

Oracle® Retail Predictive Application Server

Administration Guide for the Classic Client

Release 14.1

E59120-01

December 2014

Primary Author: Judith Meskill

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

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

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

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" 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, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

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

Oracle and Java are registered trademarks of Oracle 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.

This software or hardware and documentation may provide access to or information about 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, except as set forth in an applicable agreement between you and Oracle.

Value-Added Reseller (VAR) Language

Oracle Retail VAR Applications

The following restrictions and provisions only apply to the programs referred to in this section and licensed to you. You acknowledge that the programs may contain third party software (VAR applications) licensed to Oracle. Depending upon your product and its version number, the VAR applications may include:

- (i) the **MicroStrategy** Components developed and licensed by MicroStrategy Services Corporation (MicroStrategy) of McLean, Virginia to Oracle and imbedded in the MicroStrategy for Oracle Retail Data Warehouse and MicroStrategy for Oracle Retail Planning & Optimization applications.
- (ii) the **Wavelink** component developed and licensed by Wavelink Corporation (Wavelink) of Kirkland, Washington, to Oracle and imbedded in Oracle Retail Mobile Store Inventory Management.
- (iii) the software component known as **Access Via**™ licensed by Access Via of Seattle, Washington, and imbedded in Oracle Retail Signs and Oracle Retail Labels and Tags.
- (iv) the software component known as **Adobe Flex**™ licensed by Adobe Systems Incorporated of San Jose, California, and imbedded in Oracle Retail Promotion Planning & Optimization application.

You acknowledge and confirm that Oracle grants you use of only the object code of the VAR Applications. Oracle will not deliver source code to the VAR Applications to you. Notwithstanding any other term or condition of the agreement and this ordering document, you shall not cause or permit alteration of any VAR Applications. For purposes of this section, "alteration" refers to all alterations, translations, upgrades, enhancements, customizations or modifications of all or any portion of the VAR Applications including all

reconfigurations, reassembly or reverse assembly, re-engineering or reverse engineering and recompilations or reverse compilations of the VAR Applications or any derivatives of the VAR Applications. You acknowledge that it shall be a breach of the agreement to utilize the relationship, and/or confidential information of the VAR Applications for purposes of competitive discovery.

The VAR Applications contain trade secrets of Oracle and Oracle's licensors and Customer shall not attempt, cause, or permit the alteration, decompilation, reverse engineering, disassembly or other reduction of the VAR Applications to a human perceivable form. Oracle reserves the right to replace, with functional equivalent software, any of the VAR Applications in future releases of the applicable program.

Contents

Send Us Your Comments	xxi
Preface	xxiii
1 Introduction	
Administrator Overview	1-1
Basic Concepts of RPAS.....	1-1
Multidimensionality	1-2
Hierarchies	1-2
Measures.....	1-2
Domains and Workbooks	1-2
Measure Data.....	1-3
Administrative Workbooks and Wizards.....	1-4
Hybrid Storage Architecture	1-4
A Note about RPAS Configurability and Extensibility.....	1-4
2 Building and Upgrading Domains	
Building a Domain.....	2-1
Prerequisites.....	2-1
Client-Side Procedures	2-1
Zip the Configuration Project Folder	2-2
Server-Side Procedures	2-2
Unzip the Configuration File	2-2
Verify the Environment Variable Settings	2-3
Get the Input Files Ready	2-3
Building the Domain Manually or Through a Command-Line Interface.....	2-3
rpasInstall	2-3
rpasInstall Usage.....	2-3
Validate Domain Build Results	2-5
After Building the Domain	2-5
Upgrading and Patching Domains.....	2-5
Upgrading From Any Release of RPAS Prior to 13.3	2-6
Domain Conversion.....	2-6
Domain Preparation Before Conversion	2-6
Setting up Domain Conversion Parameters.....	2-9

Running the Domain Conversion.....	2-10
Post-Conversion Processing	2-10
Upgrading From RPAS 13.3 or Later	2-11
Updating Domain Content	2-12
Hierarchies	2-13
Measure Properties	2-13
Rule Groups	2-13
Workbooks	2-13
RPAS 14.0 Cryptography Upgrades.....	2-14
Upgrade Configurations	2-14
Command-Line Support for Configuration Upgrading	2-16
Upgrade Application Configurations	2-16
convertDomain	2-16
convertDomain Usage	2-17
convertDomain -suggestDimBitSize	2-19
Usage for -suggestDimBitSize	2-20
upgradeDomain	2-20
upgradeDomain Usage	2-20
fixDomain	2-21
fixDomain Usage	2-22

3 Domain Administration

Configuring the Classic Client Using eConfigure	3-1
The eConfigure Menu Bar.....	3-1
The eConfigure Main View.....	3-2
Advanced Settings Dialog	3-3
DomainDaemon	3-3
DomainDaemon Usage	3-4
Starting the DomainDaemon.....	3-6
Monitoring the DomainDaemon.....	3-6
Stopping the DomainDaemon.....	3-6
Losing a Client-Server Connection	3-7
Graceful Termination of the Existing Session	3-7
Autosave of Workbooks	3-7

4 SSL

Introduction	4-1
One-Way SSL	4-1
Creating a Self-Signed Root Certificate.....	4-2
Configuring the RPAS Server.....	4-3
Setting Up the Key Store	4-3
Starting the Domain Daemon.....	4-5
Configuring the Classic Client	4-5
Script for Creating Self-Signed Wallets.....	4-7
Usage for createSSLWallets.sh	4-7
User Input	4-7
Output	4-7

Error Handling	4-7
RPAS Web Launch	4-7
SSL Without Authentication	4-8
RPAS Server Configuration	4-8

5 User Maintenance

Access the User Administration Tab	5-1
Add User	5-2
Add User Group	5-4
Delete User	5-5
Delete User Group	5-6
Edit User	5-7
Managing Users Using usermgr	5-9
usermgr Usage	5-10
XML Schema	5-11
Use Cases	5-13
Exporting from an Existing Domain	5-13
Importing into a Domain	5-13
Converting Between XML and Database	5-13

6 System Administration

Auto Workbook Maintenance Wizard	6-2
Add a Workbook to the Auto Build Queue	6-2
Delete a Workbook from the Auto Build Queue	6-3
Edit Workbook Settings	6-4
Workbook Batch Category Management	6-4
Workbook Batch Category Management Wizard	6-4
Adding a Category	6-5
Deleting a Category	6-6
Changing a Category Label	6-7
Hierarchy Maintenance Workbook	6-8
Hierarchy Maintenance Example	6-9
Hierarchy Maintenance Wizard	6-9
Hierarchy Maintenance Worksheet	6-9
Access Hierarchy Maintenance	6-11
Maintain a User-Defined Dimension Within a Hierarchy	6-11
Password Policy Administration Workbook	6-12
Security Administration Workbook	6-14
Security Overview	6-14
User Logon Security	6-14
Measure Level Security	6-15
Position Level Security	6-15
Workbook Security	6-16
Security Administration Workbook	6-17
Workbook Template Rights Worksheet	6-18
Workbook Template Measure Rights Worksheet	6-18

Measure Rights Worksheet.....	6-18
Dimension Modification Rights Worksheet.....	6-19
Position-Level Security Worksheets.....	6-19
Workbook Template Limits Worksheets.....	6-19
Max Domain Session Limit Worksheet	6-20
Max User Session Limit Worksheet	6-20
Using the Security Administration Workbook.....	6-20
Access Security Administration.....	6-20
Set or Modify User Access to Workbook Templates	6-20
Set Measure Availability for Workbook Templates.....	6-20
Assign or Restrict User Access to Measures	6-21
Change User Ability to Modify Dimensions	6-21
Set or Modify Access to Positions.....	6-21
Limit the Number of Workbooks that a User Can Save.....	6-22
Set or Modify the Maximum Domain Session Limit	6-22
Set or Modify the Maximum User Session Limit	6-22
Measure Analysis Workbook.....	6-23
Building a Measure Analysis Workbook.....	6-23
Measure Analysis Worksheet.....	6-24
Review and Edit Sales or Other Registered Measure Data.....	6-24

7 Hierarchy Management

Loading Hierarchies Using loadHier.....	7-1
loadHier Usage.....	7-3
loadHier Notes	7-4
Position Label Translation	7-5
Integer Indexing	7-6
Changing the Bit Size.....	7-8
Reindexing Domains Using reindexDomain.....	7-9
reindexDomain Usage.....	7-10
ReindexDomain Option: Reindex.....	7-10
ReindexDomain Option: Analysis.....	7-11
ReindexDomain Option: Domain Properties.....	7-11
When to Reindex	7-12
Run Reports	7-14
How to Reindex a Dimension	7-15
Reindex Entire Domain as Needed	7-15
Reindex Entire Domain with Force Option.....	7-15
Reindex Dimensions.....	7-16
Prepending Calendar Dimensions	7-16
Reindex Domain Examples.....	7-17
Condition 1: Change of Bit Size	7-17
Condition 2: Using the -force Option.....	7-18
Condition 3: Fragmentation Checks.....	7-18
Optimizing Domains Using optimizeDomain.....	7-19
optimizeDomain Usage.....	7-20
Usage Examples	7-21

Adding New Dimensions to Hierarchies	7-22
Exporting Measure Data	7-23
Exporting Hierarchy Data.....	7-23
Purging Hierarchy Data	7-23
Adding New Dimensions	7-24
Reloading Formal Hierarchy Data.....	7-24
Informalizing DPM Positions.....	7-24
Reloading Measure Data	7-25
Adding New Hierarchy Dimensions Sample Script.....	7-25
Transferring Data	7-27
Masks	7-28
transferData Usage	7-29
Scenarios.....	7-29
FilterHier Utility	7-30
filterHier Usage	7-30
filterHier Notes.....	7-31
Error Conditions.....	7-31
Position Repartitioning.....	7-32
Loading RDF and Curve Parameters after Repartitioning	7-32
Syntax	7-32
Reconfiguring Partitions of a Global Domain Using reconfigGlobalDomainPartitions	7-32
reconfigGlobalDomainPartitions Usage	7-33
Using an Input File.....	7-34
Notes, Assumptions, and Limitations.....	7-35
Renaming Positions Using renamePositions.....	7-36
renamePositions Usage	7-36
Managing Dimension Attributes Using dattrmgr	7-37
dattrmgr Usage.....	7-37
dattrmgr -report Examples	7-38
Setting Properties for Dimensions Using dimensionMgr	7-39
dimensionMgr Usage	7-39
Changing the Label of a Dimension	7-40
Exporting Hierarchy Data Using exportHier	7-40
exportHier Usage	7-40
Informal Position Manager	7-41
InformalPositionMgr Usage	7-41
Basic Operations: Informalize, Formalize, Remove, and Load	7-41
Error Handling.....	7-42
Create Informal Positions in Bulk.....	7-43
Example 1: Creating Informal Positions on the STYL Dimension.....	7-44
Example 2: Creating Informal Positions on the SZ12 Dimension.....	7-45
Position Naming Convention.....	7-46
Data Slice Copying.....	7-46
Merging Informal Positions to Formal Positions	7-48
Managing Position Lists as PQDs Using pqdMgr.....	7-48
pqdMgr Usage	7-49
pqdMgr Notes.....	7-50

8 Data Management

Loading Measure Data Using loadmeasure	8-1
Load File Names and Load Behavior	8-2
Loading Multiple Measures from One File	8-3
CSV Files	8-4
Fixed Width Files	8-4
Loading Data from Below the Base Intersection of the Measures	8-4
Staging Measure Loads	8-4
Running Pre-Load or Post-Load Scripts	8-5
Purging Old Measure Data	8-5
Behavior in an Integrated Environment	8-5
Behavior in a Global Domain Environment	8-5
loadmeasure Usage	8-6
Loading Image Paths for Positions	8-8
Exporting Measure Data Using exportMeasure	8-9
exportMeasure Usage	8-9
Exporting Measure Data Using exportData	8-11
exportData Usage	8-12
The -useLoadFormat Parameter	8-15
Mapping Data Between Domains Using mapData	8-16
mapdata Usage	8-16
Moving Data Between Arrays Using updateArray	8-16
updateArray Usage	8-17
Scan Domain Data Using scanDomain	8-18
scanDomain Usage	8-18
Repair Domain Metadata Using fixDomain	8-19
Sync Measure NA Value with Array NA Value Using syncNAValue	8-20
syncNAValue Usage	8-20

9 Operational Utilities

Finding Alerts Using alertmgr	9-1
alertmgr Usage	9-2
Copying Domains Using copyDomain	9-3
copyDomain Usage	9-3
copyDomain: Format of the XML Configuration File	9-5
Moving a Domain Using moveDomain	9-6
moveDomain Usage	9-6
moveDomain: Format of the XML Configuration File	9-6
Assumptions and Requirements	9-7
Minimum Space Requirement	9-8
Setting Miscellaneous Domain Properties Using domainprop	9-8
domainprop Usage	9-8
Available Properties	9-8
Using the mace Calculation Engine	9-10
mace Usage	9-11
mace Restrictions in Global Domains	9-13
Executing mace in an Integrated Environment	9-14

Managing the Workbook Batch Queue Using wbbatch.....	9-15
wbbatch Usage.....	9-15
Managing Workbooks Using wbmgr.....	9-17
wbmgr Usage.....	9-17
Registering Measures Using regmeasure	9-18
regmeasure Usage.....	9-18
Registering Token Measures Using regTokenMeasure	9-24
regTokenMeasure Usage.....	9-25
Batch Plug-In Tasks: execPluginTask.sh	9-25
execPluginTask.sh usage.....	9-26
 10 Informational Utilities	
Retrieving Domain Information Using domaininfo	10-1
domaininfo Usage	10-1
Checking the Validity of a Domain Using checkDomain.....	10-3
checkDomain Usage	10-3
Determining RPAS Server Version Using rpasversion	10-4
rpasversion Usage.....	10-4
Listing Contents of a Database Using listDb	10-4
listDb Usage	10-4
Printing Data from Arrays Using printArray.....	10-4
printArray Usage	10-4
Printing Data from Measures Using printmeasure	10-5
printmeasure Usage.....	10-6
 11 Internationalization	
Translation.....	11-1
Translation Administration.....	11-2
Translation Administration Workbook	11-4
Hierarchy Labels Worksheet.....	11-4
Dimension Labels Worksheet	11-5
Workbook Template Group Labels Worksheet.....	11-5
Workbook Template Labels Worksheet	11-5
Measure Labels Worksheet.....	11-5
Measure Descriptions Worksheet.....	11-5
User Group Labels Worksheet	11-5
Message Labels Worksheet.....	11-5
RGRP Labels Worksheet	11-5
 12 Commit as Soon as Possible	
Using Commit ASAP	12-1
Managing the Workbook Queue Using showWorkbookQueues	12-2
showWorkbookQueue Usage.....	12-2
Commit ASAP Settings Using configCommitAsap	12-3
configCommitAsap Usage.....	12-3
Logging and Technical Information.....	12-4

13 Batch Processes and RPAS Utilities

CSV File Format.....	13-1
RPAS Utilities Logging Options	13-2
Log Levels	13-2
Utilities with Standard Logging.....	13-3
Scripts	13-3
Utilities with Multi-Process Logging	13-3
Utilities with Special Logging	13-4
Utilities and Database Locks	13-6
Using Shell Scripts to Run Batch Processes.....	13-6
A Sample Shell Script	13-6
Common Information and Parameters for RPAS Utilities.....	13-7
Configuration Tools Log Files	13-8
RPAS Intraday Enabler	13-8
ride Usage.....	13-9
Scenarios	13-10
Scenario 1	13-10
Scenario 2	13-11
Scenario 3	13-11
Scenario 4	13-12
Scenario 5	13-12
Scenario 6	13-13
Domain Lock Status Using domainStatus.....	13-13
domainStatus Usage	13-14

14 RPAS ODBC/JDBC Driver

ODBC Configuration.....	14-1
Defining the ODBC Server Configuration Settings.....	14-2
Adding RPAS ODBC Manager in the Management Console	14-2
Configuring the Use of the ODBC Manager	14-7
Importing the Configuration Changes	14-13
Defining the ODBC Client Configuration for Windows	14-13
Starting the RPAS ODBC Server Process	14-17
Testing the Connection Using Interactive SQL.....	14-17
ODBC Client Configuration for UNIX	14-17
Client Configuration	14-18
Testing the Connection.....	14-19
Installing and Using the RPAS JDBC Driver	14-19
Updating Environment Variables for the JDBC Driver on Windows	14-20
Updating Environment Variables for JDBC Driver on UNIX and Linux	14-21
Using the RPAS JDBC Driver	14-21
Enabling Spy for the RPAS JDBC Driver.....	14-21
Using the jdbcisql Utility Provided with the RPAS JDBC Driver.....	14-22
Using Oracle SQL Developer	14-22
Using Oracle JDeveloper.....	14-23
Using a Java Program	14-24
Running the Program.....	14-25

