling Multiple Jobs

2.4 How do I set up job frequency?

You can define the frequency at which an incremental job needs to run by mentioning Slice Duration and Slice Duration Frequency against the Target Entity Definition. The target entity can be a fact, or a dimension, or a materialized view, or a temporary table.

Where: ORMBA Administration UI

To define how often a job needs to run, follow the procedure below:

1. Navigate to the Target Entity Definition page under the Administration menu. You will see the list of entities that are available in the system.
2. Select the entity against which you want to set a job frequency.
3. The Scheduler Configuration group box displays if there is a schedule already defined.
4. In the Slice Duration Frequency field, select the measure in which you would like to define the job frequency. The values available are: Year, Quarter, Month, Week, Day, Hour, Minute.
5. In the Slice Duration field, enter the frequency at which the job should run. The number mentioned here along with the measure mentioned in the previous step defines the job frequency.
6. Click Save.

Scheduler Configuration

Start Date (yyyy-MM-dd HH:mm:ss)	2015-07-24 00:00:00	Maximum Retry *	1
Schedule Type *	Incremental Load	Retry Interval (Minutes) *	30
Slice Duration Frequency *	Day(s)	Analyze Table Flag *	No
Slice Duration *	1	Daily Sync Flag *	No
Maximum Parallel Count *	1		

Figure 4: Configuring Job Frequency

2.5 How can I configure job dependency?

While creating a job, you can configure if the job is dependent on any other job(s). You can define any number of levels of dependency using the Job Configuration page.

Where: ORMBA Administration UI

To define a dependent job(s) for a job, follow the procedure below:

1. Navigate to the Job Configuration page.
2. Select the job for which you want to define a dependency.
3. Click Edit.
4. Under the Dependencies tab, click Add.
5. Select the parent job to which the current job is dependent on. You can add one or more jobs as parent jobs.
6. Click Save.

Dependencies

View ▾
+ Add
✕ Delete
>>

Parent Job - Entity Name
Customer Dimension BD_CUST

Figure 5: Configuring Job Dependency

Note: Once you define one or more parent(s) to a job, the job is executed only after the parent job(s) is or are executed. Additionally, if a parent job fails, the dependent jobs are not executed and wait in the execution queue.

2.6 Can I prioritize jobs?

Yes, you can prioritize jobs at the same *peer level* (with respect to the defined job dependencies) by defining the sequence of their execution. System queues up the jobs for execution based on the execution sequence you configure. However, the actual execution depends on other factors like, whether the job is dependent on another job and has it completed.

Where: ORMBA Administration UI

To define the priority of jobs, follow the procedure below:

1. Navigate to the Job Configuration page.
2. Select the job with highest priority.
3. Click Edit.
4. Enter **1** in the Execution Sequence field.
5. Click Save.
6. Select the job that needs to be executed after the previous one.
7. Click Edit and enter **2** in the Execution Sequence field. The jobs follow the sequence thus defined.

Note: All jobs taken up by the scheduler are executed asynchronously. The scheduler arranges the jobs in the queue based on the Execution Sequence. However, the actual execution depends on whether the job has a dependency or not. Thus, configuring an execution sequence does not ensure that the jobs with higher sequence start only after the execution of one with lower sequence.

2.7 What happens when a job fails?

When a job fails, you can configure to receive an email notification (see section [6.5](#)). System also offers a re-try option for failed jobs. You can configure the number of times the system should attempt retry, as well as the interval (in minutes) between two consecutive retries.

Where: ORMBA Administration UI

To define the number of retries the system should attempt, follow the procedure below:

1. Navigate to the Target Entity Definition page.
2. Under the Scheduler Configuration group box, check the Maximum Retry field. The default value of this field is **1**. That is, system tries to execute the job once and if it fails, there will not be any retry.
3. Enter the number of times the system should attempt to execute the job (including the first attempt) in the Maximum Retry field.
4. Enter the number of minutes after which the system should attempt job retry in the Retry Interval field. This determines the interval between successive retries.
5. Click Save.

Scheduler Configuration

Start Date (yyyy-MM-dd HH:mm:ss)	2015-07-24 00:00:00	Maximum Retry *	2
Schedule Type *	Incremental Load	Retry Interval (Minutes) *	30
Slice Duration Frequency *	Day(s)	Analyze Table Flag *	No
Slice Duration *	1	Daily Sync Flag *	No
Maximum Parallel Count *	1		

Figure 6: Configuring Job Retry

When a job fails at first attempt, it appears in the Job Status View page, under the Job Execution Status tab with 'Error' status. When the job succeeds after retry, a new entry is added in Job Execution Status page with status 'Reprocessed'. When a parent job fails, all child jobs will be in 'Pending' status.

2.8 What happens when a job fails even after maximum number of retries?

When a job fails even after maximum number of retries, it appears in the Job Status View page, under the Job Execution Status tab, with status marked as 'Abandoned'. Additionally, in the Job Configuration page, the respective job appears disabled. You can see the error messages in the Job Execution Status tab. Also, check the Operator tab of ODI.

You have to troubleshoot to identify the reason why the job is failing and rectify the problem. To see the error, check the Operator tab of ODI. Once rectified, you can enable the job from the Job Configuration page. To see how you can enable a job, see section 2.2 of this document.

2.9 What happens when a parent job fails?

If a parent job fails, its dependent jobs are put on hold. Only after the parent job executes successfully, the dependent jobs are executed.

The dependent jobs that are on hold due to the parent job failure appear in the Jobs Waiting for Execution tab of the Job Status View page with wait reason 'Waiting for a change to occur in dependent entity'. To know more about the Job Status View page and the queues available, see *ORMBA_Functional_Overview.doc*.

2.10 What happens when schedules for a job overlap?

If the previous schedule is not over, the current schedule will wait. The job thus waiting appears in the Jobs Waiting for Execution tab of the Job Status View page with wait reason 'Waiting on dependent entity to be loaded'.

2.11 Which time zone will a job follow?

In case of an implementation scenario where more than one time zones are involved, the job scheduling follows the time zone of the server in which ORMBA warehouse runs.

2.12 Can I work with an external scheduler?

Yes. In place of ODI scheduler, you can use other external third-party schedulers as well.

3. Dashboard User Extensions

We deliver several out-of-box dashboards that contain several reports and charts that are required for normal analysis. However, you can customize the existing dashboards and reports to address your business-specific requirements.

There are several dashboard-level extensions that you can do to existing dashboards/reports, as well as create new ones. You can configure most of the dashboard extensions by making use of the various OBIEE features itself. The main features of OBIEE that you will commonly use for updating ORMBA dashboards are:

- Page Options button on a Dashboard Page

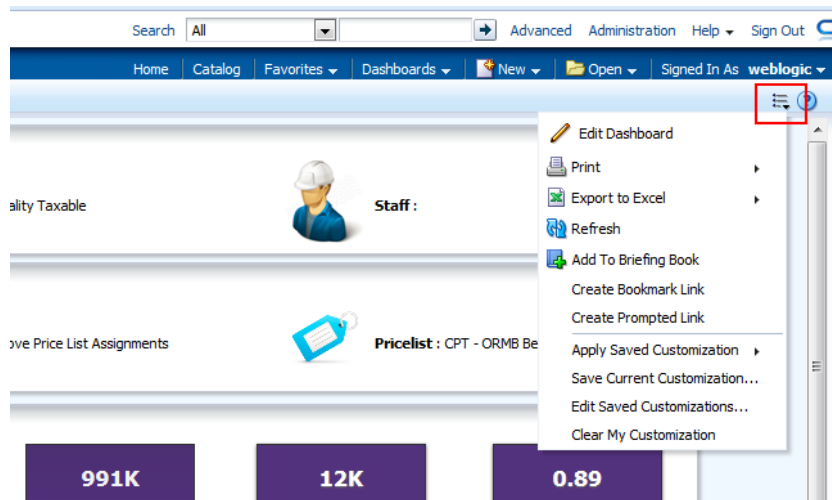


Figure 7: Page Options

- Properties button on a report

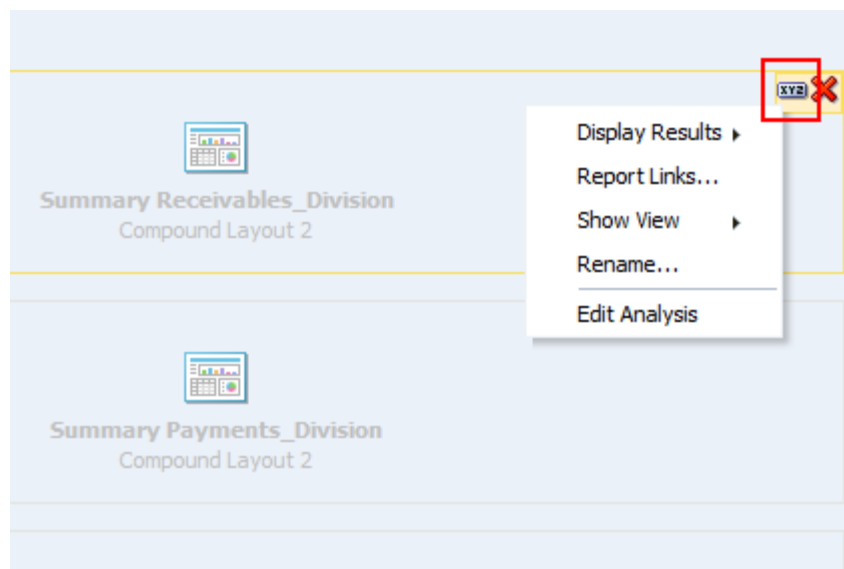


Figure 8: Properties button

- Properties button of a column/field

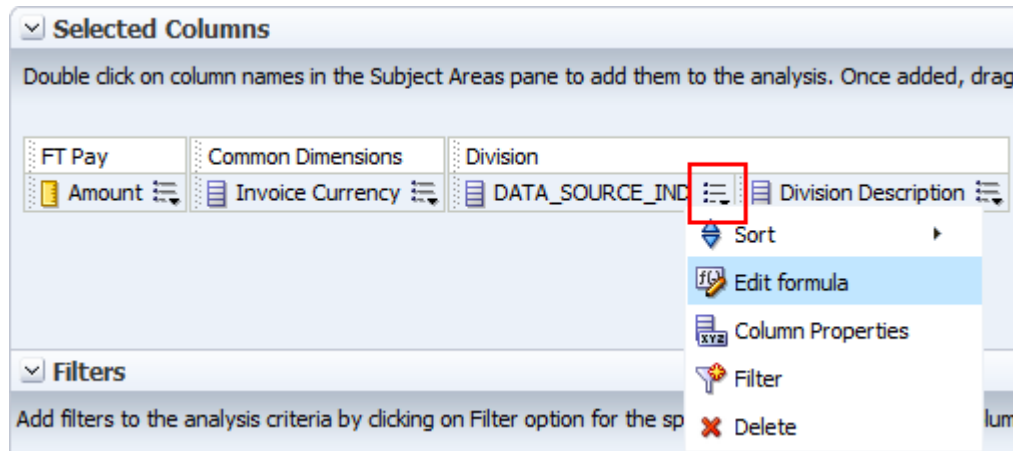


Figure 9: Properties button

Some dashboard extensions, especially when a new subject area is introduced or when a new fact/dimension is introduced in existing subject areas, require configurations using the BI Administration Tool.

The following are some frequently asked queries related to customizing the dashboard reports and charts.

3.1 How do I change labels in a report?

Where: ORMBA Dashboards

To change the label of a report:

1. Open the dashboard.
2. Click the Page Options (☰) button and select Edit Dashboard option.
3. Navigate to the page where the report appears.
4. Click the Properties (⌘) button of the report and select Edit Analysis option.
5. Navigate to Criteria tab.
6. Hover over the Properties (⌘) button of the field and select the Column Properties option.
7. Navigate to Column Format tab.
8. Edit the 'Column Heading' field to change the label.
9. You can save this as the default label.

An example scenario:

You can change the label of Invoice Account field to Billing Account by following the above procedure.

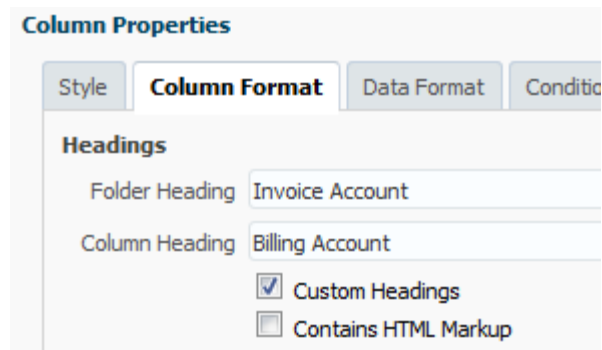


Figure 10: Changing report labels

3.2 How do I add a new attribute into a report?

Where: ORMBA Dashboards

You can add a new attribute, either pre-defined or user-defined, to a report by dragging and dropping the attribute from the right tables. To do an attribute, follow the procedure below:

1. Open the dashboard.
2. Click Page Options button (☰) and select Edit Dashboard option.
3. Navigate to the report you want to edit.
4. Click the Properties button of the report and select Edit Analysis option. All facts, dimensions and materialized views associated with the subject area are available on the left pane.
5. Expand a table to view the attributes within it.
6. Drag and drop the required attribute - either fields (📄) or measures (📊) into the right pane.

An example scenario:

You need to add Adjustment Status attribute to the Adjustments Printable Report. To do this, follow the procedure below:

1. Open the Financial Transactions dashboard and open the Adjustments page.
2. Click Export button (📄) to open the Adjustments Printable Report.
3. Click the Page Options button and select Edit Dashboard option.
4. Click the Properties button of the report and select Edit Analysis option.
5. Expand the Adjustments folder and select the Adjustment Status field.
6. Drag and drop the field to the required place on the right pane.
7. Save the analysis. The Adjustment Status now appears on the Adjustments Printable Report.

Year	2011	Month	December	Amount		
Division	Credit/Debit	Customer Segment	Adjustment Type	Contract Type	Adjustment	Cancellation
930-California Branch	Debit		Bill Correction (USD)	Demo Services CA	\$256,610.00	
Debit Total					\$256,610.00	
930-California Branch Total					\$256,610.00	
940-NY FI Brance	Debit		Bill Correction (USD)	Banking Services (NY)	\$154,570.00	
Debit Total					\$154,570.00	
940-NY FI Brance Total					\$154,570.00	
ORMB Bank - Europe Division	Debit		Bill Correction (EUR)	XYZ Bank - Volume charges	\$181,947.16	
			Bill Correction (GBP)	XYZ Bank - Volume charges	\$3,267.53	
Debit Total					\$185,214.69	
ORMB Bank - Europe Division Total					\$185,214.69	
ORMB Bank - UK Division	Debit		Bill Correction (EUR)	XYZ Bank - Volume charges	\$1,366.71	
			Bill Correction (GBP)	XYZ Bank - Volume charges	\$443,266.48	
Debit Total					\$444,633.19	
ORMB Bank - UK Division Total					\$444,633.19	
Test Data Division	Debit		Bill Correction (USD)	Demo Services	\$101,196.00	
Debit Total					\$101,196.00	
Test Data Division Total					\$101,196.00	
Grand Total					\$1,142,223.88	

Figure 11: Adding a new attribute to a report

3.3 How do I hide or remove an attribute from a report?

Where: ORMBA Dashboards

To hide or remove an attribute from a report, follow the procedure below:

1. Open the dashboard.
2. Click Page Options button (⚙️) and select Edit Dashboard option.
3. Navigate to the report you want to edit.
4. Click the Properties button of the report and select Edit Analysis option.
5. Click the Edit View button (✎) on top of the Pivot Table.
6. Scroll down to Layout section.
7. Click More Options (⋮) button against the attribute you want to hide or remove.
8. Select 'Hidden' to hide the attribute, or 'Remove Column' to remove the attribute from the report.
9. Click Done and save the analysis. The

An example scenario:

You need to remove Adjustment Status attribute from the Adjustments Printable Report. To do this, follow the procedure below:

1. Open the Financial Transactions dashboard and open the Adjustments page.
2. Click 🖨 button to open the Adjustments Printable Report.

3. Click the Page Options button and select Edit Dashboard option.
4. Click Edit View button on top of the Pivot Table.
5. Under the Layout section, click More Options button against Adjustment Status field.
6. Select Remove column option from the menu and click Done.





Figure 12: Hiding or removing an attribute

3.4 How do I create a new report?

Where: ORMBA Dashboards

To create a new report, follow the procedure below:

1. Open the dashboard.
2. Click the New button and select Analysis option.
3. Select a Subject Area from the list available. The newly created report will use attributes from the selected subject area.
4. The left side of the screen displays the list of tables available under the selected subject area. Click on a table to see the fields available.
5. Select the required field and drag it to the Selected Columns pane on the right side. You can use the Ctrl key to select multiple fields in a go.
6. Click the Results tab to preview the look and feel of the report.
7. Click Save and select the folder within which you would like to save the new report. Ideally, this should be same as the subject area you used to create the report.
8. Navigate to the dashboard where you want to include the newly created report.
9. Click Page Options button and select Edit Dashboard option.
10. Go to the Catalog section on the left side of the page.
11. Navigate within the catalog to select the newly created report.
12. Drag and drop the report to the right pane, positioning it as you desire.
13. Click Save.



3.5 How do I bring a new user defined attribute into a subject area?

Where: OBI Admin Tool

To bring a new attribute to a subject area, follow the procedure below:

1. In the physical layer, navigate and select the attribute you want to bring into the subject area.
2. Drag and drop the entities to the required subject area in the business layer.
3. If needed, edit the entity name in the business layer.
4. Drag and drop the updated entity to the corresponding folder under the presentation layer.
5. Save the changes to the RPD file.
6. Deploy the new RPD file using Enterprise Manager. The new attribute is now available for selection under the subject area.

An example scenario:

You want to bring a new user defined attribute called 'Account Region' to Invoice Account table within Financial Transaction subject area. To do this, follow the procedure below:

1. In the BI Admin Tool, expand the Invoice Account table in the Physical layer.
2. Drag the field UDF1_DESCR from the Physical layer to under the Invoice Account table in the Business layer.
3. Double-click the field to edit the name of the field to Account Region.
4. Drag the Account Region field from the Business layer to the corresponding table in the Presentation layer.
5. Save the changes to the RPD and deploy it using the Enterprise Manager tool. The field Account Region is now available for selection in the ORMBA Dashboards, under the Financial Transaction subject area.





