

Oracle Hyperion Profitability and Cost Management Configuration Guidelines for Detailed Profitability Applications

November 2015

Contents

About these Guidelines	1
Setup and Configuration Guidelines	2
Database Edition Guidelines	2
Database Configuration Guidelines.....	2
Schema Creation Guidelines	2
Database Configuration Options	2
Model Data Schema Guidelines	3
Performance Considerations	5
Tables and Indexes	5
Server Sizing Guidelines	5
Concurrency Guidelines	5
Model Design Considerations	6
Performance Optimization	7
Optimized Dimension Member Retrieval.....	7
Bulk Same As Source Performance Optimization	8
Using Detailed Profitability System Reports	8

About these Guidelines

Oracle Hyperion Profitability and Cost Management Configuration Guidelines for Detailed Profitability Applications outlines best-practice guidelines for configuring Oracle Database for use with Oracle Hyperion Profitability and Cost Management Detailed Profitability applications.

Note: Unless otherwise specified, all parts of this document apply to Detailed Profitability applications in Oracle Hyperion Profitability and Cost Management, release 11.1.2.3.000 and later.

During model calculations, SQL statements are submitted to the database to produce the model calculation results. The configuration and setup of Oracle Database significantly affects and determines the efficiency of SQL statements performing the model calculations.

Important! If the scale of your data requires a Very Large Database (VLDB), review these guidelines before installing Oracle Database. If your data does not require a VLDB, review these guidelines before installing Profitability and Cost Management.

Setup and Configuration Guidelines

Database Edition Guidelines

While any database platform described in the Oracle Enterprise Performance Management System Certification Matrix (MTAC) is supported, the Parallel SQL Execution feature for calculation rules in a Detailed Profitability application requires Oracle Database Enterprise Edition.

Database Configuration Guidelines

Redo Log Recommendation

Profitability and Cost Management Detailed Profitability applications perform a high volume of updates upon the destination table. Oracle recommends the creation of three redo log groups. Each log should be 1 GB in size.

Disable ARCHIVELOG Mode

ARCHIVELOG mode is required for advanced disaster recovery such as online backup, restore to point in time, Data Guard and other features of Oracle Database. It is not unusual for ARCHIVELOG to use 30% of available resources and significantly affect the performance of Detailed Profitability applications. A database administrator (DBA) evaluating the tradeoff between the loss of data and performance should help determine whether the ARCHIVELOG mode is enabled.

Database Parameters

Parameter	Value
CURSOR_SHARING	EXACT
OPEN_CURSORS	5000
PROCESSES	1000
SESSIONS	2000
SESSION_CACHED_CURSORS	200

Schema Creation Guidelines

For greatest efficiency:

- The Product Schema and the Model Data Schema must reside in the same database.
- The Product Schema must meet the requirements defined in the *Oracle Hyperion Profitability and Cost Management Administrator’s Guide*, especially the permissions granted to the user.
- The Model Schema must meet the requirements defined in the *Oracle Hyperion Profitability and Cost Management Administrator’s Guide*, especially the permissions granted to the user.

Database Configuration Options

The design of the model calculations for Detailed Profitability applications seeks to leverage the powerful features and performance of the Enterprise Edition of Oracle Database. Oracle recommends customer-focused performance tuning activities for Detailed Profitability applications in Oracle Database.

Low Effort Configuration

Oracle Enterprise Performance Management System Standard Deployment Guide describes the default installation and configuration for the EPM suite. This default installation is the Low Effort Configuration.

High Performance Configuration

The High Performance Configuration requires the installation of optional features for Oracle Database by the DBA. The two optional features are grid infrastructure and Automatic Storage Management.