Data Query	14-25
Limitations	14-25
Contention.....	14-25
Workbook Queries	14-25
Metadata	14-26
Fact and Dimension Tables.....	14-26
Measure Security in the ODBC Driver	14-28
Use Cases	14-28
Using Metadata Tables to Explore the Structure of a Domain or a Workbook	14-28
Querying Fact Data	14-29
Connecting to a Workbook	14-30
Requesting Additional Aggregate Tables.....	14-30
Clients	14-31
Oracle Business Intelligence Enterprise Edition (OBIEE)	14-31
Configuring the ODBC Client for OBIEE	14-31
Connecting OBIEE to an RPAS Domain	14-32
Microsoft Access.....	14-33
JDeveloper	14-33
XML Publisher	14-34
Interactive SQL (ISQL) Utility	14-35
Supported and Unsupported SQL Functions	14-35
Supported SQL Functions	14-35
Numeric Functions	14-36
String Functions	14-38
Time and Date Functions	14-42
System Functions	14-46
Aggregate Functions.....	14-46
Other Functions.....	14-47
Unsupported SQL Functions.....	14-52
Handling of NULLS.....	14-52
Schema Information.....	14-53
Data Definition Language (DDL)	14-53
Insert	14-53
SELECT Syntax.....	14-53
Value Expressions.....	14-53
Value Functions.....	14-53
Date and Time Functions.....	14-53
Advanced Value Expressions.....	14-54
Row Value Constructor.....	14-54
Predicates	14-54
Join Operators.....	14-55

15 Publishing Measure Change Events

Configuring Subjects and Measures for Monitoring	15-2
Configuring the RPAS JMS Publisher	15-3
Command Line.....	15-4
Configuration Settings.....	15-4

General Settings	15-4
Vendor-Specific Settings	15-5
File-Based JNDI Object Store.....	15-5
LDAP-Based JNDI Object Store	15-6
Configuring the RPAS JMS Subscriber.....	15-7
Command Line.....	15-7
Configuration Settings.....	15-7

16 In-Context Launch

Launching RPAS	16-1
Issuing a Launch Request from a Web Page	16-2
Embedding an RWL Applet on a Web Page.....	16-2
Using JavaScript with the RWL Applet.....	16-4
User Authentication.....	16-8
Deploying In-Context Launch.....	16-9
Use Cases.....	16-10
RPAS Launch Context Format	16-11
Client Settings.....	16-11
Server Settings	16-12
Domain Settings	16-12
Workbook Settings.....	16-12
XML Schema	16-12
Open Workbook	16-16
Specifications	16-17
Examples	16-17
Build Workbook	16-18
Specifications	16-19
Example.....	16-21
The buildWorkbookContext Utility	16-23
buildWorkbookContext Usage	16-23

17 Hybrid Storage Architecture

Data Model	17-3
Naming Convention	17-4
Dimension-Fact Entity Relationship Diagram	17-4
Sample Integration Configuration.....	17-4
Sample ERD for Dimension and Fact Tables	17-5
Table Descriptions.....	17-6
RPAS Metadata Tables	17-6
Commit Metadata Tables.....	17-9
Dimension-Based Tables.....	17-10
Translation Table.....	17-11
Index Translation Table	17-13
Fact Tables.....	17-13
Temporary and Logging Tables.....	17-14
Schema Partitioning.....	17-14
Dimension and Fact Table Partitioning.....	17-14

ITT Table Partitioning	17-15
Sequence Descriptions.....	17-15
Package Descriptions.....	17-15
Access Schema Descriptions.....	17-16
Security Considerations	17-17
External Integration	17-17
Setup	17-17
Execution.....	17-19
Configuring the RPAS Data Mart.....	17-21
RDM Repository.....	17-21
Directory Structure	17-22
RDM Properties in a Domain	17-23
Schema Information.....	17-24
Integration Configuration.....	17-26
Shared Hierarchies.....	17-26
Shared Facts	17-26
Integration Map.....	17-26
Domain and Configuration Map	17-26
Example.....	17-26
Partition Information	17-28
RDM SSL Connection	17-28
SSL Overview	17-28
Set Up SSL on Oracle Server.....	17-29
RPAS HSA Build Process	17-34
RPAS Installation Utilities	17-34
RDM Preparation Script.....	17-34
RPAS Installer.....	17-35
RDM Manager	17-35
Pre-Installation – Preparing the Integration Environment	17-39
Installation – Building the RDM Schemas.....	17-41
Post-Installation – Initialization and Registration.....	17-42
Load Dimension Data.....	17-42
Register Participating Domains	17-42
Load Shared Fact Data	17-42
Operating and Maintaining the RPAS Data Mart	17-43
upgradeRDM	17-44
Usage.....	17-44
patchRDM	17-45
Usage.....	17-46
loadDimData.....	17-47
Flat File Format	17-47
Filtered Shared Levels	17-47
Validation of Input Data	17-47
Usage.....	17-48
loadFactData	17-49
Input File Format	17-49
Input Files Location	17-49

Input File Naming.....	17-49
Input Data Validation.....	17-50
Data Loading Modes	17-50
Logging.....	17-50
Usage.....	17-51
transferFactData	17-51
Usage.....	17-51
checkRdm	17-52
Usage.....	17-52
RDM Partitioning and RDM Partition Manager	17-55
Determining the RDM Partition Level	17-55
Mapping Local Domains to RDM Partitions.....	17-56
RDM Schema Partitioning	17-57
RDM Partition Manager	17-58
Usage.....	17-58
The Format of RDM Partition Configuration	17-59
Initial RDM Partitioning	17-60
Patching RDM and Integrating One or More New Domains.....	17-61
Hierarchy Operations Affecting RDM Partitioning	17-62
Dimension Data Loading in RDM and Domain.....	17-62
Domain Reclassification.....	17-63
Dynamic Position Maintenance	17-63
Coordinating RDM Partitioning Using the RDM Status Table	17-63
Status: loadHier in Progress	17-64
Status: reconfigGlobalDomainPartitions in Progress	17-64
Status: loadDimData in Progress	17-64
Status: Repartition in Progress.....	17-65
Status: Repartition Recommended	17-65
Coordination of Domain Processes in an Integrated Environment	17-65
Recommended Process Sequencing	17-66
Parallel Execution of Processes	17-66
RDM Status Tables.....	17-68
RDM Status Messages	17-68
RDM is Missing Positions.....	17-70
Repartition Recommended.....	17-71
RDM Status Utility.....	17-71
Usage of rdmStatus Utility	17-71
RPAS HSA Translation	17-71
Flat File Requirements.....	17-72
Dimension Label Flat File	17-72
Level Label Flat File	17-72
Fact Label and Description Flat Files	17-72
Position Label Flat Files	17-72
Source Language	17-72
loadTranslation.....	17-73
Input Files.....	17-73
loadTranslation Utility	17-73

A Appendix: Curve Administration Guide

curvevalidate	A-1
curvevalidate Usage.....	A-1
curvebatch	A-3
curvebatch Usage	A-3

B Appendix: RPAS Test Driver

Support	B-1
Introduction.....	B-1
Audience.....	B-1
Finding Required Information	B-1
Compatibility	B-2
Interface.....	B-2
XML Requests	B-4
Connection Messages	B-4
<Logon>	B-4
<Logoff>	B-4
Wizard Messages.....	B-4
<WizardInitialize>	B-4
<GetNextWizardDialog>.....	B-4
<GetWizardLayout>.....	B-4
<SelectPositions>	B-5
<SetWizardData>.....	B-5
<WizardFinish>.....	B-6
Workbook Messages.....	B-6
<OpenWorkbook>	B-6
<Fetch>	B-7
<SetValues>	B-7
<CommitWorkbook>	B-8
<RefreshWorkbook>	B-8
<MenuEvent>	B-8
<CreatePOPMeasure>.....	B-9
<SaveWorkbook>.....	B-9
<CloseWorkbook>	B-9
RTD Use with HP LoadRunner	B-9
Prerequisites	B-10
Header File.....	B-10
lr_rtd.h	B-10
LoadRunner Scripting	B-10
Script Settings	B-10
Script Sections	B-11
Script Implementation	B-11
Script Section: vuser_init.....	B-11
Script Section: Action (simple test).....	B-11
Script Section: Action (real-world)	B-12
Considerations for Scenario Use	B-13

C Appendix: RPAS Test Automation

Introduction.....	C-1
Usage.....	C-1
Writing Test Cases	C-2
Example	C-4
Output.....	C-6
Schema.....	C-7
1. Top-Level Parent Tag	C-7
1.1 testscript	C-7
1.2 testsuite.....	C-7
1.3 testcase	C-7
1.4 once	C-8
1.5 finalize	C-8
1.6 setup.....	C-8
1.7 teardown	C-8
1.8 log	C-8
Sample Usage	C-8
2. Generic Tags.....	C-8
2.1 Attribute	C-8
2.2 Description.....	C-8
2.3 Fail	C-9
2.4 key	C-9
2.5 rpas-today	C-9
2.6 shell	C-9
2.7 register-measure.....	C-9
2.8 unregister-measure	C-10
3. Domain Operation Tags	C-10
3.1 Assertions.....	C-10
3.2 Domain Data Edit Operations.....	C-19
4. Workbook Operation Tags	C-21
4.1 Workbook Wizard Operations.....	C-21
4.2 Workbook Operations	C-23
4.3 Workbook Assertions	C-25
4.4 Workbook Data Edit Operations	C-33
5. Dynamic Position Maintenance (DPM) Tags.....	C-36
5.1 Assertions.....	C-36
5.2 Update Informal Position Operation	C-37
5.3 Delete Informal Position Operation	C-38
6. HSA Tags.....	C-38
6.1 assert-fact-group	C-38
6.2 assert-fact-table	C-39
6.3 assert-fact-value-eq.....	C-39
6.4 assert-fact-value-ge	C-39
6.5 assert-fact-value-gt.....	C-40
6.6 assert-fact-value-le	C-40
6.7 assert-fact-value-lt.....	C-41
6.8 assert-fact-value-ne	C-41

6.9 assert-level-contain	C-42
6.10 clear-rdm-status	C-42
6.11 set-rdm-status	C-42
6.12 set-fact-value.....	C-42
6.13 reset-seq-generator.....	C-43
6.14 assert-rdm-hier-exists.....	C-43
6.15 assert-level-exists	C-43

D Appendix: Environment Variables

Required Settings	D-3
Database Settings	D-3
RPAS_ARRAY_TAG.....	D-3
RPAS_CACHE.....	D-3
RPAS_PAGE_SIZE.....	D-3
RPAS_PAGE_SPLIT_PERCENTAGE	D-4
RPAS_LOCK_TIMEOUT	D-4
RPAS_REQUEST_TIMEOUT	D-5
RPAS_USER_MODE.....	D-5
Java Special Expression Settings	D-5
RPAS_JAVA_CLASSPATH	D-5
RPAS_JAVA_DEBUG_PORT	D-5
RPAS_JAVA_MAXHEAPSIZE.....	D-6
Log Settings	D-6
RPAS_CACHE_LM_MULTIPILER / RPAS_DISABLE_LM_DIRECT	D-6
RPAS_CLIENT_TIMEOUT / RPAS_SERVER_TIMEOUT	D-6
RPAS_COMMIT_FAILED_MAX / RPAS_COMMIT_SUCCESS_MAX.....	D-6
RPAS_JAVA_MAXHEAPSIZE.....	D-6
RPAS_LOCK_DETAIL_PATH.....	D-6
RPAS_LOCK_LOGGING.....	D-7
RPAS_LOCK_LOG_PATH.....	D-7
RPAS_LOCK_TRACK_DATABASES	D-7
RPAS_LOCK_ACCESS_LIMIT	D-7
RPAS_LOG_BACKUPS	D-7
RPAS_LOG_FD	D-8
RPAS_LOG_FLUSH.....	D-8
RPAS_LOG_LEVEL	D-8
RPAS_LOG_PATH	D-8
RPAS_REQUEST_TIMEOUT	D-8
RPAS_SHELL.....	D-9
RPAS_SPLIT_BUF_MB.....	D-9
RPAS_STATS_PATH.....	D-9
RPAS_TCPNODELAY	D-9
RPAS_TRACE_SSL	D-9
TMPDIR.....	D-9
TODAY_RETEK	D-10
Parallelization Settings	D-10
RPAS_EXPORT_MEASURES_MAX	D-10

RPAS_MIN_MEAS_PER_PROC	D-10
RPAS_PROCESSES	D-10
Profiling Logging	D-10
RPAS_PROFILING_ENABLE	D-10
RPAS_PROFILING_PATH	D-10
Date and Time Setting	D-11
RPAS_TODAY	D-11
RPAS_TODAY_STATIC	D-11
Numeric Precision	D-11
RPAS_INCAGGPC	D-12
epsilon	D-12
RPAS_SSL_WALLET	D-12
LDR_CNTRL=NAMEDSHLIB=RPASZONE	D-12
Nagle's Algorithm	D-13

Send Us Your Comments

Oracle Retail Predictive Application Server Administration Guide for the Classic Client, Release 14.1

Oracle welcomes customers' comments and suggestions on the quality and usefulness of this document.

Your feedback is important, and helps us to best meet your needs as a user of our products. For example:

- Are the implementation steps correct and complete?
- Did you understand the context of the procedures?
- Did you find any errors in the information?
- Does the structure of the information help you with your tasks?
- Do you need different information or graphics? If so, where, and in what format?
- Are the examples correct? Do you need more examples?

If you find any errors or have any other suggestions for improvement, then please tell us your name, the name of the company who has licensed our products, the title and part number of the documentation and the chapter, section, and page number (if available).

Note: Before sending us your comments, you might like to check that you have the latest version of the document and if any concerns are already addressed. To do this, access the Online Documentation available on the Oracle Technology Network Web site. It contains the most current Documentation Library plus all documents revised or released recently.

Send your comments to us using the electronic mail address: retail-doc_us@oracle.com

Please give your name, address, electronic mail address, and telephone number (optional).

If you need assistance with Oracle software, then please contact your support representative or Oracle Support Services.

If you require training or instruction in using Oracle software, then please contact your Oracle local office and inquire about our Oracle University offerings. A list of Oracle offices is available on our Web site at <http://www.oracle.com>.

Preface

Oracle Retail Administration Guides are designed so that you can view and understand the application's behind-the-scenes processing, including such information as the following:

- Key system administration configuration settings
- Technical architecture
- Functional integration dataflow across the enterprise
- Batch processing

Audience

This document is intended for the users and administrators of Oracle Retail Predictive Application Server. This may include merchandisers, buyers, and business analysts.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at

<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit

<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit

<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Related Documents

For more information, see the following documents in the Oracle Retail Predictive Application Server documentation set:

- *Oracle Retail Predictive Application Server Administration Guide for the Fusion Client*
- *Oracle Retail Predictive Application Server Installation Guide*
- *Oracle Retail Predictive Application Server Release Notes*
- *Oracle Retail Predictive Application Server User Guide for the Fusion Client*

Customer Support

To contact Oracle Customer Support, access My Oracle Support at the following URL:

<https://support.oracle.com>

When contacting Customer Support, please provide the following:

- Product version and program/module name
- Functional and technical description of the problem (include business impact)
- Detailed step-by-step instructions to re-create
- Exact error message received
- Screen shots of each step you take

Review Patch Documentation

When you install the application for the first time, you install either a base release (for example, 14.1) or a later patch release (for example, 14.1.1). If you are installing the base release or additional patches, read the documentation for all releases that have occurred since the base release before you begin installation. Documentation for patch releases can contain critical information related to the base release, as well as information about code changes since the base release.

Improved Process for Oracle Retail Documentation Corrections

To more quickly address critical corrections to Oracle Retail documentation content, Oracle Retail documentation may be republished whenever a critical correction is needed. For critical corrections, the republication of an Oracle Retail document may at times not be attached to a numbered software release; instead, the Oracle Retail document will simply be replaced on the Oracle Technology Network Web site, or, in the case of Data Models, to the applicable My Oracle Support Documentation container where they reside.