Figure 13: Adding an attribute to a subject area

3.6 How do I add a filter to a dashboard?

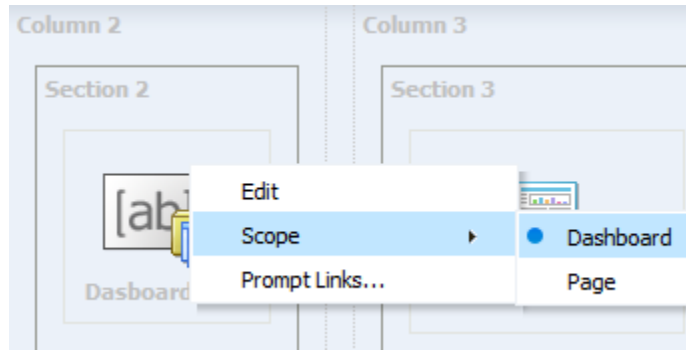
Where: ORMBA Dashboards

To add a filter to a dashboard, create a dashboard prompt with one or more columns and then create a filter using the same column(s). You need to add the filter to all reports within the dashboard so that the filter condition applies to the data. You can read the detailed procedure below:

1. Open the dashboard.
2. Click the New button and select Dashboard Prompt option.
3. In the Subject Areas menu, select the subject area under which you plan to add the dashboard prompt.
4. Click the New button and select Column Prompt.
5. Navigate and select the column you want to add as a dashboard prompt. You can see a preview of the filter in the Display section below.
6. Save the dashboard prompt.
7. Navigate to the dashboard where you want to introduce the filter.
8. Click the Page Options button and select Edit Dashboard option.




9. From the Dashboard Objects menu, drag and drop a column to the dashboard.
10. From the Catalog, browse and select the newly created dashboard prompt.
11. Drag and drop the dashboard prompt to the newly added column.
12. You can set the scope of a dashboard prompt to a dashboard or a page.



13. Save the changes. Once you set a dashboard prompt at the dashboard level, you need to add the same attribute as a filter in the required analysis of the page. To do this, follow the procedure in section [3.7](#).

3.7 How do I add a filter to an analysis?

To add a filter to an analysis within a dashboard, follow the procedure below:

1. Open the dashboard.
2. Click the Page Options button and select Edit Dashboard option.
3. Scroll down to the analysis where you want to add the filter.
4. Click Properties button and select Edit Analysis option.
5. Navigate to the Criteria tab and click the  button against Filters group box.
6. Select the column you want to add as the filter and click OK.
7. In the New Filter popup, select **is prompted** in the Operator field and click OK.
8. Save the changes.


An example scenario:


You want to add customer segment as a filter in the Receivables Printable Report. To do this, navigate to the Financial Transactions dashboard and edit the Receivables Printable Report analysis. In the Criteria tab, select Customer Segment Description column and add it as a filter. The newly added filter appears in the list and is now applied on the analysis.

3.8 How do I create a new dashboard?

Where: ORMBA Dashboards

To create a new dashboard, follow the procedure below:


1. Open the dashboard.
2. Click the New button and select Dashboard option.
3. You can now drag and drop existing analyses from the Catalog to the newly created dashboard.
4. You can also add new pages to the newly created dashboard by clicking the Add Dashboard Page () button.



3.9 How do I export data from a dashboard?

Where: ORMBA Dashboards

ORMBA allows you to export the content of most of its dashboards into various printable formats. A printable report shows aggregated data grouped under various attributes. To export the contents of a printable report, follow the procedure below.

1. Open the dashboard.
2. Click the  button on the left side. This opens the printable report.
3. Scroll down to the bottom of the report.
4. Click the Export link. You will see the following options:
 - PDF
 - Excel
 - Excel 2003
 - Excel 2007+
 - Powerpoint
 - Powerpoint 2003
 - Powerpoint 2007+
 - Web Archive (.mht)
 - Data
 - CSV Format
 - Tab delimited Format
 - XML Format
5. Save the exported data.

Note: You can also customize the format of excel columns using OBIEE's feature.

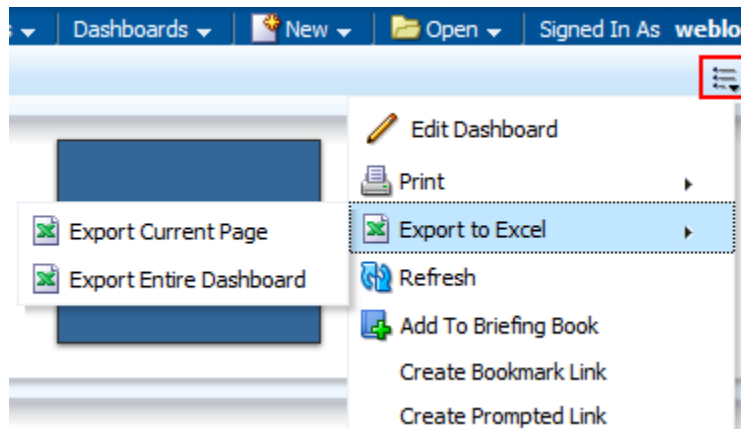
3.10 How do I share the results of an analysis?

Where: ORMBA Dashboards

To share the results of an analysis from a dashboard, you can export the contents of a page, or the entire dashboard into an excel sheet which can be shared.

To export the contents of a page or a dashboard, follow the procedure below:

1. Open the dashboard.
2. Click the Page Options button and select Export to Excel option.
3. You will see two options to export the data:
 - Export Current Page
 - Export Entire Dashboard



4. On selecting the required option, an excel sheet opens containing the data of the page or dashboard, based on your selection.

Apart from this, you can also export data from printable reports provided by ORMBA. To do this, you can use the Export link available at the bottom of a report. While trying to export a printable report, you have many options like PDF, Excel, PowerPoint, Web Archive (.mht) and Data (.csv, .xml etc.).

3.11 How do I change theme and settings?

Where: ORMBA Dashboards

To change the theme and settings of a dashboard, follow the procedure below:

1. Open the dashboard.
2. Click Page Options button and select Edit Dashboard.
3. Click Tools button and select the Dashboard Properties option.
4. In the Style field, you can select the required theme from the drop-down list and click OK. The selected theme will be applied to all dashboard pages within the selected dashboard.

The custom themes provided by ORMBA are

- **Theme 1:** Lighter scheme
- **Theme 2:** Darker scheme

The other dashboard settings you can alter are shown in the image below.

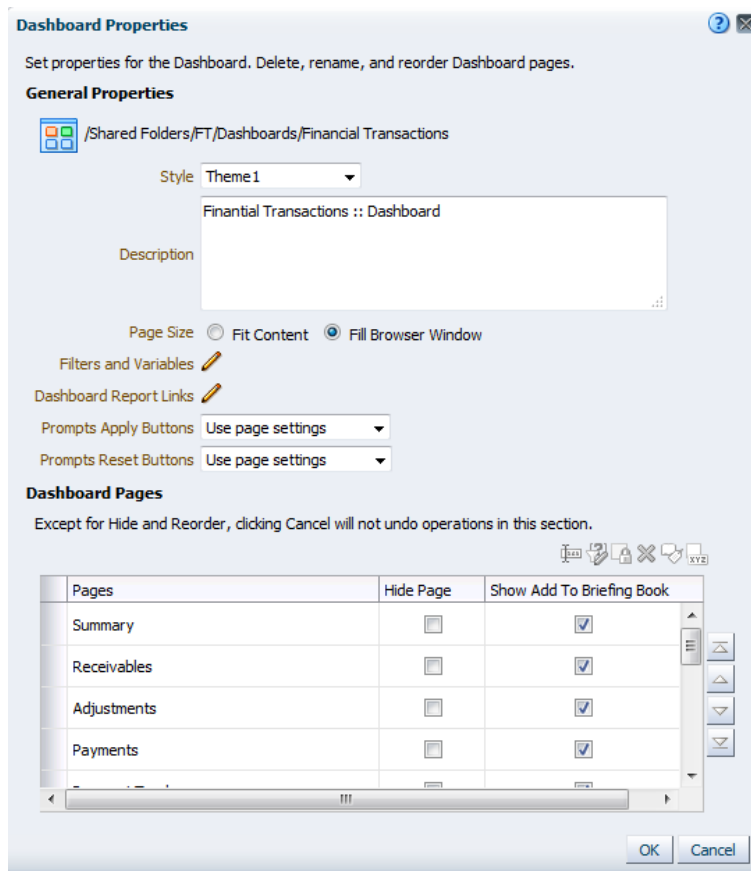


Figure 14: Dashboard Properties

4. ETL User Extensions

The following section explains how to extend the ETL layer of ORMBA to suit your requirements. Most of the below tasks require a fair working knowledge of Oracle Data Integrator (ODI).