Configuration Comparison

Configuration Method	IO Expectations (requirements)	Advantages	Disadvantages
Low Performance Configuration	<ul style="list-style-type: none"> SQL from concurrent detailed applications without contention 	<ul style="list-style-type: none"> Lower cost for initial setup of Oracle Database Less Oracle Database Administration expertise required Appropriate for model design and initial proof of concept (POC) 	<ul style="list-style-type: none"> More effort required to use all cores available to Oracle Database upon the server All segregation and load balancing of IO activities are performed manually by the DBA
High Performance Configuration	<ul style="list-style-type: none"> Degree of parallelism without contention SQL from concurrent detailed applications without contention 	<ul style="list-style-type: none"> Ability to effectively use all the cores available to Oracle Database upon the server Ability to easily segregate recovery-activity-related IO (redo logs for instance) from database file-access-related IO Automatic management of data files and IO load balancing 	<ul style="list-style-type: none"> Requires installation of advanced features of Oracle Database Requires IO tuning and optimization Requires managing resources and tuning of parallel server Performance POC to identify the correct configuration recommended

Model Data Schema Guidelines

Oracle recommends following the model data schema design rules and best practices described in this section. This section restates and extends information from the “Model Data Schema Table Structural Requirements” section of the *Oracle Hyperion Profitability and Cost Management Administrator’s Guide* for Release 11.1.2.3.000 and later.

Schema Design Rules

Model data schema design rules are as follows:

- The system supports a star schema setup where destination lookup tables are joined directly to the destination stage table.
- The system does not support access to snowflake lookup tables that are not joined directly to the destination stage table.
- The columns associated with the target measure of a driver artifact must reside upon the stage destination table.
- The columns associated with Oracle Enterprise Performance Management Architect (EPMA) dimensions must store a value identical to the name of an EPMA dimension member.
- The column type and size of a column associated with an EPMA dimension or an EPMA POV dimension must be:
 - A Varchar2 (80 Char) for Oracle Database
 - A Nvarchar (80 Char) for Microsoft SQL Server
- The join criteria for retrieving a value from a destination lookup table must return only one value per row in the destination stage table.
- The dimension member name from the lookup table must be identical to the name from the destination stage table for the driver value to be returned from the lookup table.
- Columns associated with the EPMA POV dimensions must reside on the destination stage table for model calculations to be successful.
- Lookup tables may have columns associated with the EPMA POV dimensions to support POV specific driver values.

- The destination stage table must have a constraint defined on it to enforce the uniqueness of the table. We recommend including a unique identifier and the EPMA POV dimension to allow multiple POVs to be stored in the table.
- A number column named 'WORKING' must exist in the destination stage table.

Schema Rules for Oracle Database

- Database object names must be in upper case characters, may include digits 0-9, and may include characters “_” and “\$”. Lower case and other special characters are not supported. If you must enclose the table or column name in double quotes in the CREATE TABLE command for the identifier to be accepted, then it is not supported for use in Profitability and Cost Management.

Data Quality

- Values in a column registered as EPMA dimensions must be a case sensitive match to the dimension member name for an assignment rule to select the row as part of the destination.

Performance Considerations

Tables and Indexes

Considerations for tables and indexes are as follows:

- An assignment of results may target a small portion of a very large destination table. Therefore, efficient indexes are critical to prevent unnecessary full table scans.
- Indexes defined for the destination table have a significant impact on performance because each row in the destination table likely will be updated multiple times.
- There must be a primary key defined upon the destination table to ensure that each row may be uniquely identified when matching calculated driver values to destination rows.
- When the destination table contains data from multiple points of view, you must include the columns associated with the EPMA POV dimension as part of the primary key.
- A non-unique index upon the destination table improves performance when the index includes a column associated with an EPMA dimension where one or more destination assignment rules include a member with a generation of 2 or greater.
- At least one non-unique index with at least one column associated with an EPMA dimension should be defined to limit full table scans.
- A column associated with an EPMA dimension where the number of unique values exceeds 50 (50 dimension member values with data) is a good candidate column for a non-unique index.
- If there are no good candidate columns (fewer than 50 values), then a non-unique index with multiple columns associated with an EPMA dimension may be required.
- In addition to these guidelines, consider using the SQL Tuning Advisor to provide feedback for index recommendations for Oracle Database.

Server Sizing Guidelines

Server sizing guidelines are as follows:

- The server memory should be sufficient to minimize accessing the physical disk for queries that join the destination lookup tables and the destination table. It is helpful when there is sufficient memory to cache all the row sets returned from all of the tables accessed by the query, but the size of the destination table may prevent this ideal situation.
- Since the table that stores the calculated driver values joins to the destination table, the server memory should be sufficient to support sorting two times the destination rows defined by the destination assignment rule.