This process will prevent delays in making critical corrections available to customers. For the customer, it means that before you begin installation, you must verify that you have the most recent version of the Oracle Retail documentation set. Oracle Retail documentation is available on the Oracle Technology Network at the following URL:

<http://www.oracle.com/technetwork/documentation/oracle-retail-100266.html>

An updated version of the applicable Oracle Retail document is indicated by Oracle part number, as well as print date (month and year). An updated version uses the same part number, with a higher-numbered suffix. For example, part number E123456-02 is an updated version of a document with part number E123456-01.

If a more recent version of a document is available, that version supersedes all previous versions.

Oracle Retail Documentation on the Oracle Technology Network

Documentation is packaged with each Oracle Retail product release. Oracle Retail product documentation is also available on the following Web site:

<http://www.oracle.com/technetwork/documentation/oracle-retail-100266.html>

(Data Model documents are not available through Oracle Technology Network. These documents are packaged with released code, or you can obtain them through My Oracle Support.)

Documentation should be available on this Web site within a month after a product release.

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
<code>monospace</code>	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Introduction

This guide is for administrators of the RPAS Server and the RPAS Classic Client.

Administrator Overview

After the RPAS Server and Classic Client have been installed, administrators must set up the RPAS Classic Client and complete several administration activities before they can begin using RPAS and RPAS applications. The activities include the following:

- [Domain Administration](#)
- [User Maintenance](#)
- [System Administration](#)
- [Hierarchy Management](#)
- [Data Management](#)
- [Translation Administration](#)
- [Hybrid Storage Architecture](#)

Before you start any of these activities, you should understand the basics of RPAS: domains, workbooks, worksheets, hierarchies, and measures.

Basic Concepts of RPAS

Retail Predictive Application Server (RPAS) is a configurable platform with a proven scalability for developing multidimensional forecasting and planning based solutions. This platform provides capabilities such as a multidimensional database structure, batch and online processing, a configurable slice-and-dice user interface, a sophisticated configurable calculation engine, user security and utility functions such as importing and exporting, all on a highly scalable technical environment that can be deployed on a variety of hardware.

This section introduces you to the following RPAS concepts:

- [Multidimensionality](#)
- [Hierarchies](#)
- [Measures](#)
- [Domains and Workbooks](#)

Multidimensionality

In RPAS, information is stored and represented based on the multidimensional framework. In a multidimensional database system, data is presented as a multidimensional array, where each individual data value is contained within a cell accessible by multiple indexes.

Multidimensional database systems are a complementary technology to entity relational systems and achieve performance levels above the relational database systems. Applications that run on RPAS identify data through dimensional relationships. In RPAS, dimensions are called hierarchies. Hierarchies are qualities of an item (such as a product, location, or time) or components of a dimension that define the structure and roll up within the dimension.

Hierarchies

Hierarchies describe the top-to-bottom relationship between the levels or positions of the hierarchies in RPAS. They reflect the hierarchies set up at your business and being used by the merchandising solutions.

RPAS supports many alternative hierarchies that provide different roll ups and help analyze the data from a different perspective.

Measures

Measures represent the events or measurements that are recorded, while the positions in the dimensions provide a context for the measurement. Measures are defined based on the business rules set in the application. The dimensionality of a measure is configured through the definition of its base intersection, which is the collection of levels (one per appropriate dimension) defining the lowest level at which the information is stored for the measure.

Measure names are completely configurable and typically named using a convention that identifies each component and the meaning of the measure.

Domains and Workbooks

RPAS stores information in a persistent multidimensional data cache that is optimized for large volumes and dimensional or time series data access requirements, typically required by multidimensional solutions. This central repository is called a domain. The domain also includes central definitions of metadata for the solution and provides a single update point.

When you use an RPAS solution, you interact with the solution through a personal data repository called a workbook. A workbook contains the subset of the data (and metadata) from the domain and its scope is constrained by the access rights available to a user. Workbooks are stored on the RPAS server, and can be built using an online wizard process or scheduled to be built in a batch process automatically. Workbooks are made up of one or more worksheets. These worksheets display the hierarchy and measure data of the domain.

Although the data and metadata in the workbook are copied from the domain, the data remains independent of the domain.

Domains can be built in one of two methods:

- Simple domain: This is the traditional, stand-alone domain that has no visibility to other domains.

- Global domain: This is a domain environment that contains two or more local domains (or subdomains) and a master domain that has visibility to all local domains that are part of that environment.

A global domain is a type of domain structure that provides users with the ability to view data from multiple domains and to administer common activities of an RPAS domain and solution.

Using a global domain environment has two primary functional benefits. The first feature allows users to have a global view of data in workbooks. Users can build workbooks with data from local domains, refresh global workbook data from local domains, save global workbooks, and commit the data from global workbooks to the individual local domains.

Local domains are typically organized, or partitioned, along organizational structures that reflect user roles and responsibilities. Most users only work within the local domains that contain their area of responsibilities, and they may not need to be aware of the global domain environment. For performance and user contention reasons, global domain usage should be limited to relatively infrequent processes that require data from multiple local domains.

The other primary feature of global domain is centralized configuration and administration. Most of the mechanisms that are required to build and administer a domain have been centralized and they need only be run in the master domain, which either propagates data to the local domains or stores the data centrally so that the local domains reference it in the master domain.

Note: For a global domain environment to function properly, all local domains must be structurally identical.

Measure Data

In a global domain environment, measure data can be physically stored in two different ways:

- Across the local domains
- In the master domain

Measure data that is stored in local domains is split across domains based on a pre-determined level of a given hierarchy. This level is defined during the configuration process, and it is referred to as the partition level.

The base intersection of a measure (for instance, what dimensions a measure contains) determines whether data is stored in the local domains or in the master domain. The data is stored in the master domain if the base intersection of a measure is above the partition level or if it does not contain the hierarchy on which the global domain environment is partitioned. This type of measure is referred to as a global domain measure or a higher base intersection measure.

Consider a global domain environment where the partition-level is based on the Department dimension in the Product hierarchy. Data for measures that have a base intersection in the Product hierarchy at or below Department are stored in the local domain based on the Department to which the underlying position in the Product hierarchy belongs. Other hierarchies are irrelevant for this discussion.

However, measures that have a higher base intersection in the Product hierarchy than Department (for instance, Division) or measures that do not contain the Product hierarchy (such as a measure based at Store-Week) cannot be split across the local

domains. These measures will reside in the master domain and will be accessed from there when these measures are required in workbooks.

All measures are registered in the master domain, and they are automatically registered in all local domains. RPAS automatically determines where the measure needs to be stored by comparing the base intersection of the measure against the designated partition-level of the global domain environment. The physical locations of the measure data are invisible to the user after the measure has been registered.

Administrative Workbooks and Wizards

Using the administration workbooks, designated employees manage other employees' use of the Oracle Retail Predictive Solutions. System administrators use the administration workbooks to perform the following:

- Set up and maintain users and user groups.
- Manage user access to specific workbook templates and individual measures.
- Edit the contents of translation tables to support multiple-language use of the application.
- Specify the type, frequency, and format of workbooks in the automatic build queue.

Note: If a solution is built in a global domain environment, most administrative activities can only be performed in the master domain. This applies to RPAS administrative workbook templates and wizards as well as RPAS utilities that are run on the backend against the domain. See each workbook or workbook wizard section in this guide for details about the domain access.

Hybrid Storage Architecture

This document contains an overview of the RPAS Hybrid Storage Architecture (HSA). The target audience includes business executives, planners, integrators, and pre-sales consultants with exposure to Retail Predictive Application Server (RPAS).

Numerous corporations adopt Oracle database for enterprise data storage and often leverage it as the system of record in integrated environments. By contrast, RPAS applications store data in a so-called domain that is a proprietary multi-dimensional database. One natural question is why RPAS does not use Oracle database.

The answer is simple. While Oracle database is a feature-rich enterprise data storage solution, it does not provide the agile and high-performing procedures for some of the access/insert patterns needed for multi-dimensional data processing. By contrast, the RPAS multi-dimensional persistence layer is optimized for fast data access and random write patterns used during multi-dimensional data processing. RPAS operations such as workbook commit and batch calculations would not perform if the entire multi-dimensional data set was to be stored in Oracle database.

A Note about RPAS Configurability and Extensibility

RPAS is a configurable and extensible platform providing a significant amount of flexibility in design and implementation. Lack of specific verbiage around a platform component does not guarantee a commitment to support its use in a non-documented

way. In case of any ambiguity in the documentation, the customer is expected to ask for clarification with Oracle Customer Support.

Building and Upgrading Domains

This chapter describes how to build and upgrade domains. It contains the following sections:

- [Building a Domain](#)
- [Building the Domain Manually or Through a Command-Line Interface](#)
- [After Building the Domain](#)
- [fixDomain](#)

Building a Domain

After a fully defined configuration is created, an RPAS domain can be installed. Since building an RPAS domain is a manual process, it is expected that this process is supported by UNIX administration if the domain is installed on a UNIX platform. If the domain is being installed on Windows servers for prototyping and demonstration purposes, it can be built using the RPAS Configuration Tools GUI installer.

Prerequisites

The following are the prerequisites for building a domain:

- Installation of RPAS on the server that will store the domains.
- Installation of the Configuration Tools on the server that will store the domains.
- A configuration built using the Configuration Tools.
- A collection of hierarchy input files that contain positions for the domain. A hierarchy data file (name.dat) is required for each defined hierarchy.
- Cygwin installed (for prototyping and demonstration on a Windows server only).

Client-Side Procedures

The following client-side procedures must be completed to build a domain.

To begin the domain build process, a configuration project built using the Configuration Tools is required. This can be a packaged template or a configuration created with the customer's specific hierarchies, measures, and workbooks. If using a new configuration, be sure to note the path where the configuration is saved on the local disk.

Note: The remainder of this section assumes that the domains are being built on a different server than Windows while the configuration is created on the Windows platform. If the domain is being built on the same server as the configuration, the steps for adding the configuration to the zip file and transferring to a different server can be eliminated.

Zip the Configuration Project Folder

1. Find the location where the configuration project is saved.
2. Using Windows Explorer, go to the path of the configuration project.
3. Right-click the **Configuration** folder and select **Add to Zip**. Package the entire contents of the project beginning with the configuration project root folder so that the zip file includes all solutions. It is important to zip the entire configuration project for the entire directory structure and not just the specific .xml files. Do not change the name of the configuration project folder or alter the contents of the folder in any way.

In [Figure 2–1](#), TPGA is the configuration selected to create the TPGA.zip.

Figure 2–1 Zipping the Configuration Project Folder



4. Using FTP, transfer the .zip file over to the server in binary mode. This can be placed in the home directory for now.

Server-Side Procedures

The following server-side procedures must be completed to build a domain.

Note: Though the RPAS Configuration Tools are supported only on the Windows platforms, the installation tools are supported on all platforms. However, they require Java 1.6. Ensure that the server being used for the domain installation has the correct version of Java.

Unzip the Configuration File

1. Find the location where you are going to save the configuration project file.

Note: Always put an updated configuration project in a new directory path. Do not overwrite an existing configuration project.

2. Move the <Configuration Project>.zip file to this location.
3. Unzip the <Configuration Project>.zip file using the UNIX command:

```
unzip -a <Configuration Project>.zip
```

4. Do not change the directory name for the configuration project or alter the contents in any way.

Verify the Environment Variable Settings

Prior to this step, RPAS and the Configuration Tools must be installed on the server that will store the domains. During that process, the necessary environment variables for RPAS and the Configuration Tools must be defined. Refer to the "[Appendix: Environment Variables](#)" if the environment variables below have not been set up.

Log in to the server. Use the commands below to verify the environment settings:

```
echo $RPAS_HOME
echo $RIDE_HOME
echo $JAVA_HOME
echo $PATH
```

Note: The path for the RPAS_HOME variable may change with each new RPAS release.

If any changes are made to the environment variable settings, remember to exit and restart the UNIX session in order to execute the corrected .profile. This step must be done before continuing with the remaining steps.

Get the Input Files Ready

1. Designate a directory for the location of the input files and move the files into this directory.

Note: As a recommendation, use the directory name to_rpas as a standard for the location of input files. At a minimum, the hierarchy files (product, location, and calendar hierarchy files) are needed to build the domain. At this time, a calendar file must be loaded.

2. If necessary, rename the hierarchy files to match the name of the configured hierarchies. The files must end in either .dat or .csv.dat. For example, a file for a configured product hierarchy named prod in the Configuration Tools should be either prod.dat or prod.csv.dat. When using the .dat extension, the format of the files must match the hierarchy configuration specified using the Configuration Tools. If using the .csv.dat extension, the files will contain fields that are comma separated. See "[Loading Hierarchies Using loadHier](#)" for more details.

Building the Domain Manually or Through a Command-Line Interface

Run the rpaInstall utility to build the domain. This executable is located in the path to \$RIDE_HOME/bin.

rpaInstall

This section provides the details required for the upgrade of rpaInstall.

rpaInstall Usage

```
rpaInstall <arguments>
```

[Table 2–1](#) provides descriptions of the arguments used by the `rpasInstall` utility.

Table 2–1 Arguments Used by the `rpasInstall` Utility

Argument	Description
<code>[-fullinstall -patchinstall -testinstall]</code>	Indicates the type of installation to be performed, where: ▪ <code>-fullinstall</code>: Builds a full domain and loads the hierarchy data files. ▪ <code>-patchinstall</code>: Patches an existing domain. Updates or unregisters/registers measures that have changed (as necessary). ▪ <code>-testinstall</code>: Used for testing only. Only generates configuration files. Required.
<code>-ch config_home</code>	Indicates path to the directory containing the configuration file. Required.
<code>-cn config_name</code>	Indicates the name of the configuration. Required.
<code>-in input_home</code>	Indicates the path that includes the directory containing the input files for the domain to be created. If this directory includes a <code>users.db</code> file, then the users and groups within that file are added upon domain creation. For more information, see "Managing Users Using <code>usermgr</code>". Required.
<code>-log log_name</code>	Indicates the path that includes the name of the log file to be created or updated. Required.
<code>-configdir config_directory</code>	Indicates the path to the directory containing the xml files used by RPAS. This is a required argument if the user wants to supply <code>globaldomainconfig.xml</code> or <code>calendar.xml</code>.
<code>-dh domain_home</code>	Indicates the path to directory in which the domain will be created. Use if and only if a <code>globaldomainconfig.xml</code> is not used.
<code>-p dim_name</code>	Indicates the partitioning dimension. Use if and only if the global domain is being implemented without the use of <code>globaldomainconfig.xml</code>.
<code>-rf function_name</code>	Indicates the filename of the function to be registered. This pairing may be repeated for multiple functions. This argument is not required if there are no functions to be registered.
<code>-updatestyles</code>	Imports configured style information into the domain. This option is automatically set in a full install. If you do not use this flag in a patch install, changes to configured styles will not be imported into the domain.

Note: When you build a domain for the first time, an installation directory is created inside the domain. The installation directory is essential for the patch process and must not be removed, moved, or renamed.

When submitting an issue to Oracle Retail Customer Support, if you are asked to provide a domain, be sure to provide the installs directory as well. The following information must be provided in order to help Customer Support better diagnose the issue:

- The configuration
 - The script used to run the rpasInstall script
 - The domain
 - The log output file
-

Validate Domain Build Results

After the domain build process is complete, the logfile should be reviewed to verify that the process executed successfully. Search for the words "ERROR," "FAILURE," and "exception" inside the logfile. The end of the logfile should look similar to the following output example:

```
Time: 58.451  
COMPLETE
```

After Building the Domain

After you build the domain, it can be accessed by the RPAS Classic Client.

Note: Building a domain creates the shell of the domain. All measures, rules, and workbooks are created, but the measures are not populated. Measures are populated with the loadmeasure utility. For more information, see ["Loading Measure Data Using loadmeasure"](#).

In order to connect to a domain, the domain information must be set with the RPAS EConfigure utility and the RPAS DomainDaemon must be running. See the *RPAS Installation Guide* for details on setting up the domain information. See ["DomainDaemon"](#) for information.

Before you can log into the new domain, you must create at least one user with the usermgr utility. After you have created one user, you can use the ["Add User"](#) workbook to create others. For more information about the usermgr utility and creating users, see the ["User Maintenance"](#) chapter.

After logging into the domain and creating the appropriate users, ensure that the appropriate permissions to workbooks and measures are set. See ["Security Administration Workbook"](#) for more information.

Upgrading and Patching Domains

RPAS supports the upgrade of RPAS 11.0.x, 11.1.x, 12.x, and 13.x environments to RPAS 14.x.