4.1 How can I extend the replication layer?

Where: ORMBA Administration UI

To extend the replication layer, you can add tables to the existing replication groups by following the procedure below:

1. Log on to ORMBA Administration UI and open the Source Table Definition page.
2. Enter the table name that you want to add, and associate it with an existing replication group. You can find more information on configuring the Source Table Definition page in the ORMBA Functional Overview document.

Note: Ensure that the table has a primary key, before adding to the replication group.

3. Save the changes.
4. Execute the following shell scripts which are available as part of ORMBA Media Pack ETL Component:
 - a. createSourceModel.sh
 - b. importData.sh

For more information on executing the scripts, see ORMBA Install Guide.

Note: We recommend you to create new replication groups for customizations, as updating an existing replication group with a new table involves stopping and re-starting the existing jobs of the replication group.

4.2 How can I extend the replication layer to add a new replication group?

Where: ORMBA Administration UI

To extend the replication group, add a new replication group by following the procedure below:

2. Log on to ORMBA Administration UI and open the Source Table Definition page.
3. Enter the name of the table to be added to the new replication group.
4. In the Replication Group field, enter the name of the new replication group. You can find more information on configuring the Source Table Definition page in the ORMBA Functional Overview document.
5. Save the changes.
6. Repeat step 2 to add the required tables to the newly created replication group. **Note:** Once created, the replication group will be available in the list of values of the field and you can select the same as explained in section [4.1](#).

7. Go to Oracle GoldenGate Home in source database server and create a new folder within the **dirdat** directory. The name of the folder should be same as the newly created replication group.
8. Repeat the previous step in the target database server also.
9. Execute the following shell scripts which are available as part of ORMBA Media Pack ETL Component:
 - a. createSourceModel.sh
 - b. importData.sh

For more information on executing the scripts, see ORMBA Install Guide.

4.3 What are the different ways to extend an ETL layer?

ORMBA is a highly flexible product platform that can be customized based on specific business requirements. You can extend the ETL layer of ORMBA through user extensions across the following entity levels:

- Dimensions
 - 50 User Defined Fields (UDFs)
 - 25 Characteristic Data Fields (CDFs)
- Facts
 - 50 User Defined Measures (UDMs)
 - 5 User Defined Dimensions (UDDs) per subject area
 - 50 User Defined Degenerate Dimensions (UDDGENs)

You can write ODI custom procedures or edit the replication views to populate data on any of the above custom data fields.

4.4 How do I populate one or more UDFs in a dimension table?

ORMBA supports two methods for populating the custom data fields: Edit a replication view or create an ODI procedure.

Editing a Replication View

You can edit an existing replication view to populate a custom data field. This method is typically used in scenarios where the changes are minimal and you do not have sufficient expertise in ODI. However, you should have a good knowledge of schemas and table structures, and should be comfortable with SQL.

Creating an ODI Custom Procedure

You can write an ODI custom procedure utilizing one or more custom fields of a target entity and associate it to the same target entity using ORMBA Administration UI. This method is comparatively simpler if you have good knowledge of ODI.

ORMBA provides a simple process to extend the out-of-the-box dimensions and facts to meet installation-specific business requirements. This makes use of a configurable procedure (CM procedure) with a predefined signature.

An example scenario:

Suppose you want to populate the invoice collection mode in one of the UDFs of the dimension table - BD_ACCT. To do this, you need to first create a custom procedure to populate each account's payment mode details, and then associate the procedure to the jobs configured for the corresponding dimension.

Follow the procedure below to create a User Exit procedure for **BD_MSG** dimension:

1. Log on to ODI Client and navigate to the Designer tab.
2. Navigate to **Customizations** project and open the **Dimensions** folder.

Note: The best practice is to maintain a separate project for all customizations in ODI. Also, ensure that ODI objects (Facts, Dimension and Materialized Views) are organized separately within the project.
3. Right-click the **Procedures** item and select the **New Procedure** option to create the custom procedure.
4. Enter the name of the procedure and a meaningful description on the right pane. The naming convention for the procedure is CM_<Entity being customized>_UDX. In this scenario, the name of the procedure would be **CM_BD_ACCT_UDX**.
5. Un-check Multi-Connections check box and select **Oracle** in the Target Technology field.
6. Select **Tasks** tab and click **+** button to add a new task.
7. Double-click the Task Name field and enter a name for the task.
8. Double-click the Target Command field to open the Edit Expression window.
9. Write the custom procedure logic in the window and click OK.

```

create or replace procedure <%=odiRef.getSchemaName( )%>.cm_bd_acct_udx (v_sess_no in number,
v_slice_beg_ts in date,
v_slice_end_ts in date) as
begin
  update udx_bd_acct udx
  set (udf1_cd,udf1_descr) = (select (nvl((select distinct 'A' from CI_ACCT_APAY
  where ACCT_ID = udx.SRC_ACCT_ID and
  (trunc(sysdate) between start_dt and nvl(end_dt,to_date('01/01/2020','dd/mm/yyyy'))))
  , 'M')), 'Mode of Payment' from dual)
  where sess_no = v_sess_no;
Exception
  WHEN OTHERS THEN
  raise_application_error(-20001, 'An error was encountered - '||SQLCODE||' - ERROR - '||SQLERRM);
end;

```

Figure 15: Custom Procedure

10. Select **RMB1** in the Target Context field and **Replication** in the Target Logical Schema field.
11. Save the custom procedure and click **▶** button to execute the custom procedure.

- Go to the Operator tab of ODI and navigate to Date, Today, <custom procedure> to view the status of execution. On successful execution of the procedure, it creates the same custom procedure in the replication schema also.

Follow the procedure below to configure the user extension procedure for the dimension:

- Log on to ORMBA Administration UI and open the Job Configuration page.
- Search and open the job configured for the target entity BD_ACCNT for editing.
- Enter **CM_BD_MSG_UDX** in the User Execution Procedure field and click Save.
- Select **Yes** in the Enabled field and save the configuration.
- Monitor the job execution and verify the data in the dimension.

Job Configuration	
Description	Account Dimension
Instance Number	1
Target Entity	BD_ACCT
Package Name	
Slice Start Time	
Enabled	* Yes <input type="button" value="v"/>
User Extension Procedure	CM_BD_ACCT_UDX
User Extension Procedure (Post Job)	
Execution Sequence	

Figure 16: Job Configuration for User Extension

4.5 How do I populate one or more CDFs in a dimension table?

Where: ORMBA Administration UI

To populate a Custom Data Field (CDF), follow the procedure below:

- Log on to ORMBA Administration UI and open the Characteristic Map page.
- Click Create and select a dimension in the Entity Name field. Only those dimensions that have Characteristic Entity defined against them (using the Target Entity Definition page) are available for selection.
- Search and select the Characteristic Code to be mapped to the CDF. You can select a Characteristic Code configured against the selected dimension in the source system.
- Save the definition.

An example scenario:

Suppose you want to populate the invoice mailing option chosen by an account against the BD_ACCT dimension. To do this, you need to map the MAILBILL to a characteristic field of BD_ACCT dimension table by following the procedure below:


Characteristic Map	
Entity Name BD_ACCT	
View ▼	 Detach
Mapping Field	Characteristic Code
Characteristic 1	MAILBILL
Characteristic 2	No Mapping

Figure 17: Characteristic Mapping

4.6 How do I populate UDDGENs or UDMs in a fact table?

To populate UDDGENs or UDMs in a fact table, you can follow the same procedure explained in section [4.4](#).

4.7 How do I create a new Mview over an existing or a newly created fact?

1. Create a new materialized view using SQL Developer.
2. Log on to ODI and bring the newly created materialized view to the target model using reverse engineering.
3. Create a new ODI Interface & Package to populate data to the newly created materialized view.
4. Log on to ORMBA Administration UI and open the Target Entity Definition page.
5. Create a new definition for the materialized view.
6. Open the Job Configuration page and create a job for the newly created MView.

4.8 Can I make an extension procedure specific to a particular RMB instance?

Yes, you can create an extension procedure specific to a source instance.

Where: ORMBA Administration UI

To make an extension procedure specific to an RMB instance, associate the custom procedure to a job rather than a target entity. The job configurations are instance-specific and hence custom procedures associated with a job will be specific to a source instance.

Job Configuration

Description	Message Dimension
Instance Number	1
Target Entity	BD_MSG
Package Name	
Slice Start Time	2015-11-20 15:38:01
Enabled	* Yes
User Extension Procedure	CM_BD_MSG_UDX
User Extension Procedure (Post Job)	
Execution Sequence	

Figure 18: User extension specific to an RMB instance

5. Handling Globalization

5.1 How do I handle multiple divisions of business?

The term division in ORMBA is synonymous with the division in ORMB (the source system), where it could be different lines of business, or different jurisdictions.

ORMB Analytics caters to multiple divisions of business and **Division** features as a common slice and dice factor on almost all ORMBA dashboards.

ORMBA also offers data level security based on division in the dashboards. That is, you can restrict or allow a user's access to data from one or more divisions. To know more about this, see the ORMBA Security Guide.

It is also possible to generate ORMBA reports with values indicated in the division currency.