Concurrency Guidelines

Concurrency between Applications

Guidelines are as follows:

- The server processing resources (CPU resources) should be sufficient to support all concurrent POV model calculations (all Detailed Profitability applications running model calculations).
- The queries and DML operation issued by the Detailed Profitability applications depend upon efficient IO resources and effective IO calibration by your DBA.
- The server IO resources should be sufficient to support queries and DML operations for the all concurrent POV model calculations (all Detailed Profitability applications running model calculations).

Oracle Parallel Server

Guidelines are as follows:

- Parallel Auto SQL Execution uses all parallel server resources allowed for the user as defined for the Oracle Database configuration. This may degrade the performance of other activities, such as reporting, that occur in the same database.
- Parallel User Specified Execution uses parallel servers equal to two times the value specified by the user.
- The Oracle Database Resource Manager described in the *Oracle Database Administrator's Guide* provides a mechanism for managing multiple workloads that are contending for system and database resources.
- Oracle Parallel Server writes directly to the physical disks to achieve the best performance. IO Contention will occur if Parallel User Specified Setting or Parallel Auto Setting exceeds the number of direct writes that may be supported simultaneously.

Model Design Considerations

Source Considerations

Considerations are as follows:

- The number of source pools assigned to the destination has a direct relationship with the amount of time to perform model calculations.
- Reducing the number of source pools assigned reduces the amount of time spent performing model calculations.
- Sources that share a driver definition and share destination assignment rules may be combined without affecting results.
- Vertical source stage tables store one source per row.
- Horizontal source stage tables store multiple sources per row.

Destination Considerations

Considerations are as follows:

- The number of rows updated in the destination table by an assignment to the destinations affects the performance of model calculations.
- When a large number of destinations are targeted by an assignment, a higher percentage of time for model calculations is spent in the relational database management system (RDBMS).
- When a small number of destinations are targeted by an assignment, a significantly reduced percentage of time for model calculations is spent in the RDBMS.

Calculation Rule Considerations

Considerations are as follows:

- Calculated Measure rules make a single pass to update the destinations.
- Single-source rules and multi-source rules make two passes on the destination for each assigned source pool, a first pass to calculate driver measures and a second pass to update the destinations.
- Single-source calculation rules assign one source cost pool identified by the level zero source member combination to the destinations identified by the destination assignment rule.
- Multi-source calculation rules without source-side driver values or same-as-source destination rules assign one aggregate source pool to destinations.
- Multi-source calculation rules with same-as-source destination rules assign one aggregate source pool for each combination of member values from the same-as-source dimensions to the destinations.
- Multi-source calculation rules with source-side driver values assign one source pool for each source member value combination to the destinations.

Driver Definition Considerations

The driver value is calculated from the driver formula once for each row in the destination table, so custom SQL where elapsed time is short may become a large cost when it is issued a million times for a million-row table.

Performance Optimization

Detailed Profitability performance is significantly affected by time taken for member traversal. An enhancement is included to make certain styles of allocation faster. Performance optimization is activated in a user-transparent fashion, so users are not required to do anything special other than making sure hierarchy tables are generated before running certain allocations. Users can modify the allocation logic, if possible, to take advantage of optimization.

Optimized Dimension Member Retrieval

Adding the ability to retrieve dimension members from dimension hierarchy tables instead of the dimension catalog reduces the elapsed time for qualifying calculations.

Retrieving Dimension Members from Dimension Catalog

The table below is a logical representation of some of the dimensional information from the Dimension Catalog. All of the dimension members across all of the dimensions from each application are stored in the Dimension Catalog. The time required to retrieve a set of dimension members increases as the number of applications increases, as the number of dimensions increases, and as the number of dimension members increases.