Note: See the solution-specific Installation Guide, Implementation Guide, and Operations Guide for potential solution upgrade limitations.

For details on upgrading to this release of RPAS from any release of RPAS prior to 13.3, see ["Upgrading From Any Release of RPAS Prior to 13.3"](#).

For details on upgrading to this release of RPAS from RPAS 13.3 or later, see ["Upgrading From RPAS 13.3 or Later"](#).

Upgrading From Any Release of RPAS Prior to 13.3

If you are upgrading to the current release of RPAS from any release prior to RPAS 13.3, complete the following steps.

You must first upgrade the domain to the latest 13.2.3 hotfix prior to starting this upgrade process.

Note: For information on upgrading to the latest RPAS 13.2.3 hot fix, see the *13.2.3 RPAS Administration Guide for the Classic Client*, section "Upgrading and Patching Domains." Note that it is optional to perform Step 8 in that section. Step 8 states that you should run the `rpasInstall` utility with the `-patchinstall` argument. This is not necessary if you are upgrading to this release because you will run that utility and argument at the end of this section.

Domain Conversion

In RPAS 13.3, a new methodology called Integer Indexing was introduced to simplify the management of multidimensional data in an RPAS domain. Integer indexing aims at eliminating a lot of inefficiencies in the addition, removal, and reclassification of positions in a hierarchy and improving the performance of hierarchy operations.


Since integer indexing uses different internal data structures to manage the hierarchy, a pre-13.3 domain must be converted to the new format before it can be recognized by the newer version of RPAS. This conversion is accomplished by using the `convertDomain` utility.

Domain conversion is an elaborate and complex process that runs on the assumption that the source domain is in a good state free of system level inconsistencies. In reality, an existing domain may be in an undesirable state due to various reasons. This can cause the domain conversion process to fail.

The information in this chapter can help you when upgrading to RPAS 13.3 or later. It provides general guidelines for RPAS domain conversion. It covers the preparation of the domain, setting up the conversion parameters, and carrying out the conversion and post-conversion processing.

Domain Preparation Before Conversion

[Figure 2–2](#) shows an overview of the domain preparation that is necessary before running domain conversion. All these utilities must be run under the RPAS 13.2.3 hotfix version 13.2.3.36 or later.

Figure 2–2 Overview of Domain Preparation

The procedure is as follows:

1. Upgrade the domain to the latest 13.2.3 hot fix.

RPAS supports the upgrade of RPAS 12.x environments to RPAS 13.x. If the domain is not currently on the latest 13.2.3 hot fix, it must be upgraded to that hot fix first before upgrading to 13.3.

To upgrade to the latest RPAS 13.2.3 hot fix, see the 13.2.3 version of the *RPAS Administration Guide for the Classic Client*, "[Upgrading and Patching Domains](#)". Note that it is optional to perform Step 8 in that section. Step 8 states that you should run the `rpasInstall` utility with the `-patchinstall` argument. This is not necessary if you are upgrading to 13.3 because you will run that utility and argument at the end of this section.

In addition, refer to the solution-specific Installation Guide, Implementation Guide, and Operations Guide, as appropriate, for potential solution upgrade limitations.

2. Run `scanDomain` to repair any array corruption.

The `scanDomain` utility can be used for repairing data corruption in an RPAS database. Data corruption can occur if an external program modifies the RPAS database files or an unforeseen defect occurs in the processes using the RPAS database.

To run the utility from the command line, do the following:

```
scanDomain -d domainPath -repairCorruption -processes N
```

For more information about scanDomain, see the 13.2.3 release of the *RPAS Administration Guide for the Classic Client*.

3. Run fixDomain to detect metadata errors.

The Fix Domain utility is a tool used to detect and fix data inconsistencies in a domain. It supports a few different types of "tasks" to analyze different areas of data in the domain. Task "RepairMetadata" is used to detect and repair metadata inconsistencies.

To run the utility from the command line, do the following:

```
fixDomain -d domainPath -task RepairMetadata
```

By default, this utility only reports errors without making changes to the domain. All errors are reported to standard output, which can be directed to a file.

4. Run fixDomain to fix metadata errors.

Run fixDomain again in the commit mode to actually fix the metadata errors in the domain.

To run the utility from the command line, do the following:

```
fixDomain -d domainPath -task RepairMetadata -commitChanges
```

Some metadata discrepancies can be too severe to fix. For example, discrepancies between the system arrays dim_hier and hierinfo cannot be fixed by the utility. It is recommended that you run fixDomain in report mode (without the -commitChanges flag) again to detect any remaining errors. Contact RPAS technical support if there are still errors.

Caution: Note that fixDomain may remove measures from the domain if the metadata for those measures are corrupted. This may prevent some workbooks from being built. If this happens, the missing measures should be re-registered in the domain.

5. Run fixDomain to check and rebuild change arrays.

Change arrays are system arrays used by the reshapeArrays utility. They should be rebuilt in case they are out of sync with other domain data, which can cause the failure or partial failure of the reshapeArrays.

To run the utility from the command line, do the following:

```
fixDomain -d domainPath -task RepairChangeArrays -commitChanges
```

6. Run reshapeArrays to fix dimension inconsistencies.

This cleans up potential dimension inconsistencies between data arrays and dimension dictionaries.

To run the utility from the command line, do the following:

```
reshapeArrays -d domainPath -registered -force -processes N
```

7. Run scanDomain to generate a list of likely unused databases.

An existing RPAS domain can accumulate a lot of unused databases during the course of its history. If those databases are not excluded from the domain conversion process, at a minimum they can slow down the conversion considerably, and, worse yet, they can cause the conversion to fail if any of the

data arrays in them is corrupt. The scanDomain utility has been enhanced to list all likely unused databases of the domain. Best efforts have been applied to generate a complete and accurate list. However, some databases may have been created by custom RPAS extension libraries or third party plug-ins and can be listed as unused by mistake. This list should be examined and corrected if necessary.


To run this utility from the command line, do the following:

```
scanDomain -d domainPath -listUnusedData outputFile
```

Setting up Domain Conversion Parameters

The convertDomain utility can take three configuration input files: a hierarchy configuration file, an optional global domain configuration file, and an optional database exclusion file. Figure 2–3 shows an overview of setting up these files. All utilities should be run under RPAS 13.3 hotfix version 13.3.0.17 or later, RPAS 13.3.1 hotfix version 13.3.1.5 or later, or RPAS 13.4 or later.

Figure 2–3 Setting Up Domain Conversion Parameters



The procedure is as follows:

1. Run convertDomain to generate a suggested bit size for each user dimension.

To run the utility from the command line, do the following:

```
convertDomain -src domainPath -suggestDimBitSize outputFile
```

See "[convertDomain](#)" for more information on using the suggestDimBitSize option.

2. Run ConfigTools to upgrade the domain configuration and to set the bit sizes.

Use the RPAS Configuration Tools to upgrade the 13.2.3.x domain configuration to 13.3, setting the bit size for each dimension. (Bit size configuration has an impact

on performance. See "Defining Dimension Properties" in the *RPAS Configuration Tools User Guide* for instructions on setting the bit size and reindex threshold.)

At this point, no additional changes should be made to the configuration; those changes must be processed separately at a later stage in the process. Solutions that include plug-ins that are normally run as part of the domain upgrade process will also be handled later.

3. Export the hierarchy.xml file from the updated configuration.

This file can be created by running the hierarchy.xml report from within the Report Generator of the RPAS Configuration Tools.

4. Construct the database exclusion file (optional).

An existing domain may contain a lot of stale or obsolete databases in it that can hinder the domain conversion progress. During the domain preparation, the scanDomain utility can be used to generate a list of likely unused databases. After confirmation, these databases can be put in an XML file that can be named exclusionList.xml.

See "[convertDomain](#)" for more information on the use of exclusion lists.

5. Construct the global domain XML configuration file (optional).

If you are converting a global domain, and not all local domains reside under the global domain path, then an XML configuration must be constructed to specify the source and destination path for each local domain.

See "[convertDomain](#)" for more information on domain.xml configuration files.


Running the Domain Conversion

Once the domain has been prepared and all input parameters have been generated, the domain can be converted using the convertDomain utility. This utility will create a new integer indexing domain rather than updating the existing pre-integer indexing domain. This new domain can then be used to finish the upgrade process.

See "[convertDomain](#)" for more information on using the convertDomain utility.

Post-Conversion Processing

After the domain has been successfully converted to integer indexing, some post-conversion processing steps, as shown in [Figure 2-4](#), are required to bring the domain up to date.

Figure 2–4 Post-Conversion Processing

To run post-conversion processing, do the following:

1. Run **fixDomain** to fix unlinked arrays.

If any arrays are listed in the `unlinkListFile` during the domain conversion, they must be fixed by running the **fixDomain** utility. The "FixArrays" task fixes unlinked arrays that contain one or more dimensions that are not linked to the dimension registry.

See "[fixDomain](#)" for more information on fixing unlinked arrays.

2. Upgrade the domain to the current RPAS version.

The **upgradeDomain** utility is used to upgrade only the RPAS version of the domain. It does not update the configuration or any other aspects of the domain itself. Refer to "[upgradeDomain](#)" for more information.

3. Upgrade the domain configuration.

Run the new version of the Configuration Tools to upgrade existing configurations to the latest version. Those changes should be made to the configuration modified in the Step 2 of "[Setting up Domain Conversion Parameters](#)". Refer to "[Upgrade Configurations](#)" for more information.

4. Patch the domain.


Running the patch installation through **rpasInstall** over the converted domain is required even if there are no configuration updates. Refer to "[rpasInstall](#)" for more information.

Upgrading From RPAS 13.3 or Later

If you are upgrading to the current release of RPAS from RPAS 13.3 or later, complete the following steps.

Figure 2–5 shows an overview of the upgrade and patching process. The detailed steps are provided after the figure.

Figure 2–5 Upgrade Process from RPAS 13.3.x



1. Acquire the latest 14.0.x RPAS release and install the following components:

- RPAS Server
- RPAS Client (Classic or Fusion)
- RPAS Configuration Tools

Note: If you are using `rsp_manager` to install an RPAS 14.0 patch, you should use the `-no_domain` option.

2. Run the `upgradeDomain` utility to complete the conversion of the domain to the current release. See "[upgradeDomain](#)" for more information.
3. Make any additional configuration changes to the domain (such as the re-execution of plug-ins for applications that require the plug-ins to be run and applied through a patch install as part of an upgrade).
4. If any additional changes were made to the domain in step 3, those changes should now be applied to the upgraded domain. To apply these changes, run the `rpasInstall` utility with the `-patchinstall` argument. See "[rpasInstall](#)" for more information.

Note: After you run the `upgradeDomain` utility, the patch installation (through `rpasInstall`) must be run over the domain even if there are no configuration updates. This is done to ensure that configuration within the domain is synchronized with any changes made as the result of the RPAS and Configuration Tools upgrade.

When using the `-updatestyles` flag, it is not necessary to delete existing styles.

Updating Domain Content

When moving to a new version of RPAS and RPAS solutions, you can also update the content of the domain. For example, the content of workbooks can be modified or the business logic represented by a rule group can be updated.

To update the domain content, make the desired changes in the configuration using the RPAS Configuration Tools. Then, when you run the `-rpasInstall` utility with the `-patchinstall` argument, the changes are applied to the domain.

The sections below specify the restrictions and special cases for updating hierarchies, measures, rule groups, and workbooks.

Hierarchies

When patching a domain, you can add new hierarchies. A hierarchy file must be present in the input directory specified in the call to `rpasInstall` for that hierarchy to be added.

Note: Although a new hierarchy can be added, the order of existing hierarchies cannot be changed. Therefore, new hierarchies must be added after existing hierarchies when modifying an existing domain.

Patching a domain upgrades some, but not all, hierarchy and dimension attributes. For existing hierarchies, only the security dimension is updated. For dimensions, the following attributes are updated: user-defined dimensions, labels, and the state of image support (enabled or disabled). In addition, a dimension can have DPM support and translation support enabled but not disabled. You can also change the bit size and reindex threshold of an existing dimension (see ["Reindexing Domains Using reindexDomain"](#)) as well as add or rename dimensions (see ["Adding New Dimensions"](#)).

Patching a domain does not patch changes to the hierarchy purge age, or change the multi-language setting for a domain.

Note: If updating the hierarchy purge age inside an existing domain, use the `loadHier` utility in batch mode to update the current settings. See the ["Loading Hierarchies Using loadHier"](#) section for more information.

Measure Properties

Certain measure properties cannot be modified without un-registering and re-registering the measures (which results in the loss of measure data). The measure properties that can be modified without being re-registered are listed in ["Registering Measures Using regmeasure"](#). If all of the measure properties that are changed are able to be modified without re-registration, the measure data will not be lost in the domain during the patch process.

If even one of the measure properties that change cannot be modified without re-registration, the patch process results in the un-registration and re-registration of the measure. Therefore, the data that was in the domain for that measure before the patch process will not be there after the patch process is complete. For more information, see ["Registering Measures Using regmeasure"](#).

Rule Groups

As part of the patch installation process, rule groups can be deleted, created, or modified with no restrictions. During the patching process, the rule sets, rule groups, and rules are completely rebuilt.

Workbooks

As part of the patch installation process, workbooks can be deleted, created, or modified with no restrictions.

Note: If you remove a template from the configuration and then patch the domain, the formatting for that template will be deleted.

RPAS 14.0 Cryptography Upgrades

As a part of the 14.0 release, RPAS incorporated a number of improvements to its cryptography functionality. These changes involved the ability to select a desired cipher suite and hash iteration count for the encryption of user credentials. As a result of these upgrades, the user dictionary of any domain will be automatically updated by the upgradeDomain utility when it converts a pre-14.x domain to version 14.x or later.

upgradeDomain can perform the conversion using a default cipher suite and iteration. However, it is possible to specify to upgradeDomain to use values other than the default when performing the upgrade. This is accomplished by setting two domain properties prior to running upgradeDomain. These properties are:

Table 2–2 Domain properties for Cryptography upgrade

Domain Property	Description
r_upgrade_hash_alg	This property sets the value for r_crypt_hash_alg in the upgraded domain. Because the list of algorithms can change with each release, it is not feasible to provide a list here. To get the list of values that can be used here, go to the rpasHome of the version to which you are upgrading and run: usermgr -listHashAlgs
r_upgrade_hash_iter	This property sets the value for r_crypto_hash_iter in the upgrade process. It is the number of iterations of hashing to perform on the password.

Note: The default cipher suite and hash iteration count will vary depending on the version of RPAS the domain is being upgraded to. In every release, RPAS will use the suite and iteration count deemed most secure at the time of that release.

For details on cypher suites, see "That new section Andy gave you"

For details on setting domain properties, see "Setting Miscellaneous Domain Properties using domainprop"

Upgrade Configurations


After the RPAS Configuration Tools have been installed and the domains have been upgraded, you may need to upgrade existing configurations to version 13 or the latest version 13 patch.

Note: Configuration upgrade is not always required. When it is required, the Configuration Tools automatically notifies you that a configuration upgrade is needed.

1. Launch the new version of the Configuration Tools.
2. Open the existing configuration by selecting **Open** from the File menu.

3. Select the configuration and click **OK**.
4. A message box appears prompting you to convert the configuration. Click **Yes**.

Figure 2–6 Converting the Configuration



5. The Backup Location message box appears. Click **OK** to continue.


Figure 2–7 Choosing the Backup Location



6. The administrator must now specify a location to store the backup of the current (non-upgraded) configuration. Use the browser to identify a location and enter the back-up configuration name. The backup may be stored in the same location as the configuration being upgraded only if the name of the backup is changed from its original name.

The Successful Conversion dialog box appears if the configuration was converted without any issues.

Figure 2–8 Confirmation of Conversion



7. The upgraded configuration now opens in the Configuration Tools. Changes can now be made to the configuration. Note that, if changes are made, the domain must be updated through the normal patch install process.

Note: If upgrading RDF, Grade, or Curve, see "[Upgrade Application Configurations](#)" for additional steps to complete the application upgrade.

Command-Line Support for Configuration Upgrading

The RPAS Configuration Tools also provides command-line support to upgrade configurations. The Configuration Converter is a standalone utility that converts a configuration that was originally created and saved in a prior release of the Configuration Tools. Only configurations created in a prior major release need to be converted. Configurations saved in previous versions of the same major release, but in different minor releases, do not need to be converted. See the *RPAS Configuration Tools Guide* for more information on using the Configuration Converter (RpasConverter.exe) through the command-line.

Note: If upgrading RDF, Grade, or Curve, see "[Upgrade Application Configurations](#)" for additional steps to complete the application upgrade.