5.2 How do I work with multiple currencies?

ORMBA handles data in three different currencies:

- **Corporate Currency:** This is the base currency of an ORMBA installation. All cross-divisional reports display values in the corporate currency. You can configure this using the Global Settings page of ORMBA Admin UI.
- **Division Currency:** This is the currency defined against a division in ORMB. You can generate reports to display the data in division currency. For example, the China division reports indicate the revenue in Yuan (CNY).
- **FT Currency:** This is the currency in which a rate or price is defined against a financial transaction. For example, in case of an *invoice* transaction, the flat rate is defined as 5 Euros.

You can use the Indicative FX page of ORMBA Admin UI to define indicative exchange rate of other currencies to corporate currency.

In all ORMBA dashboards, wherever amounts are mentioned, it clearly indicates the currency in which it is displayed, either by symbol or currency code.

5.3 How do I work with multiple languages?

ORMBA supports any language supported by Oracle databases. You can set the default language for an ORMBA installation using the Global Settings page of ORMBA Admin UI.

ORMBA dashboards display descriptions in any of the selected language, as available in the source system (ORMB). In case of multiple instances, ensure that you select a common language.

Note: Only data and descriptions appear in the selected language. All labels and fields across the dashboards are, by default, displayed in English.

5.4 How do I handle data without division?

A division is an important slice and dice factor for ORMBA dashboards and reports. In case of data that arrives without a division, ORMBA assigns a default division based on the configuration in ORMBA Administration UI > Global Settings page > Identifier for Global Division field. Default value of the field is **GLB**.

If you would like to provide a user with access to data without division, add **GLB** to the division list against the user in the Data Level Security page of Administration UI.

For example, if a transaction feed arrives without information on the division, ORMBA assigns **GLB** as its division. All users with **GLB** in their division list will be able to see the data. However, if the user does not have a division-specific data level security definition, they are able to view the data.

5.5 How do I work with multiple ORMB instances?

ORMBA can have a global implementation scenario, where it receives source data from separate ORMB systems in different geographical locations; that is, from different source instances. Data from these different instances is loaded to a single warehouse; thereby enabling ORMBA to perform Global/Enterprise level analytics.

By default, an ORMBA installation holds a single, default, source system instance – namely **1**. However, it is possible to work with multiple ORMB instances.

The data warehouse is not instance-specific, whereas the replication layer (schema) is instance-specific, and therefore requires a synchronization to align the multiple source system instances.

In case of a new source system being introduced into an ORMBA installation, follow the high-level procedure below:

1. Using the Source Instance page of the ORMBA Administration UI, create a new source instance definition. The status appears **New**.
2. Configure the database tables of the new source instance using the Source Table Definition page of ORMBA Administration UI.
3. Review and edit the `ormba.properties` file to include details of the new source instance. To do this, edit the properties under the following heads:
 - `#ormba source configuration`
 - `#ODI Agent details`
 - `#Initial Load Properties`

Note: For more information, you can refer to the ORMBA installation guide.

4. Execute the scripts given below in the order mentioned:
 - i. `addDefaultData.sh`
 - ii. `addInstance.sh`
 - iii. `configureGG.sh`
 - iv. `checkConfiguration.sh`

- v. createSourceModel.sh
- vi. importData.sh
- vii. createReplicationModel.sh

Note: You can find the above scripts in the ORMBA Media Pack ETL Component.

5. (Optional) Edit and schedule the jobs for the new instance.
 - You can configure different user exec procedures against the replicated jobs.
 - You need to re-schedule the jobs for the new instance.
 - You need to update the GG configurations for the new instance.
6. Using the Data Level Security page of the ORMBA Admin UI, allow access to data from the newly created source instance to the required users.
7. Configure global settings using the Administration UI.

6. Monitoring & Troubleshooting

6.1 How do I monitor replication?

Where: Oracle GoldenGate Monitor

You can monitor replication via the Oracle GoldenGate Monitor. Oracle also offers a plug-in to integrate the Oracle GoldenGate Monitor with Oracle Enterprise Manager for centralized monitoring.

The screenshot displays the Oracle GoldenGate Monitor web interface. The top navigation bar includes links for 'Purge Data', 'User Profile', 'About', 'Help', and 'Logout'. The main content area is divided into several tabs: 'Data and Alerts View', 'Logs', 'Configuration', 'Problem Summary', 'Historical Data', 'Alert History', and 'Alert Definition'. The 'Data and Alerts View' tab is active, showing a tree view of the system hierarchy: System: system > Hosts: (1) > Host: localhost > Instance: localhost-7809 > Delivery: I1TD. The 'Delivery: I1TD' process is highlighted, and its status is shown as 'Stopped'. The right-hand pane provides an 'Overview' of the process, including a 'Lag Status: Stopped' indicator. Below this, the 'Attributes' section lists various metrics such as 'Checkpoint Position', 'Connected', 'Delta Deletes', and 'Delta Discards'. The 'Alerts' section shows a recent alert for 'AlertTest1' with a severity of 'Error' and a notification message: 'The alert was triggered because: Status = STOPPED_STATE.' The timestamp of the alert is '01/12/2016 12:44:53 PM'.

Figure 19: Oracle GoldenGate Monitor

6.2 How do I stop replication?

You can stop replication by executing the shell script `stopGG.sh`, available in the ORMBAs Media Pack ETL Component. This shell script stops the GoldenGate processes and / or cleans the data imported to replication tables during initial load, depending upon the mode of execution.

The various options available for executing the script are:

Command	Purpose
stopGG.sh CLEAN=M1,M2	Cleans the data and stops the GG processes associated with the specified models, e.g. M1, M2
stopGG.sh CLEANSERVICE=M1,M2	Stops the GG processes associated with the specified models, e.g. M1,M2
stopGG.sh CLEANSERVICE=ALL	Stops the GG processes associated with all models of the instance in which they are running
stopGG.sh CLEAN=ALL	Cleans the data and stops the GG processes associated with all models of the instance in which they are running

Note: The ALL option is available only in ONLINE mode. If you are running the script in OFFLINE mode, you have to specify the model names, separated by comma.

6.3 Where can I view replication status?

Where: ORMBA Administration UI

You can view the replication status of a table using the Source Table Definition page in Administration GUI. The field 'Replication Status' indicates if the replication is done or not.

Source Table Definition

Product Flag	Revenue Management and Billing
InstanceNum	1
Table Name	CI_CHG_TYPE_L
History Type	Override <input type="button" value="v"/>
Effective Date Column	<input type="text" value=""/>
Replication Flag	No <input type="button" value="v"/>
Purge Flag	No <input type="button" value="v"/>
Replication Retention Days	<input type="text" value="60"/>
Replication Status	Yes <input type="button" value="v"/>
Model Code	COMMON3 <input type="button" value="q"/>

Figure 20: Viewing Replication Status

6.4 How do I enable automated consistency checks?

Where: ODI Studio

You can enable automated consistency check for a fact, dimension, or materialized view using the ODI Studio.

1. In the ODI Studio, open the project.
2. From the folder structure, select either Facts/Dimensions/Materialized Views node.
3. Expand the Mappings node within.
4. From the folder structure, select the fact or dimension or materialized view for which you want to enable the consistency checks.
5. In the bottom pane, select the Physical tab.
6. On the right pane, you will see the properties of the entity you selected.
7. Expand the Integration Knowledge Module section and select the HEALTH_CHECK option.
8. To enable the consistency check, set the value of the option as **True**.
9. Save the changes.
10. Navigate to the Packages folder in the left pane and re-generate the scenario. The consistency check is now enabled for the selected fact, dimension, or materialized view.




Figure 21: Enabling Consistency Check

Note: We recommend you to disable consistency checks in a working environment to avoid performance issues.

6.5 Will I receive a notification if an ETL job fails?

Where: ODI Studio

Yes, if you have configured email notification for ETL job failures, you will receive email notification when an ETL job fails.

To enable email notification for ETL job failure, follow the procedure below:

1. Open ODI Studio.
2. Select the Designer tab and navigate to Project > Facts > Procedures.
3. Select `SendOraMailTLS` and open the definition to edit the attributes.
4. On the right pane, under the Options menu, edit the following fields: To, From, CC, SUBJECT, and BODY.
5. Set the value of `ENABLE_EMAIL` field to 'True'.