Application	Dimension	Dimension Member
App01	Cost Center	Marketing
App02	Product Line	Vans
App01	Product	Router
App03	Region	West Coast
App01	Channel	Online

Retrieving Dimension Members from Dimension Hierarchy Tables

The table below is a logical representation of some of the dimensional information from the Cost Center hierarchy table for the Detailed Profitability sample model. Since each hierarchy table stores the dimension members from a single dimension from one application, the time required to retrieve a set of dimension members is dependent on the number of Cost Center dimension members. This retrieval method is not affected by adding more applications or more dimensions to the system.

Level2	Level1	Level0
AllCostCenters	NoCostCenter	NoCostCenter
AllCostCenters	CC1001	CC1001
AllCostCenters	CC1010	CC1010
AllCostCenters	CC2001	CC2001
AllCostCenters	CC3000	CC3100
AllCostCenters	CC3000	CC3200
AllCostCenters	CC4000	CC4100
AllCostCenters	CC4000	CC4200

Calculation Rule Qualification Table

The rows in the table below represent attributes of Calculation Rules. If a Calculation Rule contains an attribute where the background is lightly shaded (cells labeled “Dimension Catalog Member Selection”), then the rule does not qualify for using Hierarchy Table Member Selection.

Attributes	Calculated Measure	Rate	Ratio
Name Filter	Dimension Catalog Member Selection	Dimension Catalog Member Selection	Dimension Catalog Member Selection
Attribute Filter	Dimension Catalog Member Selection	Dimension Catalog Member Selection	Dimension Catalog Member Selection
UDA Filter	Dimension Catalog Member Selection	Dimension Catalog Member Selection	Dimension Catalog Member Selection
Dest Assignment Rule Non-SAS	Hierarchy Table Member Selection	Hierarchy Table Member Selection	Hierarchy Table Member Selection
Dest Assignment Rule SAS	Hierarchy Table Member Selection	Hierarchy Table Member Selection	Hierarchy Table Member Selection
Dimensions Registered on Stage Table only	Hierarchy Table Member Selection	Hierarchy Table Member Selection	Hierarchy Table Member Selection
Dimension Registered on Lookup Table	Dimension Catalog Member Selection	Dimension Catalog Member Selection	Dimension Catalog Member Selection

Bulk Same As Source Performance Optimization

The goal of Bulk Same As Source Performance Optimization is to perform all calculations of results for the multi-source calculation rule within a single update of the destination stage table. The system chooses the Bulk Same As Source Performance Optimization when the following requirements are met:

- The calculation rule type is Multi-source.
- The business dimensions set as Same As Source within the destination assignment rule are registered directly upon a stage destination table (not registered upon a destination lookup table).
- The database platform is a supported version of the Oracle database.
- Cross-dimensional Same As Source is not used within the destination assignment rule.

Note: The Bulk Same As Source Performance Optimization was introduced in Oracle Hyperion Profitability and Cost Management versions 11.1.2.4.110 and 11.1.2.3.701.

Using Detailed Profitability System Reports

The system reports delivered with HPCM 11.1.2.4.00 facilitate tuning performance.