Upgrade Application Configurations

Applications such as RDF, Curve, and Grade are configured using a plug-in architecture in the RPAS Configuration Tools. This architecture allows for the automation of most configuration activities for the solution. The plug-in requests specific information from the configuration administrator and the solution auto-generation tool automatically generates the solution configuration. Prior to each patch or upgrade to a major release, the auto-generation tool should be executed to ensure the solution configuration is updated with the base configuration changes for the application.

Update the solution configuration for each application in the following order:

1. Curve
2. RDF
3. RDF Clone
4. Promote
5. Grade

See the RDF, Curve, or Grade Configuration Guides for more information on the auto-generation process for each application.

convertDomain

The convertDomain utility creates a new domain that has all the dimensions and measure arrays converted to integer-indexed arrays.

The conversion process creates a new domain based on the source domain provided. Since a new domain is created from the domain and does not replace the source domain, the source domain remains unaffected.

Note: Because you are creating a second domain, you must have enough space for the original source domain and the new converted domain. Since workbooks are not converted, you only need the amount of space the source domain uses, minus the space that the workbooks in the source domain use.

The `convertDomain` utility copies all information specific to the domain at the time of conversion. It does not copy any custom information such as shell scripts, which exist under the domain root. Such information must be manually copied.

The `convertDomain` utility blindly copies the following:

- `config` (to master domain only)
- `input`
- `installs`
- `output`
- `repos`
- `scripts`
- `fusionClient`
- `WizardPQD`

The data directory is recreated based on the source domain. Only databases (.db) are converted, excluding the following:

- `changearray.db`
- `shadow.db`
- `styles.db`
- `hmaint.db` (completely rebuilt separately)
- `measdata.db` (rebuilt separately)
- `master_staging.db` (use only for loadhier)
- `alitemp*.db`

You can also use an exclusion list to exclude more databases from the conversion.

If files exist in the data directory but are not a database, they are not copied.

The users directory is recreated without all workbooks. Therefore, for all users' directory structures only contain the `<username>.db` and `<username>.db.lck` files.

The styles directory is copied conditionally. Only style files listed in `#workbook.xml` are copied. Any arrays that exist in that directory are ignored.

If your domain contains files or directories that are not copied, you must manually copy those files and directories to the converted domain.

Note: The `convertDomain` utility can be run multiple times if needed. Data already converted will be skipped after the first run.

convertDomain Usage

```
convertDomain [-src srcDomain -dest destDomain | -xmlConfigFile XMLConfigFile]
              -xmlHierFile PathToHierarchy.xml [-exclude exclusionList.xml] -genUnlinkedList
              unlinkedListFile -maxProcesses count
```

To use the `convertDomain` utility, perform the following steps:

1. Provide a pair of source/destination domain paths using the `-src` and `-dest` arguments. Or, supply an `XMLConfig` file.

- If you use `-src` and `-dest` arguments to specify the source and destination domain paths, the local domain path must be under the global domain path.
- If converting a global domain, and not all local domains reside under the global domain path, then you must use the `-xmlConfigFile` argument to specify the source and destination path for each local domain.
- The `XMLConfigFile` should follow this syntax:

```
<rpas>
  <globaldomain>
    <srcPath>/path/to/13.2.3/domain</srcPath>
    <dstPath>/path/to/13.3.0/destination</dstPath>
    <subdomain>
      <srcPath>/path/to/13.2.3/local/domain/0</srcPath>
      <dstPath>/path/to/13.3.0/destination/local/domain/0</dstPath>
    </subdomain>
    ...
  </globaldomain>
</rpas>
```

2. Use the `-xmlHierFile` to specify the path to the `hierarchy.xml` file that is exported from the upgraded configuration, including a new `BitSize` attribute for each dimension.
3. If it is desired, the `-exclude` argument may be used to specify databases that should be excluded from the conversion process. The databases specified in the `exclusionList.xml` document will not be recreated in the new integer indexing domain. The `exclusionList.xml` should use the following format. Note that if the path is an absolute path (starting with a slash "/"), it refers to only one directory. If it is a relative path, it applies to the master domain and all local domains in a global domain environment.

```
<rpas>
  <convertDomain option="exclude">
    <path>relative/path/to/database1.db</path>
    <path>relative/path/to/database2.db</path>
    <path>/absolute/path/to/database.db</path>
    ...
  </convertDomain>
</rpas>
```

4. Under certain conditions, an array in a pre-integer indexing domain can become corrupted such that the embedded dimension information of the array does not match the dimension information contained in the dimension dictionary of the domain. These arrays are called unlinked arrays. When `convertDomain` is run with the `-genUnlinkedList` argument, it creates a file in the specified location containing any arrays that are corrupted. The `fixDomain` utility can then be used to bring these unlinked arrays into conformity with the domain. See ["fixDomain"](#) for more information on restoring unlinked arrays.
5. Use the `-maxProcesses` argument to specify the number of processes that can be used.

[Table 2–3](#) provides descriptions of the arguments used by the `convertDomain` utility.

Table 2–3 Arguments Used by the `convertDomain` Utility

Argument	Description
<code>-src srcDomain</code>	Specifies the path to the source domain being converted.

Table 2–3 (Cont.) Arguments Used by the convertDomain Utility

Argument	Description
-dest <i>destDomain</i>	Specifies the path to the destination domain.
-xmlConfigFile <i>XMLConfigFile</i>	If you are converting a global domain and not all local domains reside under the global domain path, use this argument to specify the source and destination paths for each local domain.
-xmlHierFile <i>PathToHierarchy.xml</i>	Specifies the path to the hierarchy.xml file that is exported from the upgraded configuration. This file should include the new BitSize attribute for each dimension.
-exclude <i>exclusionList.xml</i>	Excludes the databases listed in the specified file from the conversion process.
-genUnlinkedList <i>unlinkedListFile</i>	Creates a file in the specified location containing the names of arrays whose embedded dimension information does not match the dimension dictionary of the domain. These arrays can then be corrected using the fixDomain utility.
-maxProcesses <i>count</i>	If specified, some parts of convertDomain will run in parallel. This means that it will use a maximum of the defined processes, which are specified by <i>count</i> .

Note: The utility upgradeDomain is not called by this process. It must be called separately. This allows the user to specify additional parameters to upgradeDomain. Once the domain has been converted and upgraded, it is ready for use.

convertDomain -suggestDimBitSize

The convertDomain utility can also be used to analyze a pre-integer indexing domain in order to suggest appropriate values for the bitSize attributes of the dimensions in that domain.

The convertDomain utility enumerates all non-system dimensions in the source domain, calculates a suitable bit size based on the current number of positions (including unused buffer positions) within the domain, and prints them out in a table format. It then calculates the key size for each of the data arrays for all measures, based on those dimension bit sizes. If the key size is above 64 bits for any measure, it will print out the measure with a warning message so that the user can consider a lower bit size for any of the dimensions in the array.

The algorithm for suggesting the bit size is to find the number of bits that can fit twice the number of positions currently in the dimension, starting with 2 bits. The results of the analysis are output to a file. Below is an example of this output:

Hierarchy	Dimension	Size	BitSize	Capacity	DPM-Enabled
clnd	day	3647	13	8192	No
clnd	week	522	11	2048	No
loc	str	22	6	64	Yes
loc	dstr	12	5	32	No
prod	sku	488	10	1024	Yes
prod	stco	258	10	1024	Yes

Note: The following measures have a key size greater than 64 bits. For better performance, it is recommended that you reduce the bit size of one or more dimensions in the base intersection of these measures so that the key size is within 64 bits.

Measure Name	Base Intersection	Key Size
ACT	www_xxx_yyy_zzzz	67

Usage for -suggestDimBitSize

```
convertDomain -src domainPath -suggestDimBitSize outputFile
```

Table 2–4 provides descriptions of the arguments used by the convertDomain utility with -suggestDimBitSize.

Table 2–4 Arguments Used by the convertDomain Utility with suggestDimBitSize

Argument	Description
-src <i>domainPath</i>	Path to the domain being analyzed.
-suggestDimBitSize outputFile	Location in which to create output file containing bitSize recommendations.

upgradeDomain

The upgradeDomain utility is used to upgrade only the RPAS version of the domain. It does not update the configuration or any other aspects of the domain itself.

upgradeDomain Usage

```
upgradeDomain -d domainPath [OPTIONS]
```

Table 2–5 provides descriptions of the arguments used by the upgradeDomain utility.

Table 2–5 Arguments Used by the upgradedomain Utility

Argument	Description
-d <i>path</i>	Path to the domain being upgraded.
-verbose	Shows the detail about each change that is applied to the domain.
-n	Reports which changes are applied (without applying the changes).
-purgeWorkbooks	Purges all existing workbooks and clears the workbook batch queue.
-ignoreSharedNames	Allows upgrade even if dimensions and hierarchies share names.
-apptag	Indicates the application and version associated with this upgrade. Parameter must be APP:VERSION.

Note: An administrator can also run the `domaininfo` utility to verify the upgrade process, as shown below.

```
domaininfo -d pathtodomain -domainversion
```

fixDomain

The `fixDomain` utility is used to repair many types of corruption that can occur within an RPAS domain. This utility can be used to perform a number of tasks, which are described below. In all cases, `fixDomain` can execute in two modes. The default mode is to perform an examination to detect problems. The second mode, which is performed if the `-commitChanges` argument is passed in the call to `fixDomain`, is to take action to correct problems that are detected.

Remove inconsistent measure attributes

It is possible for inconsistencies in measure attributes between the master and local domains of a global domain to interfere with operations performed by RPAS. By executing `fixDomain` with the `CleanMeasureAttrs` task, users can analyze a domain for such discrepancies and correct them by removing attributes that appear in only a subset of the domains.

Remove partially unregistered measures

It is possible for the process of measure un-registration to exit without completely removing the measure from all local domains of a global domain. By executing `fixDomain` with the `CleanMeasure` task, users can analyze a domain for partially unregistered measures and correct them by completing the un-registration process and completely removing these measures from the domain.

Repair discrepancies in hierarchy information

It is possible, in integer-indexing domains, for the hierarchy and dimension information of an array to become out of synch with the information of the domain. It is also possible for this information to become out of synch between the local domains of a global domain. By executing `fixDomain` with the `RepairHierarchy` task, users can analyze a domain for desynchronized dimension information and correct it when found by updating arrays or local domains with the dimension information of the master domain.

Repair unlinked arrays

In pre-integer indexing domains, there may be discrepancies between the dimension information in an array and the dimension information of the domain. When these corrupted domains are converted to integer indexing, these arrays are converted to rely on internal dimension information and cannot be resolved against the dimension information of the domain. By executing `fixDomain` with the `FixArrays` task, users can analyze a domain to detect these unlinked arrays and correct them by rebuilding the array using the dimension information of the domain.

Repair metadata inconsistencies

System dimensions and related metadata define the basic structure of a domain. Any discrepancies within them can cause serious problems during domain operations. By executing `fixDomain` with the `RepairMetadata` task, users can analyze a domain for any discrepancies in system dimensions. Depending on the type of information that contains a discrepancy, `fixDomain` may or may not be able to resolve the issue.

Hierarchy and Dimension Information: In general, this type of discrepancy is a result of a domain that did not build correctly. Due to the severity of this type of problem, fixDomain cannot resolve such issues.

Measure Information: It is possible for measure information to appear in some but not all the metadata arrays of a domain. Depending on which arrays do contain information about a measure, fixDomain may be able to correct the problem. Otherwise, the measures with incomplete information will be removed from the domain and must be re-registered.

Workbook Information: It is possible for information about workbooks, workbook templates, and workbook template groups to appear in some but not all of the metadata arrays of the domain. Depending on which arrays do contain information about the workbook, template or group, fixDomain may be able to correct the problem. Otherwise, the workbooks, templates or groups with incomplete information will be removed from the domain and must be re-registered.

User Information: It is possible for information about RPAS users and user groups to appear in some but not all of the metadata arrays of the domain. Depending on which arrays do contain information about the user or group, fixDomain may be able to correct the problem. Otherwise, the user or group will be removed and must be re-registered.

Note: The RepairMetadata task is only supported in versions of 13.2.3.x after 13.2.3.36. It cannot be used in RPAS version 13.3 and later. It is described here to support the upgrade and conversion process. See "[Upgrading From Any Release of RPAS Prior to 13.3](#)" for more information on the convertDomain process.

Repair corrupted change arrays

In pre-integer indexing domains, change arrays can get out of synch with the measure metadata of a domain. By executing fixDomain with the RepairChangeArrays task, users can analyze a domain for any discrepancies between change arrays and measure metadata and repair them, based on the measure metadata of the domain.

Note: The RepairChangeArrays task is only supported in versions of 13.2.3.x after 13.2.3.36. It cannot be used in RPAS version 13.3 and later. It is described here to support the upgrade and conversion process. See "[Upgrading From Any Release of RPAS Prior to 13.3](#)" for more information on the convertDomain process.

fixDomain Usage

The following arguments are used with fixDomain:

```
fixDomain -d pathToDomain -task CleanMeasure [-commitChanges] [-loglevel level]  
fixDomain -d pathToDomain -task RepairHierarchy -hier hierName [-commitChanges]  
[-loglevel level]  
fixDomain -d pathToDomain -task FixArrays [-db dbName | -array arrayName]  
[-commitChanges] [-loglevel level]  
fixDomain -d pathToDomain -task RepairMetadata [-dim dimName] [-commitChanges]  
[-loglevel level]  
fixDomain -d pathToDomain -task RepairChangeArrays [-commitChanges] [-loglevel  
level]
```

The following table provides descriptions of the arguments used by the fixDomain utility.

Table 2–6 Arguments Used by the fixDomain Utility

Argument	Description
-d pathToDomain	Indicates the location of the domain begin repaired.
-d CleanMeasureAttrs	Detects and removes measure attributes that do not exist in the master and all local domains of a global domain.
-task CleanMeasure	Detects partially unregistered measures and completes their un-registration.
-task RepairHierarchy	Detects and repairs inconsistencies in hierarchy information in arrays within the master and local domains of a global domain.
-hier <i>hierName</i>	Used in conjunction with the Repair Hierarchy task. Specifies the hierarchy to correct within arrays.
-task FixArrays	Detects and corrects unlinked arrays.
-db <i>dbName</i>	Used in conjunction with the FixArrays task. Specifies the database to examine or repair. The FixArrays task requires the use of either the -db or the -array argument, but not both.
-array <i>arrayName</i>	Used in conjunction with the FixArrays task. Specifies the array to examine or repair. The FixArrays task requires the use of either the -db or the -array argument, but not both.
-task RepairMetadata	Detects and repairs metadata discrepancies in system dimensions. Note: the RepairMetadata task is supported only in versions of 13.2.3.x after 13.2.3.36. It cannot be used in RPAS version 13.3 and later.
-dim <i>dimName</i>	Used in conjunction with the RepairMetadata task. Specifies the system dimension to examine or repair. This is an optional argument. If not supplied, fixDomain will examine all system dimensions.
-task RepairChangeArrays	Detects problems in change arrays within the domain and repairs them by relying on the measure metadata of the domain. Note: the RepairChangeArrays task is supported only in versions of 13.2.3.x after 13.2.3.36. It cannot be used in RPAS versions 13.3 and later.
-commitChanges	If the -commitChanges argument is not supplied to any call of fixDomain, it will only examine the domain and report detected problems. The fixDomain utility only modifies the domain to repair changes if the -commitChanges argument is specified.
-loglevel <i>level</i>	Sets the logging verbosity level. Supports use of all, profile, debug, audit, information, warning, error and none as values of level.

Domain Administration

This chapter describes domain administration tasks such as configuring the Classic Client and using the DomainDaemon. It contains the following sections:

- [Configuring the Classic Client Using eConfigure](#)
- [DomainDaemon](#)
- [Losing a Client-Server Connection](#)

Configuring the Classic Client Using eConfigure

This section provides instructions for configuring the RPAS Classic Client with the eConfigure utility on a local computer using a Microsoft Windows operating system.


The eConfigure utility is a Windows application that configures the client-server communication for RPAS. The eConfigure utility lets you specify communication parameters and produces a file that is used as input to the client. These files must be in the FCF (Foundation Configuration File) format/extension. The files contain the necessary information for the client to start up the communication with the server. These files can be stored on the client machine or on the network.

When the client is executed, a file named Foundation.FCF is expected in the same directory. If the file has a different name or if it is stored somewhere on the network, the path to this file must be passed in as an argument to the client.