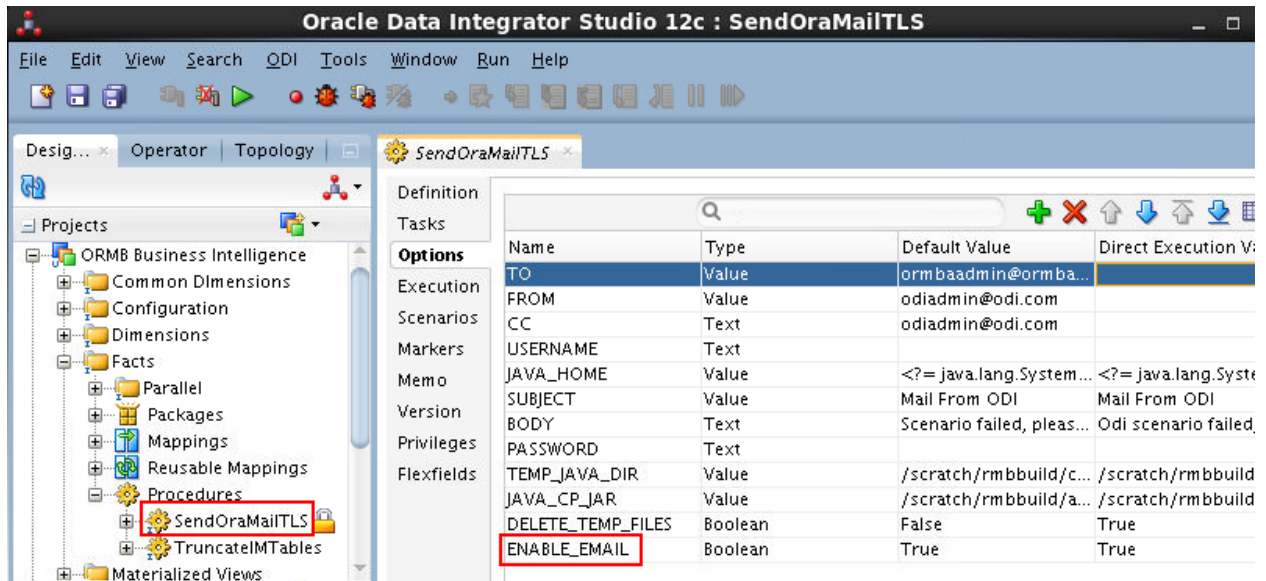


Figure 22: Configuring Email Notification

Note: By default, email notification facility is disabled for ETL jobs.

6.6 Where can I view ETL errors?

Where: ORMBA Administration UI

You can view the ETL errors in the Job Execution Status tab of Job Status View page. To view the jobs that have failed, filter the list of jobs with status **Error**. Against each failed job, you can click the **Status** link to view the error details.

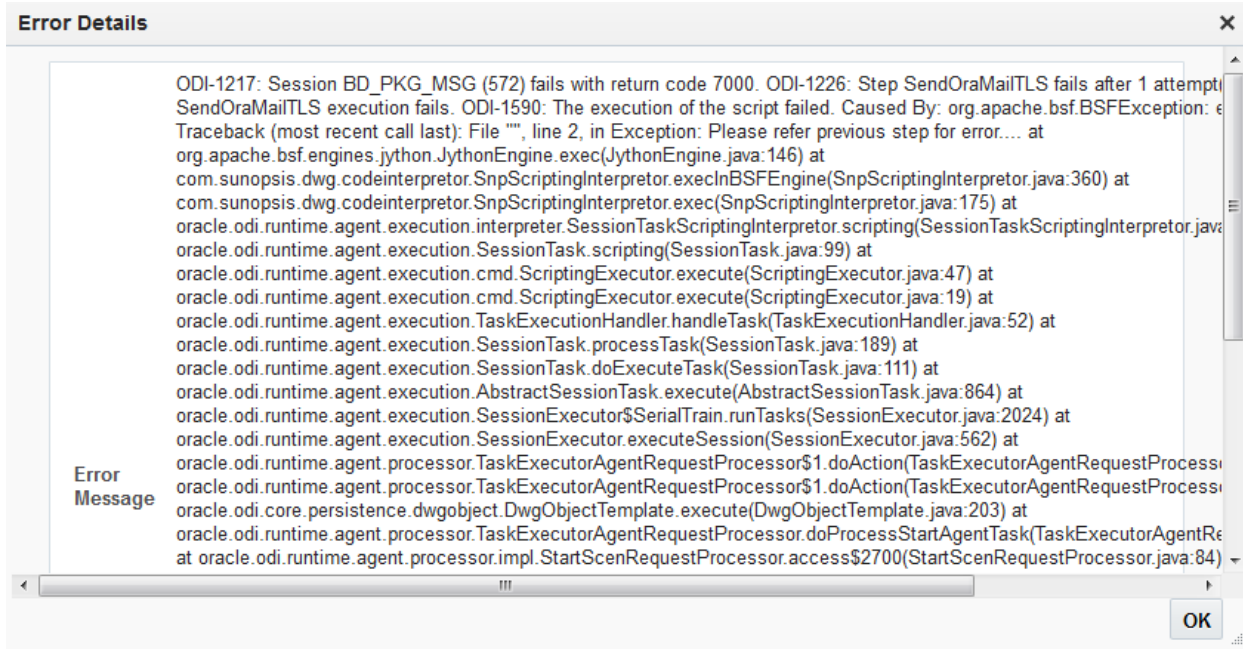


Figure 23: Viewing Error Details in ORMBA Admin UI

Where: ODI Studio

You can also view the ETL errors in the Operator tab of ODI. This view is more detailed and includes the reason for the error.




Figure 24: Viewing Error Details in ODI

6.7 How do I configure failure notification for replication?

Where: Oracle GoldenGate Monitor (GG Monitor)

You can configure failure notification for replication using the Oracle GoldenGate Monitor. Configure a mailer for failure notification using the following procedure:

1. Log on to GG Monitor.
2. Navigate to the Alert Definition tab.
3. Enter a name for the definition in the Definition Name field.
4. In the Alert Severity field, select Error.
5. You can enter the frequency at which you need the notification (in case the situation continues) in the Suppression Time field.
6. Select Delivery in the Object Type field.
7. In the Expression group box, define the condition for notification. To do this, select **Status** as attribute and create a definition. You can select one or more of the below status:
 - Abended
 - Stopped
 - Error
8. Select a user in the User Assignment group box.
9. Select the required replication group from the Available Objects field and move it to Assigned Objects field.
10. Create content for the alert by including relevant attributes like, Group Name, Message, Start Time, Last Checkpoint Timestamp, etc in the Additional Notice Attributes group box.
11. Click Submit to save the definition.

An example scenario:

You can configure a failure notification for I1TXN replication group as shown in the image below:

General Info				
* Definition Name	Error for I1TXN			
Alert Severity	Error			
Suppression Time	30	minutes		
Object Type	Delivery			
Expression				
	Attribute	Operator	Value	Unit
<input type="checkbox"/>	Status	==	Abended	
<input type="checkbox"/>	Status	==	Stopped	
<input type="checkbox"/>	Status	==	Error	

Figure 25: Alert definition for error

As per the definition, if the status of replication of I1TXN model is Abended, Stopped, or Error, you will get an alert in GG Monitor. If needed, you can configure to trigger an email alert also.

6.8 How do I configure delay notification for replication?

You can configure delay notification for replication using the Oracle GoldenGate Monitor.

Where: Oracle GoldenGate Monitor (GG Monitor)

1. Log on to GG Monitor.
2. Navigate to the Alert Definition tab.
3. Enter a name for the definition in the Definition Name field.
4. In the Alert Severity field, select Warning.
5. You can enter the frequency at which you need the notification (in case the situation continues) in the Suppression Time field.
6. Select Delivery in the Object Type field.
7. In the Expression group box, define the condition for notification. To do this, select **Lag** as attribute and create a definition.
8. Select a user in the User Assignment group box.
9. Select the required replication group from the Available Objects field and move it to Assigned Objects field.
10. Create content for the alert by including relevant attributes like, Lag, Total Inserts, Total Deletes, Total Operations, etc.
11. Click Submit to save the definition.

An example scenario:

You can configure delay notification for I1TXN replication group as shown in the image below:

Edit Alert Definition

▼ **General Info**

* Definition Name: Lag Alert for I1TXN

Alert Severity: Warning

Suppression Time: 30 minutes

Object Type: Delivery

▼ **Expression**

	Attribute	Operator	Value	Unit
<input type="checkbox"/>	Lag	>	200	sec

Figure 26: Alert definition for lag

As per the definition, if the replication lag of I1TXN model exceeds 200 seconds, you will get an alert in GG Monitor. If needed, you can configure to trigger an email alert also.

6.9 How do I change GG parameters?

Where: ORMBA Administration UI

You can alter the Oracle GoldenGate parameters as per your requirements. The currently available parameters are:

- Table-level
 - THREADRANGE
 - IGNOREUPDATES
 - IGNORETRUNCATE
 - IGNOREDELETES
 - IGNOREINSERTS
- Replication Group-level
 - MAXTRANSOPS
 - NUMFILES
 - TRIMSPACES
 - NOTRIMSPACES
 - COMPRESSDELETES
 - COMPRESSUPDATES
 - GROUPTRANSOPS
 - MAXSQLSTATEMENTS

To change the GG parameters for replication, follow the procedure below:


1. Log on to the ORMBA Administration UI and open the GG Parameters page.
2. Select the Replication Group and/or Table Name against which you want to change the GG parameter.
3. Select the Param Code and Param Value as per your requirement, and save the changes.
4. Run the **importData.sh** in RELOAD mode by giving the corresponding model name as argument.

Note: You can find the script in the ORMBA Media Pack ETL Component.

An example scenario:

For the Transaction Feed Management table, you have configured alert notification for a replication lag of more than 120 seconds. Suppose you notice a lag of 500 seconds against the replication group I1TXN. To improve the replication performance, you decide to increase the number of threads for the replication group. To do this, follow the procedure below:

1. Open the GG Parameters page to create a new configuration.
2. Select **I1TXN** in the Replication Group field and **CI_TXN_DETAIL** in the Table Name field.
3. Select the parameter **THREADRANGE()** in the Param Code field and enter **4** in the Param Value field.
4. Select **Replicat** in the Param Scope field.
5. Run the **importData.sh** in **RELOAD** mode by giving **I1TXN** as argument.