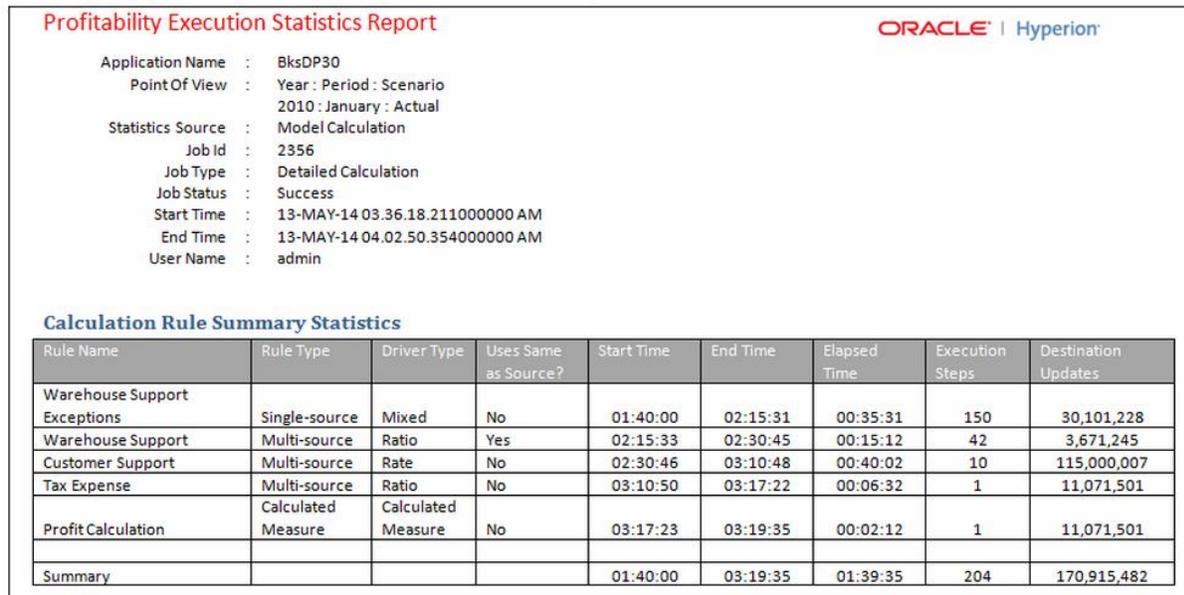
Executions Statistics Report Usage

You can use the information within the Execution Statistics Report to improve performance in the following ways:

- Identify the calculation rules with the largest elapse time. Modifications to long running calculation rules are the best opportunities for improving performance.

- Each execution step performed by a calculation rule issues a set of SQL statements. After identifying the calculation rules with the most execution steps, make changes to the calculation rules to reduce the number of steps.
 - Since single-source calculation rules have one execution step for each source, limit the total number of sources allocated within your model by all of single-source rules (use single-source rules only when necessary).
 - Identify rows where the Rule Type is Multi-source, the Driver type is Ratio, the Use Same as Source column value is Yes, and the number of execution steps are greater than 1. Same as Source rules that are using the Bulk Same as Source optimization will have a single execution step. Same as Source rules that cannot use the bulk Same as Source optimization will have multiple execution step. See example below for more information about tuning multiple execution steps.
- When a calculation rule updates more than 50,000 rows with per execution step, then altering the Rule SQL Execution Mode may reduce the elapse time. See step 4 within the Adding Calculation Rules section of the User Guide for more details about setting the Rule SQL Execution Mode.
 - When there is only a single execution step that update more than 50,000 rows, the chances for improved performance are high.
 - When at least 70-80% of the execution steps associated with a Multi-source rule update 50,000 rows or more, then the chances are good that altering the SQL Execution Mode improves performance. See the “Gathering Individual Execution Step Details” section below for more information.

Figure 11. Example of a Detailed Profitability Execution Statistics Report



Same As Source Multiple Execution Step Example

Same As Source Multi-source calculation rules that cannot use the Same As Source Bulk Optimization have multiple execution steps. An execution step exists for each unique intersection of level 0 members from the same as source dimensions. For example, the organization dimension and the products dimension are two destination stage dimensions. There are 6 level 0 members from the organization dimension and 7 level 0 members from the products dimension chosen by the source assignment rule. The calculation rule in this example would have 42 execution steps.

See the “Bulk Same As Source Performance Optimization” section above to determine why the Multi-source rule does not benefit from the Bulk Same As Source performance optimization. One of the four requirements is not met by a Same As Source calculation rule that does not use the Bulk Same As Source Performance Optimization.

Gathering Individual Execution Step Details

The POV Temporary Object Details View contains information about the temporary objects created by detailed model calculations within the database. An image of data returned by this view, named HPM_EXPDP_OV_TEMP_OBJECT, is shown below. This view is available within the Hyperion Profitability and Cost Management product schema after applying the 11.1.2.4.110 PSU or applying the 11.1.2.3.701 PSU. Detailed model calculations create a table that provides detailed information about the execution steps for the specified POV. To find this temporary object for a specified POV, find the row within the HPM_EXPDP_OV_TEMP_OBJECT view where the Object Usage Type is Allocation Execution Steps and the name of the temporary object follows the HPMD_<app_name>_ALLOC_<seq#> where POV column values match the member names for the specified POV. Once you have identified the correct Allocation Execution Steps table, query the table to obtain the number of rows updated upon the destination table by each execution step. This information is useful when determining if 70-80% of the execution steps update 50,000 rows or more as described above.