The eConfigure utility consists of a menu bar, a main view, and the advanced settings dialog window. Passwords saved in the FCF file are encrypted. To launch eConfigure, double-click the eConfigure.exe file, which is by default located in the root directory of the RPAS Classic Client.


The eConfigure Menu Bar

The files produced by eConfigure may contain multiple connections. Each connection is specific for a server with certain communication settings. Connections need to have unique descriptions, and they can be added and deleted using the menu bar.

Figure 3–1 The eConfigure Menu Bar

The eConfigure Main View

The main view has the basic connection parameters, as shown in [Figure 3–2](#).

Figure 3–2 The eConfigure Utility

- **Database Server:** The hostname or the IP address of the server, for example, `atlddev03` or `10.2.1.23`. This value must be `localhost` when running the RPAS Server on a Windows machine.
- **Daemon Port:** The port number on which the domain daemon is listening. This must be an integer between 1025 and 65535 (for example, 55278).
- **Oracle Wallet for SSL:** The Oracle Wallet stores the SSL credentials. You must provide the path where the wallet is stored. This information is only required if one-way SSL is used. If SSL is disabled or if anonymous SSL is being used, then the wallet is not required. For more information, see [Chapter 4, "SSL."](#)
- **Domain:** The name of the domain that is displayed to the user when logging in. Select a domain from the list or type the name of a new domain and click **Add Domain**. You can delete a domain from the list by selecting it and then clicking **Delete Domain**.

- **Domain Path:** The full path to the directory containing the domain, for example, /root/testenv/domain/Sample_Project
- **User:** Provide the user ID if you do not want to force the user to provide it when logging in. The user ID must be defined in the associated domain. Note that the password cannot be entered here. The user must enter the password manually each time the RPAS client starts the connection.

Advanced Settings Dialog

Click the **Advanced** button shown in [Figure 3-3](#) to access the Advanced Settings dialog window.

Figure 3-3 Advanced Settings Dialog

- **User:** The database user that is used by the client if a domain-specific user has not been entered, for example, adm.
- **Database Port Range Start and End:** Port range is used to specify the range of ports on which the RPAS Server processes is started by the DomainDaemon (the rpasDbServer processes). The port Start and port End fields are the lower and upper limits of this range, respectively.
These fields must be integers between 1025 and 65535, which are also the default values if values are not specified, for example, Start: 40000, End: 45000.
- **Compression Threshold:** The number of bytes above which client and server are using compression. Only advanced users should manipulate this number.
- **Web Tunneling:** The configuration of Web tunneling.
- **Proxy Settings:** The configuration of the RPAS Classic Client to support a proxy server is not completed in this utility.

DomainDaemon

The RPAS DomainDaemon is a process that is used to enable the communication channel between RPAS Clients and RPAS domains.

The DomainDaemon runs on the server side and waits for requests from RPAS Clients on a given port. Once the DomainDaemon receives a request from a client, it starts a server process that the client connects to. From this point, the client and server communicate directly. The system administrators may choose to have one single DomainDaemon process for all of the users, or they may choose to have separate processes per domain, per enterprise, and so on.

The DomainDaemon is installed in the [RPASDIR]/bin directory. [RPASDIR] represents the full path to the directory where the RPAS Server is installed.

Note: On rare occasions, the Domain Daemon can throw the following exception when SSL is enabled: Failed to initialize Oracle SSL environment (nerror 43120: NZERROR_LX_ERROR).

This is caused by an unmatched language setting of the operating system. Setting the environment variable NLS_LANG using the following command before starting the Domain Daemon resolves this problem.

```
export NLS_LANG=AMERICAN
```

DomainDaemon Usage

Table 3–1 provides descriptions of the arguments used by the DomainDaemon utility.

Note: DomainDaemon commands that talk to a running DomainDaemon (stop, ping, stopActiveServers, and so on) now need to specify the wallet information. This applies to SSL types 1, 2, and 4, but not 3. This can be set in the environment variable RPAS_SSL_WALLET to avoid having to pass it on the command line.

Table 3–1 DomainDaemon Arguments

Argument	Description
-version	Prints the RPAS version, revision, and build information of the utility.
-start	Starts a DomainDaemon on the specified port.
-port <i>portNum</i>	Defines the <i>portNum</i> , which must be between 1025 and 65535 (inclusive). If <i>portNum</i> is set to auto, it will find any free port.
-loglevel	Enables additional logging. Note: The -debug option has been deprecated. Instead of using -debug, use -loglevel to add additional logging.
-timeout <i>milliseconds</i>	Specifies the number of milliseconds to wait for the server to start. A value of -1 means no timeout.
-server <i>serverProgramName</i>	Specifies the name of the RPAS database server program. Defaults to RpasDbServer.
-no_auto_add	Disables the registering of domains in response to client requests to start a RPAS database server.
-stop	Stops the DomainDaemon on the specified port.
-ping	Reports the status of a DomainDaemon process.
-showDomains	Shows all domains managed by this daemon.

Table 3–1 (Cont.) DomainDaemon Arguments

Argument	Description
-add pathToDomain	Adds the specified domain to the list of domains managed by a DomainDaemon.
-activate <i>pathToDomain</i>	Reactivates a previously deactivated domain. Specify the port number and the complete path to the domain.
-deactivate <i>pathToDomain</i>	Marks a domain as temporarily unavailable. Deactivating a domain also terminates all user sessions in that domain. A message will be displayed in the client to notify users when this occurs. Domains are most commonly deactivated before beginning a routine nightly or weekly batch process. This ensures that no users make updates to the system during these processes. Specify the port number and the complete path to the domain.
-remove <i>pathToDomain</i>	Removes the specified domain from the list of domains managed by a DomainDaemon.
-showActiveServers	Shows all active server processes. Specifying a port number is required. For each active server, the DomainDaemon shows the process ID, domain, and user ID.
-stopActiveServers	Stops all active servers. Specify a port number and a process ID.
-stopServer <i>processId</i>	Stops the server using the specified processed.
-stopUser <i>userId</i>	Stops the server using the specified userId.
-ssl	Specifies whether the messages exchanged between the Classic Client and the RPAS database server are encrypted using a password generated during the initial SSL handshake. Here are the options for SSL: 1: This option provides one-way authentication and encrypted data transfer. 3: This options provides encrypted data transfer without authentication.
-CipherSuite	Specifies the encryption algorithms used by SSL with the -cipherSuite field. A cipher suite consists of three components: - Asynchronous: The asynchronous algorithm is used in the SSL handshake to validate the client/server certificates and exchange the session key. - Synchronous: The synchronous algorithm uses the session key to encrypt data transfers between the client and server. - MAC: The MAC is a hash of the message contents appended to the end of the message to validate that the contents of the message were not altered. Because cryptography requirements evolve over time, this list of supported algorithms will change from release to release. o see which cipher suites are supported by your version of RPAS, use the -listCipherSuites option. If -cipherSuite is not specified, a default suite will be provided. The default suite will be whichever one is considered the most secure at the time of release. Conversely, if at the time of release a supported cipher suite is considered insecure, RPAS will deprecate it. Using a deprecated cipher suite will result in a warning being logged when starting DomainDaemon

Table 3–1 (Cont.) DomainDaemon Arguments

Argument	Description
-listCipherSuites	Lists the cipher suites supported by the current version of RPAS. Deprecated cipher suites are marked with a (X), and the default suite is marked with a (*). Because the list of available cipher suites can vary based what type of SSL you are using, you must specify the -ssl argument.
-wallet	-wallet Specifies the location of the server wallet folder. The wallet contains the RPAS server's private key and certificate and the client's certificate. Format: -wallet file:<path to the server wallet folder> Example: -wallet file:d:\wallets\server [on Windows] on Unix: -wallet file:/u00/admin/wallets/server The RPAS_SSL_WALLET environment variable can be used to specify the wallet location in place of this argument.

Starting the DomainDaemon

To start the DomainDaemon, execute the DomainDaemon utility. The port number where the DomainDaemon is running must be passed in as an argument. The port number must be between 1025 and 65535. If **auto** is specified instead of a number, the DomainDaemon is started on any available port.

Issuing the following command from a UNIX shell starts a DomainDaemon on port 55278:

```
DomainDaemon -port 55278 -start
```

If the command is successful, the DomainDaemon returns the following message:

```
Daemon listening on port '55278'.
```

Note: The proper environment variables must be set in order to execute the DomainDaemon. See the ["Verify the Environment Variable Settings"](#) for more information about setting RPAS_HOME and updating the PATH variable.

Monitoring the DomainDaemon

The -ping argument can be used to see whether a DomainDaemon is active. The port number must also be passed as an argument. If the DomainDaemon is active on the port, a message will be printed and the script will return true. Otherwise, the script will return false.

```
DomainDaemon -port 55278 -ping
DomainDaemon on port 55278 is alive.
```

Stopping the DomainDaemon

Use the -stop argument to stop the DomainDaemon running on a given port.

```
DomainDaemon -port 55278 -stop
```

Losing a Client-Server Connection

The connection between the RPAS Client and the RPAS Server can be lost for any number of reasons, but occurs most commonly when the user's computer crashes or the network connection is lost.

If an RPAS Client-Server connection is lost, the user's work is guaranteed to be saved up to the last calculation; all deferred calculations are lost. Upon subsequent login, the user can access either the last version of the original workbook explicitly saved by the user or the auto-saved version of the workbook that was being worked on when the connection was lost.

If the user tries to log in after recently losing the connection or from a different instance of the RPAS client, the user is prompted to either terminate all existing sessions and start a new session, or start a new concurrent user session. If the user chooses to terminate the existing session, the RPAS server gracefully terminates the existing session and then logs the user in and starts up a new session.

Graceful Termination of the Existing Session

If a connection has been lost and the previous session has not timed out or a session for the user is running from a different instance of the RPAS client, the user can use the RPAS client to gracefully terminate the existing session and log into a new session. Graceful termination includes the completion of any pending processes and custom menu or rule group processing in the existing session, followed by an auto-save of the workbook and subsequent termination of the server process.

Graceful termination of an existing session can take an arbitrarily long time. This time is equal to the sum of the time taken to complete any running calculations and the time taken to save the workbook.

When the RPAS server times out waiting on a request from the RPAS Client, it gracefully terminates, irrespective of whether the user tried to log in again.

If for any reason the client server connection is lost before the new login occurs, RPAS asks users if they want to resume the previous session or terminate it and start a new one. However, it will not start another session for the user until the user regains connectivity and tries to log in again.

There is a limit on the number of login sessions per user. By default, the maximum number of concurrent login sessions per user is five. This limit can be altered by an administrator using the Security Administration workbook.

Autosave of Workbooks

When an RPAS Server session terminates automatically, either due to re-login or a server timeout, it auto-saves the workbook that is currently open by the session. An auto-saved workbook includes all of the user's work up to the last calculation. If the user had some pending calculations, that is, edits were made in the client but calculation was not performed, all the edits are lost.

When a workbook is auto-saved, the original workbook is kept in the same state it was in the last time it was explicitly saved by the user. A new workbook is created for the user with the name of the original workbook suffixed with '_autosave' and added to the user's Most Recently Used list. Upon subsequent login, the user can view both the original workbook and the auto-saved workbook in the workbook list.

RPAS administrators can impose a limit on the number of workbooks a user can create for a particular workbook template. When such a limit has been imposed, the auto-save feature allows the user to exceed that limit by one. This allows the user to

operate at the limit without the fear of losing any work because of a connection failure or computer crash. However, when the limit is exceeded due to an auto-save, the auto-save feature is disabled for the user for the workbook template on which the limit has been imposed. The auto-save feature is disabled until the user deletes one or more of the workbooks for that template in order to bring the user at or below the limit. It is not required that the user delete the original or the auto-saved workbook. After the user is back at or below the limit, the auto-save feature is automatically enabled. Note that if a user exceeds the limit, the RPAS Client will inform the user of this situation every time the user attempts to open a workbook for the given template.

Secure Socket Layer (SSL), a protocol for securing network connections, is used by RPAS to provide secure communication between RPAS client and server processes. RPAS Classic Client supports one-way SSL and SSL without authentication.

This chapter contains the following:

- [Introduction](#)
- [One-Way SSL](#)
- [SSL Without Authentication](#)

Introduction

In order to use SSL with RPAS, you must configure both the server and the client. For one-way and two-way SSL, this involves setting up key stores and trust stores used to manage certificates. The key store contains a private key and its corresponding public certificate chain. The trust store contains trusted public certificates and certificate chains. SSL without authentication does not require digital certificates on either the client side or the server side.


In order to enable the use of SSL on RPAS, you must use the Domain Daemon to set up the SSL type and the Oracle wallet location (which is optional for certain SSL types).

Oracle Wallet is managed by a command line utility called orapki, which is included in the RPAS installation.

A Certificate Authority (CA) certificate is the public certificate from the CA that issues the server certificate.

One-Way SSL

In one-way SSL, the client asks to verify the identity of the server. The server does not ask to identify the client, but does require application-level authentication.

Figure 4–1 RPAS One-Way SSL Configuration

One-way SSL only authenticates the server on the client side. Only the server needs to have a key store for its private key and digital certificate. RPAS server uses Oracle Wallet for this purpose.

The trust store on the client side is an Oracle wallet for the Classic Client. The trust store needs to contain the server CA certificate so that the client can recognize and authenticate the server certificate.

Creating a Self-Signed Root Certificate

A self-signed root certificate can be used instead of a third-party CA certificate (for example, for testing). The same root certificate can be used to generate both server and client certificates.

Note: Under a Cygwin environment on Windows, the path of the wallet must be in Windows format, not Cygwin format. For example, it must be C:/wallets/root, not /cygdrive/c/wallets/root.

To create a self-signed root certificate, complete the following steps:

1. Create an Oracle wallet.

```
orapki wallet create -wallet {root_wallet_directory} -pwd {root_wallet_password}
```

Example:

```
orapki wallet create -wallet C:/wallets/root -pwd rootpass1
```

2. Generate a private key and a self-signed root certificate.

```
orapki wallet add -wallet {root_wallet_directory} -keysize {key_len} -dn {root_dn} -self_signed -validity {validity_days} -pwd {root_wallet_password}
```

Example:

```
orapki wallet add -wallet C:/wallets/root -keysize 2048 -dn "cn=rpas_qa_
ca,dc=us,dc=oracle,dc=com" -self_signed -validity 3650 -pwd rootpass1
```

3. Export the root certificate chain to a file for later use.

```
orapki wallet export_trust_chain -wallet {root_wallet_directory} -certchain
{root_cert_chain_file} -dn {root_dn} -pwd {root_wallet_password}.
```

Example:

```
orapki wallet export_trust_chain -wallet C:/wallets/root -certchain
C:/wallets/root_chain.txt -dn "cn=rpas_qa_ca,dc=us,dc=oracle,dc=com" -pwd
rootpass1
```

Configuring the RPAS Server

To configure the RPAS server for one-way SSL, you must set up a key store and then start the Domain Daemon.

Setting Up the Key Store

To set up the key store, complete the following steps:

1. Create an Oracle wallet with auto login. This type of wallet can be read without a password. The default file permissions for the wall are 0300 (accessible by owner only).

```
orapki wallet create -wallet {server_wallet_directory} -auto_login -pwd
{server_wallet_password}
```

Example:

```
orapki wallet create -wallet C:/wallets/server -auto_login -pwd serverpass1
```

2. Generate a private key.

```
orapki wallet add -wallet {server_wallet_directory} -keysize {key_len} -dn
{server_dn} -pwd {server_wallet_password}
```

Example:

```
orapki wallet add -wallet C:/wallets/server -keysize 2048 -dn "cn=rpas_qa_
server,dc=us,dc=oracle,dc=com" -pwd serverpass1
```

3. Export a certificate request.

```
orapki wallet export -wallet {server_wallet_directory} -dn {server_dn} -request
{server_cert_req_file} -pwd {server_wallet_password}
```

Example:

```
orapki wallet export -wallet C:/wallets/server -dn "cn=rpas_qa_
server,dc=us,dc=oracle,dc=com" -request C:/wallets/server_cert_req.txt -pwd
serverpass1
```

4. The certificate request must be digitally signed before it can be used. If the deployment of the RPAS Fusion client and the RPAS server are both under the control of the customer, then a self-signed SSL certificate may be used as the signer. Otherwise, the certificate request must be sent out for signing by a well-known certificate authority.

A self-signed SSL certificate used in this manner may be referred to as a self-signed root certificate or as a private certificate authority.