Golden Gate Parameter Configuration

Product *	Revenue Management and Billing
Instance Number *	1
Replication Group *	I1TXN
Table Name	CI_TXN_DETAIL
Param Code	1
Param Value *	4
Param Scope *	Replicat

Figure 27: Changing GG Parameters

6.10 Where can I view the currently active job queue?

Where: ORMBA Administration UI

The active jobs can be of any of the following status:

- Submitted
- Running
- Pending
- Ready for Execution

You can find the jobs with Submitted, Running or Pending status in the Job Execution Status tab of Job Status View page. To view jobs in Ready for Execution status, you can check the Jobs Ready for Execution tab of the Job Status View page.

6.11 Where can I view the pending job queue?

Where: ORMBA Administration UI

You can view the pending jobs in Jobs Waiting For Execution tab of Job Status View page. There are several reasons why a job could be waiting in the job queue. Some of them require you to take actions to push them into execution, whereas others will be automatically executed. You can see the list of common wait reasons below, along with whether it would require user intervention or not.

Reason	Require user intervention (Y/N)
Entity is deactivated	Yes
Configured scenario <scenario-name> does not exist	Yes
This job is scheduled to run at <time>	No
This one time load interface has already executed successfully	No
Reached max retries allowed	Yes
Approximately <mins> minutes till next retry	No
Maximum number of parallel MV execution reached	No
Waiting for a change to occur in dependent entity	No
Waiting on dependent entity to be loaded	No
Waiting for at least one minute of incremental data to arrive	No
Cut off time has been enabled. All entities will be loaded only up to <time>. (This reason is just a general info)	No

7. Generic

7.1 Is there a risk of processing partial transactions?

No.

7.2 Why do we need Oracle GoldenGate when ODI can extract data?

Oracle GoldenGate (OGG) in combination with the Oracle Data Integrator (ODI) is used to implement an efficient data replication strategy. This ensures minimal impact to the source systems

Oracle GoldenGate 12c offers a real-time, log-based Change Data Capture (CDC) and replication software platform to meet the needs of today's transaction-driven applications. The software provides capture, routing, transformation, and delivery of transactional data across heterogeneous environments in real time. Oracle GoldenGate only captures and moves committed database transactions to insure that transactional integrity is maintained at all times. The application carefully ensures the integrity of data as it is moved from the source database or messaging system, and is applied to any number of target databases or messaging systems.

The latest release sets Oracle GoldenGate further apart from competition by bringing extreme performance and advanced capabilities such as intelligent and integrated data delivery and cloud-based real-time replication, while simplifying product deployment significantly

7.3 How do we replicate historical data during initial load?

In Oracle GoldenGate installation, the initial load setup is used to perform a batch load to populate tables initially. This is a part of initial synchronization run with Oracle GoldenGate. To initiate initial run, you need to execute the B1_SYNC_CONTEXT package in the RMBBI_Automation (refer to the Replication folder) project. Before executing the B1_SYNC_CONTEXT package, set the name of instance in the package.

8. Performance and Scalability

ORMBA handles a huge volume of data for analysis and aims to offer near-real-time results. Hence, it is imperative that we keep the response time to the minimum, thereby ensuring high system performance.

The system performance depends mainly on the volume of data being handled, service window being used and the infrastructure available. While tuning the performance and scalability of the ORMBA system, we should consider the current volume of data AND the rate at which the data is growing.

8.1 What are the recommendations for best performance?

To ensure best performance of the system, we make the following recommendations:

- Ensure that the data is evenly spread across replication groups. We recommend having a replication group for each subject area. You can also configure multiple threads against a table or replication group, using the GG Parameters page in ORMBA Administration UI.
- You can configure up to 25 common tables (configuration tables where updates are less frequent) under a replication group. However, in case of heavy-duty tables, do not group more than 2 or 3 under a replication group.
- Check and ensure that all consistency checks are disabled, and turn them ON only for testing.
- Make use of the three levels of scalability offered by the ORMBA system. For example, if the volume range is above 10 million and you need the upload to finish in 10 to 15 minutes, we recommend you to use ALL three levels of scalability. However, ensure that the recommended infrastructure facilities (No of processors) are available.
- Turn OFF un-wanted subject areas and disable the corresponding jobs.
- Prior to configuring the job schedules, consider the following aspects:
 - While deciding the jobs to be run in parallel, consider the data handled by each job to ensure that the load is balanced.
 - While scheduling the job, consider how frequently data arrives from the source and how often we require the data in the warehouse.
 - We recommend frequent cycle of job with lesser data, when compared to less-frequent jobs with huge data.
- Purge data from replication layer as early as possible.
- After a load cycle (that is, after all jobs are run once), gather index and tables statistics, including partitioning for optimizer. You can decide if you need to alter the parallelism and if yes, make changes wherever required.
- Create new MViews or utilize the existing ones wherever possible, so that the dashboards query from MViews; rather than hitting the underlying fact tables.

- As far as possible, make use of the ODI Knowledge Modules delivered as part of the product for new or custom facts, dimensions and MViews, as they are specially designed for optimal performance.
- Configure the number of ODI threads as multiples of two.
- Ensure that all tables are composite partitioned and ETL SQLs are partition-aware, so that they can leverage the partitioning. While adding tables, ensure that they are also composite partitioned.
- Recommended thumb rule for deciding the number of ODI threads and DB threads are:

$$\text{DBI Parallel Threads} * \text{No of ODI Threads} = \text{Total No of Threads Available}$$
- The default INITRANS for dimensions is set as '4' and for fact is set as '16'. If the number of threads concurrently updating the tables is more than the default values and the ITL Waits are significant, you can change the parameters as required.
- Storage size of the table spaces of the Replication schema and Warehouse should be greater than or equal to the storage size of table spaces in the source schema.
- Use the DBMS Data Pump feature for initial load of data from the Source System to ORMBA Replication Layer. This ensures high-speed transfer of data from one database to other.
- You can improve the performance of Data Pump feature by enabling parallelism. To do this, ensure that the source tables are partitioned and the CPU count is more than one.

8.2 What are the different levels of scalability?

Due to the huge data being handled by the system, we offer three levels of scalability:

1. **Multiple parallel instances:** You can configure this using the 'Maximum Parallel Count' field in the Target Entity Definition page of the ORMBA Administration UI.
2. **Multiple ODI threads within an instance:** You can configure each job instance of an ETL job to run in multiple threads. To configure the number of threads, use the field 'Thread Count' in the Target Entity Definition page of ORMBA Administration UI.
3. **Multiple DB threads within an ODI thread:** You can configure the database statements in an ETL job thread to execute in parallel database threads. You can enable this by configuring database hints in the replication views. The default degree of parallelism is 8. However, if CPUs are not available, you can change the hint.

8.3 How do I utilize Application level parallelism?

Where: ORMBA Administration UI

We offer two types of application-level parallelism:

1. **Multiple parallel instances:**

You can configure the system to run multiple parallel instances of a job. Prior to deciding the count of parallel instances, you need to consider the CPU capacity and the expected volume of data. Based on this count, system divides the total data into equal slices.

To configure the number of parallel instances at which the ETL job should run, use the 'Maximum Parallel Count' field in the Target Entity Definition page.

Scheduler Configuration

Start Date (yyyy-MM-dd HH:mm:ss)	2015-09-18 00:00:00	Maxi
Schedule Type	* Incremental Load	Retry
Slice Duration Frequency	* Day(s)	Anal
Slice Duration	* 1	Daily
Maximum Parallel Count	* 1	

Figure 28: Configuring Multiple Parallel Instances of a Job

2. Multiple ODI threads within an instance:

You can configure each of the job instances to run in multiple ODI threads. Each ODI thread works on a different slice of data. We recommend you to configure the number of ODI threads as multiples of two.

To configure the number of ODI threads to be spawned, use the 'Thread Count' field available in the Target Entity Definition page.

Identifier for parent child relationship in Customer Hierarchy	CUST_HIER_FLAG	PARENT
Number of threads for ODI load job	ODI_NUM_THREADS	4
Date from which all ETL jobs for the specific product instance will be configured to end the initial ...	BM_EXTRACT_END_DT	01-Jan-2015

Figure 29: Configuring Multiple ODI Threads within an Instance

8.4 How do I utilize DB level parallelism?

Where: ODI

To utilize database level parallelism, configure database Hints in the required replication views.

By default, database parallelism is enabled for the following heavy-duty facts:

- BF_PAY_TNDR
- BF_PAY_EXCP
- BF_FEED
- BF_TXN
- BF_TXN_ERROR
- BF_TD
- BF_SA
- BF_ADJUSTMENT
- BF_PAYMENT
- BF_GL
- BF_BILL
- BF_PRICE

8.5 How do I enable parallelism for a fact/dimension?

You can enable parallel processing for facts or dimensions. To do this, create a package for parallel processing following the procedure below:

1. Bring driving tables `slicing_ts` and `slicing_group` into the source view and add the below filter condition in ODI:

```
Default.BF_GL_VW.JRN_SLICING_GRP IN
(#ORMB_BI.SLICING_GRP_VALUES) AND
Default.BF_GL_VW.JRN_SLICING_TS BETWEEN TO_DATE
(#ORMB_BI.BM_SLICE_BEG_TS, 'YYYYMMDD') and TO_DATE
(#ORMB_BI.BM_SLICE_END_TS, 'YYYYMMDD')
```

where `ORMB_BI.SLICING_GRP_VALUES` is a project variable.