	APPLICA...	POV_DIM1...	POV_DIM2...	POV_DIM3...	POV_D...	OBJECT_NAME	OBJECT_TYPE	OBJECT_USAGE_TYPE
1	BksDP31	2012	January	Actual	(null)	HPMD_BKSDP31_CDV_4642	TABLE	Calculated Driver Values
2	BksDP31	2012	January	Actual	(null)	HPMD_BKSDP31_CLC_V_4643	VIEW	Calculation View
3	BksDP31	2012	January	Actual	(null)	HPMD_BKSDP31_CDV_4640	TABLE	Calculated Driver Values
4	BksDP31	2012	January	Actual	(null)	HPMD_BKSDP31_CLC_V_4641	VIEW	Calculation View
5	BksDP31	2012	January	Actual	(null)	HPMD_BKSDP31_CDV_4638	TABLE	Calculated Driver Values
6	BksDP31	2012	January	Actual	(null)	HPMD_BKSDP31_CLC_V_4639	VIEW	Calculation View
7	BksDP31	2012	January	Actual	(null)	HPMD_BKSDP31_CDV_4636	TABLE	Calculated Driver Values
8	BksDP31	2012	January	Actual	(null)	HPMD_BKSDP31_CLC_V_4637	VIEW	Calculation View
9	BksDP31	2012	January	Actual	(null)	HPMD_BKSDP31_CDV_4634	TABLE	Calculated Driver Values
10	BksDP31	2012	January	Actual	(null)	HPMD_BKSDP31_CLC_V_4635	VIEW	Calculation View
11	BksDP31	2012	January	Actual	(null)	HPMD_BKSDP31_CDV_4632	TABLE	Calculated Driver Values
12	BksDP31	2012	January	Actual	(null)	HPMD_BKSDP31_CLC_V_4633	VIEW	Calculation View
13	BksDP31	2012	January	Actual	(null)	HPMD_BKSDP31_CDV_4630	TABLE	Calculated Driver Values
14	BksDP31	2012	January	Actual	(null)	HPMD_BKSDP31_CLC_V_4631	VIEW	Calculation View
15	BksDP31	2012	January	Actual	(null)	HPMD_BKSDP31_CLCRL_4629	TABLE	Calculation Rule Source Dimension Expansio
16	BksDP31	2012	January	Actual	(null)	HPMD_BKSDP31_ALLOC_4628	TABLE	Allocation Execution Steps
17	BksDP31	2012	January	Actual	(null)	HPMD_BKSDP31_ALC_V_4627	VIEW	Allocation View
18	BksDP31	2012	January	Actual	(null)	HPMD_BKSDP31_IDV_4625	TABLE	Source Pools and Total Driver Values
19	BksDP30	2012	January	Actual	(null)	HPMD_BKSDP30_CLCRL_4685	TABLE	Calculation Rule Source Dimension Expansio
20	BksDP30	2012	January	Actual	(null)	HPMD_BKSDP30_ALLOC_4684	TABLE	Allocation Execution Steps
21	BksDP30	2012	January	Actual	(null)	HPMD_BKSDP30_ALC_V_4683	VIEW	Allocation View
22	BksDP30	2012	January	Actual	(null)	HPMD_BKSDP30_IDV_4681	TABLE	Source Pools and Total Driver Values
23	cgp2	2010	January	Actual	(null)	HPMD_CGP2_CLCRL_4665	TABLE	Calculation Rule Source Dimension Expansio
24	cgp2	2010	January	Actual	(null)	HPMD_CGP2_ALLOC_4664	TABLE	Allocation Execution Steps
25	cgp2	2010	January	Actual	(null)	HPMD_CGP2_ALC_V_4663	VIEW	Allocation View
26	cgp2	2010	January	Actual	(null)	HPMD_CGP2_IDV_4661	TABLE	Source Pools and Total Driver Values