This is the method for creating a self-signed root certificate:

1. Generate a private key for the root certificate:

```
orapki wallet add -wallet {root_wallet_directory} -keysize {key_len} -dn  
{root_dn} -self_signed -validity {validity_days} -pwd {root_wallet_  
password}
```

Example:

```
orapki wallet add -wallet C:/wallets/root -keysize 1024 -dn  
"cn=root,dc=us,dc=example,dc=com" -self_signed -validity 3650 -pwd  
rootpass1
```

2. Export the root certificate chain to a file:

```
orapki wallet export_trust_chain -wallet {root_wallet_directory} -certchain  
{root_cert_chain_file} -dn {root_dn} -pwd {root_wallet_password}
```

Example:

```
orapki wallet export_trust_chain -wallet C:/wallets/root -certchain  
C:/wallets/root_chain.txt -dn "cn=root,dc=us,dc=example,dc=com" -pwd  
rootpass1
```

In the above example, the file containing the root certificate chain is
c:/wallets/root_chain.txt.

- 5.** If using an external root certificate, create the server certificate signed by a third-party CA. (Alternatively, see step 6.) To do this, send the server certificate request file ({server_cert_req_file}) to the CA. The CA sends back a certificate ({server_certificate_file}) for the server, along with the public certificate ({root_cert_chain_file}) of the CA.
- 6.** If using a self-signed root certificate, the server certificate file is signed as follows:

```
orapki cert create -wallet {root_wallet_directory} -request {server_cert_req_  
file} -cert {server_certificate_file} -validity {validity_days} -pwd {root_  
wallet_password}
```

Example:

```
orapki cert create -wallet C:/wallets/root -request C:/wallets/server_cert_  
req.txt -cert C:/wallets/server_cert.txt -validity 3650 -pwd rootpass1
```

The file c:/wallets/server_cert.txt is the signed server certificate, in the form of a file.

- 7.** Import the CA into the server wallet. Note that the public certificate chain may contain more than one certificate. In such cases, the certificates must be imported one by one, starting at the top of the chain.

```
orapki wallet add -wallet {server_wallet_directory} -trusted_cert -cert {root_  
cert_chain_file} -pwd {server_wallet_password}
```

Example:

```
orapki wallet add -wallet C:/wallets/server -trusted_cert -cert  
C:/wallets/root_chain.txt -pwd serverpass1
```

- 8.** Import the server certificate into the wallet.

```
orapki wallet add -wallet {server_wallet_directory} -user_cert -cert {server_  
certificate_file} -pwd {server_wallet_password}
```


Example:

```
orapki wallet add -wallet C:/wallets/server -user_cert -cert C:/wallets/server_
cert.txt -pwd serverpass1
```

Starting the Domain Daemon

To start the domain Daemon using one-way SSL, execute the following command, where <walletLocation> is the absolute path to the directory of the Oracle wallet.

```
DomainDaemon -port portNum -ssl 1 -wallet file:<walletLocation> -start
```

Configuring the Classic Client

To configure the Classic Client for one-way SSL, you must set up a trust store. To do this, create an Oracle wallet on the client. Then, import the CA into the wallet. Since the trust store does not contain any private keys, it can be shared by all Classic Clients. It can either be copied to every client machine or placed in a shared network location that is accessible by the client machines.

Once you have set up the trust store, you must configure the Classic Client connection so that the wallet can be used by the client.

1. Create an Oracle wallet with auto login.

```
orapki wallet create -wallet {client_wallet_directory} -auto_login -pwd
{client_wallet_password}
```

Example:

```
orapki wallet create -wallet C:/wallets/client -auto_login -pwd clientpass1
```

2. Import the CA from the server into the wallet. Note that the public certificate chain may contain more than one certificate. In such cases, the certificates must be imported one by one, starting at the top of the chain.

```
orapki wallet add -wallet {client_wallet_directory} -trusted_cert -cert {root_
cert_chain_file} -pwd {client_wallet_password}
```

Example:

```
orapki wallet add -wallet C:/wallets/client -trusted_cert -cert
C:/wallets/root_chain.txt -pwd clientpass1
```


3. Configure the Classic Client connection by entering the path for the location of the wallet into the Oracle Wallet for SSL dialog box, which is part of the EConfigure configuration utility (shown in [Figure 4-2](#)).

The wallet can be a directory on a local drive or on a network computer.


Examples:

A wallet on the C drive: C:/wallets/client.

A wallet on a shared directory called "shared" of a server called "server123::
//server123/shared/wallets/clients.

Figure 4–2 EConfigure Utility

Use the dialog box shown in [Figure 4–3](#) to browse for the wallet location.

Figure 4–3 Select Oracle Wallet Location Dialog Box

Once this configuration is complete, you can start the Classic Client as usual. It automatically detects the SSL type and uses the wallet specified in the connection configuration (Step 3 above).

When you are using SSL, you see a secure lock icon on the status bar at the bottom of the client window.

Script for Creating Self-Signed Wallets

The createSSLWallets.sh script is provided to automate the creation of Oracle wallets for self-signed certificates. The script creates both server and client wallets for RPAS server and classic client.

Usage for createSSLWallets.sh

The createSSLWallets.sh script is an interactive one that prompts the user for inputs. It does not take any command line arguments. The shell to start the script must have access to the orapki utility that is deployed under the bin directory of RPAS home.

User Input

The script requires the following user input:

The directory for the wallets

This is the location for the wallets. A directory named "rpaswallets" is created to store all wallets under this given directory. If a directory "rpaswallets" already exists, the script will stop and ask the user to remove that directory. This is to ensure that some old stale wallets do not become mixed up with the new wallets.

By default, the current directory is used if not given.

Organization name

The user is asked to enter an organization name that is used as part of the DN (Distinguished Name) of the certificate identities.

By default, the organization name is "My Company, Inc.".

The passwords for root wallet, server wallet, and client wallet

Each password must be entered twice. Basic validation is done to ensure that the user has entered matched passwords and that the password is not empty.

Output

The script displays every step it takes to create the wallets and certificates. There are a total of 17 steps. The orapki command line for each step is displayed with password masked out. The output of the orapki command is displayed to the user and thus any error message from orapki is also displayed.


If the script finishes successfully, it will display the locations of the root wallet, server wallet, and client wallet. Since RPAS Domain Daemon only takes absolute path for the wallet, these locations will show as absolute paths.

Error Handling

The script checks the output of every orapki command and stops when it encounters an error. Appropriate error message are displayed to the user.

RPAS Web Launch

To use RPAS Web Launch with one-way SSL support, you must specify the location of the wallet in the field labeled Oracle Wallet for SSL. Since this domain configuration is shared by all users, a network location must be used for the wallet or the wallet must reside at the same location across all client PCs.

Figure 4–4 RPAS Web Launch for Classic Client


The wallet location can also be specified in the <Client> section of the XML file used for In-context Launch, in the setting called sslWalletLocation. For example:

```
<Client>
  <version>13.3.1</version>
  <sslWalletLocation>//ntserver02/shared/wallets/client</sslWalletLocation>
</Client>
```

SSL Without Authentication

SSL without authentication (also known as anonymous SSL) does not require a digital certificate on either the client side or the server side.

An anonymous Diffie-Hellman (DH) key exchange protocol is used to generate and exchange a secret session key for use in the encryption and decryption of the traffic to achieve confidentiality.

No additional configuration is needed for RPAS Classic Client. It automatically recognizes the type of SSL and chooses the corresponding cipher suite.

SSL without authentication is considered to be a security risk. If you are using SSL without authentication, you will see the following warning message regarding this risk:

Warning: The client/server network communication is using SSL without authentication, which is prone to man-in-the-middle attacks.

RPAS Server Configuration

Use the following command to start the Domain Daemon using SSL without authentication:

```
DomainDaemon -port portNum -ssl 3 -start
```

User Maintenance

User administration is the process by which administrators add or delete authorized system users, create or delete user groups, and edit user profiles. These tasks are performed through completion wizards on the User Administration tab.

The following user administration procedures are discussed in this chapter:

- [Add User](#)
- [Add User Group](#)
- [Delete User](#)
- [Delete User Group](#)
- [Edit User](#)

These procedures can be performed through the RPAS Classic Client by accessing the User Administration tab in the New Workbook window ([Figure 5-1](#)).

In addition, the `usermgr` utility is also described in this chapter. This utility allows you to manage users and users groups through a command line interface. For more information, see "[Managing Users Using usermgr](#)".


After users and user groups are set up, access permissions to workbook templates and measures within workbooks can be assigned through the Security Administration workbook. The Security Administration workbook also supports modification of the label, default workbook template, and admin status associated with individual users.

Access the User Administration Tab

User administration workbooks are available only in a master domain of a global domain environment. To access the User Administration workbooks, do the following:

1. Select **New** from the File menu. The New dialog box appears.
2. Select the **User Administration** tab.

Figure 5–1 User Administration Tab in New Workbook Window



Add User

To add a user, perform the following steps:

1. From the File menu, select **New**.
2. Click the **User Administration** tab.
3. Select **Add User** and click **OK**.
4. The Add User Wizard appears.

Figure 5–2 Add User Wizard

The screenshot shows the 'Add User Wizard' dialog box. It contains the following fields and controls:

- User name:** Text box containing 'Wonka'.
- User label:** Text box containing 'WonkaE'.
- Default group:** Dropdown menu with 'admin' selected.
- Other groups:** List box containing 'admin'.
- Password:** Password field (masked with dots).
- Password verification:** Password field (masked with dots).
- Administrator:** Check box, checked.
- Force password change:** Check box, checked.
- Lock user account:** Check box, unchecked.
- Buttons:** 'Cancel', '<Back', 'Next>', 'Finish', and 'Help'.

Enter relevant information in the following fields:

- **User name:** Type the user name to be used for login.

Note: Each user name must begin with a letter and contain only alphanumeric characters and underscores. It cannot have spaces. User names are case-sensitive.

- **User label:** Type a label that describes the user (for example, the user's full name). This identifying label appears in various locations throughout the application. For example, labels appear on the Open dialog box to identify the owner of a given workbook and on some worksheets to specify which user performed a task.
 - **Default group:** Select the user group to which the user belongs.
 - **Other groups:** If a user will belong to more than one group, select the additional groups from the list in the **Other groups** field.
 - **Password:** Type a password for the user.
 - **Password Verification:** Re-type the same password.
5. If the user requires Admin status, check the **Administrator** box.

Note: Admin status enables users to perform the Format menu option Save Format/Admin, which creates new system-wide default styles for workbook templates. You can modify a user's Admin status at any time on the Users worksheet of the User and Template Administration workbook.

Note: Granting users Admin status gives them access to all workbook templates, but it does not automatically give them access to all workbooks.


6. If the user must change his or her password when logging on for the first time, check the **Force Password Change** box.
7. Check the **Lock User Account** box to temporarily disable the user's account.
8. Click **Finish** to add the new user to the database.

Workbook template and measure access rights can now be assigned to the user. To do this, access the User and Template Administration workbook. For more information, see "[Security Administration Workbook](#)".

Add User Group

User groups provide an intermediate level of security to workbooks that were created and saved by specific users. When new users are assigned to the system, they must be assigned to existing user groups. User groups should consist of individuals with similar job functions or responsibilities. In the Oracle Retail Predictive Planning Suite, the user group corresponds to the user's planning role.

1. Select **New** from the File menu.
2. Click the **User Administration** tab.
3. Select **Add User Group** and click **OK**.
4. The Add User Group Wizard appears.

Figure 5–3 Add User Group Wizard

The screenshot shows a standard Windows-style dialog box titled "Add User Group Wizard". Inside the dialog, there are two text input fields. The first is labeled "Group name:" and the second is labeled "Group label:". Both fields contain the text "RetailPlanners". At the bottom of the dialog, there are five buttons arranged horizontally: "Cancel", "<Back", "Next>", "Finish", and "Help". The "Cancel" button is highlighted with a blue border.

Enter information in the following fields:

- In the **Group Name** field, type a name for the group.


Note: Each group name must begin with a letter and contain only alphanumeric characters and underscores. It cannot have spaces. User group names are case-sensitive.

- In the **Group Label** field, type a descriptive label for the group. This label is displayed when referring to the group throughout RPAS.
5. Click **Finish** to add the user group to the database.

Delete User

If a user profile is no longer needed, it should be deleted from the system in order to maintain system security.

1. From the File menu, select **New**.
2. Click the **User Administration** tab.
3. Select **Delete User** and click **OK**.
4. The Delete User Wizard appears. The user names and labels for all users appear. Select the name of the user to delete.

Figure 5–4 Delete User Wizard


5. Click **Finish** to delete the user from the system.

Delete User Group

If a user group no longer exists, the group should be deleted from the system as soon as possible to maintain system security.

Note: Before you can delete a user group, you must remove all users from the group. For each user in the group, you must either delete the user or change the default user group assignment for the user.

1. From the File menu, select **New**.
2. Click the **User Administration** tab.
3. Select **Delete User Group** and click **OK**.
4. The Delete User Group Wizard appears. Select the user group to delete.


Figure 5-5 Delete User Group Wizard

5. Click **Finish** to delete the user group from the system.

Edit User

To edit a user's profile, perform the following steps:

1. From the File menu, select **New**.
2. Click the **User Administration** tab.
3. Select **Edit User** and click **OK**.
4. The Edit User Wizard appears. The User ID field displays the user names and labels for all users. Select the user to edit and click **Next**.

Figure 5–6 Edit User Wizard

5. The Edit User Wizard appears. Make the necessary changes to the user's profile. You can change everything except the User Name. See "[Add User](#)" for details.

Figure 5–7 Edit User Wizard, Page 2

User name:
 User label:
 Default group:
 Other groups:
 Change password:
 Password verification:
 Administrator: ☒
 Force password change: ☐
 Lock user account: ☐
 Cancel <Back Next> Finish Help

6. Click **Finish** to save the changes.

Managing Users Using usermgr

Use the usermgr utility to add and remove users and groups, copy user and group information to other domains, edit user and group information, and convert that information from XML files to database files and vice versa.

With the usermgr utility, you can create user and administrator accounts using the `-addGroup` command. This allows you to add many accounts at once. When you create accounts, you must supply a temporary password that is used for all user accounts and a temporary password that is used for all administrator accounts. These temporary passwords expire the first time the user or administrator logs in.

Note that you can only create temporary passwords; you cannot create real ones. This protects the security of the user and administrator accounts. For the same reasons, the default accounts `adm` and `usr` are no longer available.

To create users and groups, you must create a database file called `users.db`, which contains user and group information. After you have created and imported the `users.db` file into RPAS, that file contains all user, administrator, and group information, including the true passwords for each account. This file is an encrypted binary file that you cannot edit.

To edit any account information, you must convert the database file to an XML file using the `-convertDbToXml` command. This creates an editable XML file that contains all the information in the `user.db` file except for the true passwords of the accounts. Again, this is to ensure the security and safety of the account information. After you edit the XML file with the changes you need to make, you must convert it back to a

database file in order to import it into RPAS. To convert it, use the `-convertXmlToDb` command.

After the `users.db` file is created, it can be shared across multiple domains. To automatically import the user, administrator, and group information every time a domain is created, place the `users.db` file in the `rpasInstall` input directory of the domain.

usermgr Usage

```
usermgr -d domainPath -add [userName] -label [label] -group [groupName] {-admin}
usermgr -d domainPath -addGroup [groupName] -label [label]
usermgr -d domainPath -remove [userName]
usermgr -d domainPath -removeLabel [label]
usermgr -d domainPath -removeGroup [groupName]
usermgr -d domainPath -list
usermgr -d domainPath -print -user [userName]
usermgr -d domainPath -print -group [groupName]
usermgr -d domainPath -importDb {-replace}
usermgr -d domainPath -exportXml [path]
usermgr -d domainPath -exportDb [path]
usermgr -d domainPath -lock [userName]
usermgr -d domainPath -unlock [userName]
usermgr -convertXmlToDb -src [pathToXml] -dest [pathToDb]
usermgr -convertDbToXml -src [pathToDb] -dest [pathToXml]
```

Note: `-convertDbToXml`, `-convertXmlToDb`, and `-listHashAlgs` do not require a domain.

Table 5–1 provides descriptions of the arguments used by the `usermgr` utility.