2. Create a copy of the available parallel processing package and change the child scenario to your desired scenario.
3. Change the execution mode of the child scenario to 'Asynchronous' from the Properties tab of the same.
4. Select the Command tab from the Properties window and pass the additional variables that are used in the child scenario.

For example,

```
OdiStartScen    "-SCEN_NAME=SCBL_BF_PKG_GL"
                "-SCEN_VERSION=001"
                "-SYNC_MODE=2"
                "-ORMB_BI.SLICING_GRP_VALUES=#ORMB_BI.SLICING_GRP_VALUES"
                "-ORMB_BI.BM_JOB_ID=#ORMB_BI.BM_JOB_ID"
                "-ORMB_BI.BM_SLICE_BEG_TS=#ORMB_BI.BM_SLICE_BEG_TS"
                "-ORMB_BI.BM_SLICE_END_TS=#ORMB_BI.BM_SLICE_END_TS"
                "-ORMB_BI.BM_LOAD_TYPE=#ORMB_BI.BM_LOAD_TYPE"
```

5. Now, regenerate the parent scenario and run the same. This will run multiple child scenarios in parallel depending on the variable values.

Note: Transaction related facts are not supported for running in parallel.

8.6 What are the recommended DB settings?

The parameters in the initialization file acts as a starting point for database tuning. We have listed the recommended values for the parameters in the table below. However, the actual or optimal value differ from one installation to another.

Parameter	Recommended Value
MEMORY_MAX_TARGET	50 to 60% of the total available RAM on the node.

Parameter	Recommended Value
MEMORY_TARGET	Value should be less than or equal to the value set for the MEMORY_MAX_TARGET parameter and at the same time it should be greater than or equal to the sum of SGA_TARGET and PGA_AGGREGATE_TARGET .
SGA_TARGET	50 to 70% of the value defined for the MEMORY_TARGET parameter.
SGA_MAX_SIZE	Minimum 8 GB.
DB_CACHE_SIZE	Minimum 4 GB.
PGA_AGGREGATE_TARGET	Minimum 4 GB.
STATISTICS_LEVEL	TYPICAL
DB_BLOCK_SIZE	8 K or 16 K.
LOG_CHECKPOINT_INTERVAL	0
LOG_CHECKPOINT_TIMEOUT	0
OPTIMIZER_SECURE_VIEW_MERGING	TRUE
OPTIMIZER_MODE	ALL_ROWS
OPEN_CURSORS	3000
PROCESSES	500 - 1000
SESSION_CACHED_CURSORS	500 or higher
PARALLEL_FORCE_LOCAL	TRUE
PARALLEL_DEGREE_POLICY	AUTO
PARALLEL_ADAPTIVE_MULTI_USER	TRUE
PARALLEL_MAX_SERVERS	4 or 8
DB_FILE_MULTIBLOCK_READ_COUNT	32
CURSOR_SPACE_FOR_TIME	TRUE
DISK_ASYNC_IO	TRUE

Parameter	Recommended Value
QUERY_REWRITE_ENABLED	FALSE
RECYLCEBIN	OFF

8.7 How do I perform DB tuning for performance?

To improve the overall performance of ORMBA, we recommend you to tune the database. The database tuning includes the following tasks:

- **Set optimal values for database initialization parameters**

For each installation of ORMBA, it is important to set optimal values for database initialization parameters in the initialization file. The optimal values vary from installation to installation.

- **Set appropriate storage size for table spaces**

Initial storage size of default table spaces of Replication schema and Warehouse should be greater than or equal to the storage size of table spaces in Source schema.

- **Set INITTRANS value for Dimensions and Facts**

The default INITTRANS value for Dimensions is set as '4' and for Facts it is set as '16'. When thread count of jobs for a dimension or a fact is higher than the default INITTRANS setting, then the INITTRANS for the corresponding fact or dimension and its indexes should be set accordingly.

- **Review AWR and ASH reports**

Monitor the database performance by periodically reviewing the AWR and ASH reports. These reports are designed to assist you with database tuning.

- **Partition the database tables**

All warehouse tables are partitioned by range, and each partition is further sub-partitioned by hash. The range partition is based on the column ILM_DTM and the hash sub-partition is based on columns that are used frequently, typically division or customer segment. In case of customization, ensure that you follow the same standards.

8.8 What is the average lag while replicating data?

ORMBA replicates the source system changes at **near-real-time** basis. The typical lag is in order of seconds, between 10 to 100 seconds. However, the exact lag also depends on the infrastructure (No of CPUs) used and the data being handled.

You can find out the replication lag using GG Monitor and to reduce the lag, you can either alter the GG Parameters of the table, or move it to a separate replication group. To know more about this, see section [6.9](#).

9. Archival

You can configure ORMBA data archival in three layers:

- Replication Schema
- Staging Schema
- Data Warehouse

To know more about archival in each of these layers, see the sections below.

9.1 How do I purge data in the replication schema?

Each source table that has been marked for replication has a replica table created in the replication schema. Thus, a replication schema contains a replica of source tables (along with history) that corresponds to dimensions and facts. Over a period, the volume of data in the replication area can grow huge and this could result in operational overhead and performance issues.

All product tables in the replication schema have day-wise partition. As per your business requirements, you can decide and configure how long (in days) you want to retain data in the replication schema. Based on this configuration, the corresponding day-wise partitions are dropped and corresponding (global) indexes are automatically updated. Dropping day-wise partitions improves the performance of the jobs.

Note: System retains the first partition as reference, and hence does not drop this partition. Instead, system truncates the data in this partition.

You can configure the number of days up to which system has to retain data using the Source Table Definition page of ORMBA Administration UI. By default, system retains data for 10 days. You can also exclude a table from purging using the Source Table Definition page. To do this, you can either set the Purge Flag to 'N' or configure Replication Retention Days as '0'.

You can manually execute the ODI procedure `BM_PURGE_REPLICATION` when you would like to purge data from replication schema; or you can schedule it based on your requirements. We recommend you to schedule this at off-peak hours, as the heavy load at database could possibly have an impact on system performance.

If any of the tables in the replication schema are not partitioned, then the data is deleted from the tables in batches of 5000 rows. However, the delete operation has a significant impact on the performance and you will see that it takes longer. Moreover, deletion of data results in fragmentation of tables and you have to de-fragment the tables manually. Hence, we recommend you to partition the custom tables based on the field `JRN_SLICING_TS`.

If there are multiple target entities dependent on the same source table, then system checks the time stamp on all entities. Only when the latest execution date is older than the retention period it drops the table.

Note: We recommend you to retain data in the replication layer for only as long as it is necessary.

9.2 How do I purge data from the staging schema?

Where: ORMBA Administration UI

Data in the staging schema is partitioned job-wise and the partitions are dropped when the job execution date exceeds the configured retention period. Purging of data from staging schema is handled differently for facts and dimensions:

- Facts: There is no retention and every time data is loaded, the tables are dropped and then re-created.
- Dimensions: Type 2 data cumulates until the 'Stage Retention Period' is reached and then the day-wise partitions are dropped.

You can configure the stage retention period (in days) using the Target Entity Definition page. While running a job for a dimension, system checks if the retention period is reached, and drops partitions automatically.

By default, system retains data in staging schema for **7** days. If you configure '0' in the Stage Retention (Days) field, then the corresponding target entity is never purged.

Target Entity Definition

Entity Name	BF_ADJUSTMENT
Entity Type	Facts
Entity Sub Type	Regular Fact
Subject Area	Financial Transaction ▼
ODI Package Name	* BF_PKG_ADJUSTMENT
Stage Retention (Days)	7
Warehouse Retention (Months)	24
User Extension Procedure	<input type="text"/>
User Extension Procedure (Post Job)	<input type="text"/>
Characteristic Entity	<input type="text"/> ▼

Figure 30: Configuring Stage Retention Period

9.3 How do I archive data from the warehouse?

Where: ORMBA Administration UI

Due to legal reasons, a warehouse normally retains data for a period of seven years. However, in ORMBA you can set a period for date retention in the warehouse, after which the data is purged.

To define the data retention period, navigate to the Target Entity Definition page in Administration UI and configure the **Warehouse Retention (Months)** field to a desired value. The default value of the field is **0**, in which case the data is never purged from warehouse. System drops partitions based on the value in ILM_DTM column and the retention period.

Target Entity Definition

Entity Name	BD_ACCT
Entity Type	Dimensions
Entity Sub Type	Slowly changing dimension Type II
Subject Area	<input type="text"/>
ODI Package Name	* BD_PKG_ACCT
Stage Retention (Days)	7
Warehouse Retention (Months)	0
User Extension Procedure	<input type="text"/>
User Extension Procedure (Post Job)	<input type="text"/>
Characteristic Entity	ACCT <input type="text"/>

Figure 31: Configuring Warehouse Retention Period