Table 5–1 Arguments Used by the usermgr Utility

Argument	Description
<code>-d domainPath</code>	Specifies the path to a domain to add, remove, or get information about a user.
<code>-add userName</code>	Adds a user with a specified name. Use the other arguments specified in the usage to add those attributes for that user. If <code>-noPassword</code> is specified, then this user can only be logged in via the Fusion Client. Otherwise, this command will prompt for a temporary password for use with the Classic Client only.
<code>-label label</code>	Specifies the label of the user or group to add to the domain. Use this argument with <code>-add</code> and <code>-addGroup</code> .
<code>-group grp</code>	Specifies the user group of the user to add to the domain.
<code>-admin</code>	Specifies that the user to add to the domain has administrative rights.
<code>-addGroup groupName</code>	Adds a group with a specified name. Use <code>-label</code> to specify the label for the group.
<code>-remove userName</code>	Removes the user with the specified name from the domain.
<code>-removeLabel label</code>	Removes all users with this label.
<code>-removeGroup groupName</code>	Removes a group with this groupName.

Table 5–1 (Cont.) Arguments Used by the usermgr Utility

Argument	Description
-list	Lists all the users registered to the specified domain.
-listHashAlgs	Lists the different hashing algorithms supported in this release. The list will show an integer id, the details of the algorithm, and a status field indicating whether the algorithm has been deprecated 'X' or is the default '*'.
-showHashSettings	Shows the selected hashing algorithm and the selected number of hash iterations.
-hashAlg	The given algorithm id parameter defines the algorithm to use for new users or users whom update their password. Use the string 'default' to always chose the default algorithm.
-hashIter	The given iterations parameter is used to set the number of iterations the hash algorithm will used for new users or users whom update their password.
-forcePasswordChange	The given boolean parameter either 'true' or 'false' determine the setting of the password must change variable. If 'true' then the next login will request the password to be updated. If the optional -user parameter is provided the effect of this value will only affect that user.
-changePassword	For the given user application will prompt for the new password.
-lock <i>userName</i>	Locks the specified user. This prevents the user from logging in the domain.
-unlock <i>userName</i>	Unlocks the specified user. This allows the user to log in the domain.
-exportXml <i>path</i>	Creates an XML file that contains all users and groups in the selected domain. Passwords and password histories are not exported.
-exportDb <i>path</i>	Creates a database that contains all users and groups in the selected domain. This prompts you for new temporary passwords for admin and user accounts.
-convertDbToXml	Converts a user database to a user XML file. Passwords are not included in the conversion.
-convertXmlToDb	Converts a user XML file to a user database. This prompts you for temporary passwords for admin and user accounts.
-src <i>path</i>	Specifies the source file that is used in -convertDbToXml and in -convertXmlToDb.
-dest <i>path</i>	Specifies the destination file that is used in -convertDbToXml and -convertXmlToDb.

XML Schema

The XML schema contains information for all groups and users that are imported into the domain. Since this file can be edited, it does not contain any password information. Each group and user contains an XML attribute with the group or user name as well as the following inner tags:

Table 5–2 XML Schema

Outer Tag	Inner Tag	Description
GROUP	LABEL	The group's label.

Table 5–2 (Cont.) XML Schema

Outer Tag	Inner Tag	Description
USER	LABEL	The user's label.
USER	DFLT_GRP	The user's default group.
USER	OTHER_GRPs	A comma-separated list that contains all other groups that the user is associated with.
USER	ADMIN	If this value contains T, this user is an admin.
USER	LOCKED	If this value contains T, the user is locked when the file is imported.

Note: All XML tags must be in all caps.

Here is a sample users.xml file.

```
<?xml version="1.0" ?>
<VERSION>1.0</VERSION>
<USER_LIST version="2.0">
  <GROUPS>
    <GROUP NAME="grp1">
      <LABEL>Administrators</LABEL>
    </GROUP>
    <GROUP NAME="grp1">
      <LABEL>Group 1</LABEL>
    <?GROUP>
    <GROUP NAME="grp2">
      <LABEL>Group 2</LABEL>
    </GROUP>
    <GROUP NAME="grp3">
      <LABEL>Group 3</LABEL>
    </GROUP>
  </GROUPS>
  <USERS>
    <USER NAME="adm1">
      <LABEL>admin 1</LABEL>
      <DFLT_GRP>adms</DFLT_GRP>
      <ADMIN>T</ADMIN>
    </USER>
    <USER NAME="adm2">
      <LABEL>admin 2</LABEL>
      <DFLT_GRP>adms</DFLT_GRP>
      <ADMIN>T</ADMIN>
      <LOCKED>T</LOCKED>
    </USER>
    <USER NAME="usr1">
      <LABEL>user_1</LABEL>
      <DFLT_GRP>grp1</DFLT_GRP>
      <OTHER_GRPs>grp2</OTHER_GRPs>
    </USER>
    <USER NAME="usr2">
      <LABEL>user_2</LABEL>
      <DFLT_GRP>grp2</DFLT_GRP>
```

```

        <OTHER_GRP>grp1,grp3</OTHER_GRP>
    </USER>
    <USER NAME="usr3">
        <LABEL>user_3</LABEL>
        <DFLT_GRP>grp1</DFLT_GRP>
        <OTHER_GRP>grp1</OTHER_GRP>
    </USER>
</USERS>
</USER_LIST>

```

Use Cases

This section provides a discussion of common use cases for the XML and database files.

Exporting from an Existing Domain

You can export from an existing domain using `-exportDb` or `-exportXml`. Exporting is useful for sharing users with another domain or for creating backups.

Use `-exportXml` if you need to edit the users or groups. This can be useful when you are making bulk updates that apply to many users or groups. After you have made changes to the file, you must convert it back to a database by using `-convertXmlToDb` and then import the updated file using `-importDb -replace`.

Use `-exportDb` if you do not need to edit the users or groups. This method produces a binary file that is ready for import.

Importing into a Domain

You can import existing users.db by using `-importDb`. Importing is useful for bulk insertion or updates of users and groups.

By default, existing users and groups are skipped. However, if `-replace` is used, existing users are updated with the following information: user label, groups that the users belong to, admin status, and account lock status. Group labels are updated for existing groups.

Password information for existing users is not affected by the update. New users receive temporary passwords that are stored in users.db.

Note that new users may not be created if the following conditions exist:

- A group already exists with the user name.
- The user's default group does not exist.

New groups are skipped if the following conditions exist:

- A user already exists with the group name.

Note: If a users.db is placed inside the rpaInstall input folder, the users within that file are automatically imported upon domain creation.

Converting Between XML and Database

Since only users.xml can be edited and only users.db can be imported, it is necessary to convert between the two formats when you need to edit and import users and groups. For example, if you are creating users for the first time, you must first create those users in an XML file. Then you must convert that XML file to a database file

using `-convertXmlToDb` in order import the file. If you have an existing `users.db` that you need to edit, convert it to xml using `-convertDbToXml`.

When converting from an XML to a database file, follow these validation rules:

- The file must be a properly formatted XML file, otherwise the operation fails.
- Groups and users without a NAME attribute are skipped.
- Duplicate groups and users are skipped. Duplicate groups and users are those that share the same NAME attribute of an existing group or user within the same XML file. For instance, a new group called Sales cannot be created if a user or group named Sales already exists.
- Users without a DFLT_GRP field are skipped.

Note: If a user or group is skipped, a warning is logged. Since warnings are not included in the default log level, you should run this utility with warnings visible.


System Administration

This chapter describes the following system administration workbooks and wizards:

- [Auto Workbook Maintenance Wizard](#)
- [Workbook Batch Category Management](#)
- [Hierarchy Maintenance Workbook](#)
- [Password Policy Administration Workbook](#)
- [Security Administration Workbook](#)
- [Measure Analysis Workbook](#)

These workbooks and wizards are found in the Administration tab and Analysis tab of the New workbook window.

Figure 6–1 System Administration Workbooks and Wizards



Note: The Alert Manager Wizard is described in the *RPAS User Guide for the Classic Client*. The Translation Administration workbook is described in the "[Internationalization](#)" chapter.

In a global domain, some administration workbooks are available only in the master domain. These workbooks are the following:

- Hierarchy Maintenance
- Password Policy Administration
- Security Administration
- Translation Administration
- Workbook Batch Category Management

All other administration workbooks are available in both the local and master domains.

Auto Workbook Maintenance Wizard

The Workbook Auto Build feature allows users to have workbooks built by the wbmgr utility. Workbooks built in this way are added to the auto build queue. When workbooks are built in this manner, users are spared the processing time of making selections in the wizard and waiting for the workbook to build. The Workbook Auto Build feature works through the Auto Workbook Maintenance Wizard.

Add a Workbook to the Auto Build Queue

Workbooks in this queue are designated to be built automatically on a specified regular basis as part of the nightly batch run.

1. Select **New** from the File menu.
2. The New workbook window appears. Select the **Administration** tab.
3. Select **Auto Workbook Maintenance** and click **OK**.
4. The next wizard page appears. Select **Add a Workbook** and click **Next**.
5. Select a workbook template type, and click **Next**.
6. Select an owner for the workbook, and click **Next**.
7. The last page of the wizard appears.

Figure 6–2 Auto Workbook Maintenance Wizard

Enter the following information:

- **Build Label:** A relevant name for the auto workbook build.
 - **Build frequency (in days):** The frequency that the workbook should be build in days.
 - **Next Build Date:** The next date that the workbook should be built.
 - **Category:** Select the category of workbooks that this workbook should be assigned to. For more information about categories, see "[Workbook Batch Category Management](#)".
 - **Saved Access:** The level of access for this workbook. Choose User, Group, or World.
 - **Group:** Select the group that owns the workbook.
8. Click **Next** to initialize the wizard for the workbook template selected in step 5. The choices made are saved under the name specified for the Build Label.

Delete a Workbook from the Auto Build Queue

To delete a workbook from the Auto Build Queue, perform the following steps:

1. Select **New** from the File menu.
2. The New workbook window appears. Select the **Administration** tab.
3. Select **Auto Workbook Maintenance** and click **OK**.
4. The next wizard page appears. Select **Delete Workbooks** and click **Next**.
5. Select the workbook or workbooks to delete from the auto build queue.
6. Click **Finish** to delete the workbooks from the Auto Workbook Build queue.

Edit Workbook Settings

To edit the settings of an auto build workbook, perform the following steps:

1. Select **New** from the File menu.
2. The New workbook window appears. Select the **Administration** tab.
3. Select **Auto Workbook Maintenance** and click **OK**.
4. The next wizard page appears. Select **Edit Workbook Settings** and click **Next**.
5. Select **OK** and click **Finish** to build the Auto Workbook Maintenance workbook.
6. The Edit Auto Workbook Settings worksheet opens. Edit the settings as necessary. When finished, save and commit the workbook.

Figure 6–3 Edit Auto Workbook Settings Worksheet

	Category	Frequency (in days)	Next Build date	Owner	Save Group	Template	Workbook Name	World Access
PlannerWorkbooks	Planner Workbooks	1	1/30/2010	adm	Admin	Sales Planning	PlannerWorkbooks	☒

Workbook Batch Category Management

A category is defined as a group of related workbooks for batch processing purposes. As an administrator, you can create new categories that auto workbook queue entries and batch workbook refresh entries can be assigned to. Each entry can be assigned to only one category. By default, each entry is assigned to a master category called Default. If a category is deleted, the entries assigned to that category are reassigned to the Default category.


When setting up auto workbooks, users can assign a category to an auto workbook and then run the build of a category that a group of entries have been assigned to. This is useful because it allows users to build only the workbooks they need to work with. These categories can also be used when entering workbooks into the refresh queue with the wbbatch utility. Similar to building workbooks, users can refresh select workbooks based on the defined category. These categories are created with the Workbook Batch Category Management wizard or with the wbbatch utility. For more information about the assignment of categories, see "[Auto Workbook Maintenance Wizard](#)" and "[Managing the Workbook Batch Queue Using wbbatch](#)".

In a global domain, the workbook batch category collection is global and stored in the master domain. It is shared by all local domains.

Workbook Batch Category Management Wizard

The Workbook Batch Category Management wizard allows users to add and delete categories as well as edit the labels of the categories. To access this wizard, perform the following steps:

1. Click the **New** icon in the toolbar or the **New** option in the File menu. The New workbook window opens.
2. Select **Workbook Batch Category Management** and click **OK**.

Figure 6–4 Workbook Batch Category Management Wizard

The Workbook Batch Category Management Wizard opens.

Adding a Category

To add a category, select **Add Category** and enter a name and label for the category.

- **Name:** The category name is restricted to standard alphanumeric characters. It cannot contain spaces. This name is used when specifying the category in the wbbatch utility. After a category has been created, this name cannot be changed.
- **Label:** The category label is displayed on the workbook template wizard pages. It can be in any language and can contain spaces. Category labels are case sensitive and must be unique.

After you have entered a name and label, click **Finish**. The category is created.

Figure 6–5 Adding a Category

The screenshot shows a Windows-style dialog box titled "Workbook Batch Category Management Wizard". The main text inside says "Please choose a workbook batch category operation:". There are three radio button options: "Add Category" (which is selected), "Delete Category", and "Change Category Label". Under "Add Category", there are two text input fields: "Name:" with the value "AdminWorkbooks" and "Label:" with the value "Admin Workbooks". Under "Delete Category", there is a "Select" dropdown menu showing "Test". Under "Change Category Label", there is a "Select" dropdown menu showing "Test" and a "New Label:" text input field which is currently empty. At the bottom of the dialog, there are five buttons: "Cancel", "<Back", "Next>", "Finish", and "Help".

Deleting a Category

To delete a category, select **Delete Category** and select the category you want to delete from the list. Click **Finish**.

Figure 6–6 *Deleting a Category*

Workbook Batch Category Management Wizard

Please choose a workbook batch category operation:

☐ Add Category

Name:

Label:

☒ Delete Category

Select

Admin Workbooks
Planner Workbooks
Produce Workbooks

☐ Change Category Label

Select

New Label:

Cancel <Back Next> Finish Help

Changing a Category Label

To change the category label, select **Change Category Label**. Select the label you want to change and enter the new one. Click **Finish**.

Note: Category names cannot be changed.

Figure 6–7 Changing a Category Label

Workbook Batch Category Management Wizard

Please choose a workbook batch category operation:

☐ Add Category

Name:

Label:

☐ Delete Category

Select:

☒ Change Category Label

Select:

New Label:

Cancel <Back Next> Finish Help

Hierarchy Maintenance Workbook

Oracle Retail Predictive Solutions provide the ability to set up and maintain user-named and user-defined dimensions within hierarchies. Hierarchy Maintenance is the means by which custom-created dimensions within a hierarchy can be established and maintained through the application interface to meet individual business needs.

When Oracle Retail Predictive Solutions are installed, implementation scripts define the dimensions and hierarchical structures specific to the customer's organization. For example, the system can be built to recognize that SKUs roll up into styles, styles roll up into product classes, and so on within the product hierarchy. Occasionally, you might want to group products according to some ad hoc personal design to suit a particular business need. You can group arbitrary items in a hierarchy to use in functions such as forecasting, replenishment, and measure analysis. These user-defined groupings act as normal dimensional levels. In other words, they allow the user to roll data up from lower levels of aggregation along the hierarchical paths that you define.

For example, suppose experience has shown that the accuracy of forecasts for your top 50 products (A products) reflects the relative accuracy of all forecasts. Therefore, you would like to group elements within a user-defined dimension as the top 50 products by designating them 'A Products.' Then, when you select products in a wizard or look at data in a worksheet, you can change the rollup to your user-defined dimension to see your top 50 products grouped together.

Note: Your collection of 50 products may comprise elements from a wide range of product classes or departments, and your grouping scheme may have little to do with the normal dimensional relationships of these items in the product hierarchy.

The group of items you designate as 'A Products' may change over time as consumer preferences change. From this example, you see that user-defined dimensions can be used to create any ad hoc groupings to provide additional support in analyzing, selecting, or summarizing data in Demand Forecasting. The Hierarchy Maintenance interface allows you to change the nature of the groupings as required.

The number and names of user-definable dimensions are set by your company when an RPAS-based solution is initially installed. The positions within each dimension and their associated labels can be altered and maintained through the hierarchy maintenance process.

Remember that any hierarchy in RPAS can have user-defined dimensions within it as long as they are set by your company at the time of installation. The following examples refer to the product hierarchy, but other hierarchies can be maintained in the same way.

Hierarchy Maintenance Example

Suppose you want to designate SKUs in your product hierarchy as either A, B, or C products so that you can group these items together when you view information, such as forecasting, replenishment, or measure analysis reports.

To do this, you need to maintain a user-defined dimension that will allow you to map the SKUs to the various positions of your classification scheme (A, B, or C). The user-defined dimension used in the following example is named Product Status. To maintain this user-defined dimension, use the Hierarchy Maintenance Wizard.

Hierarchy Maintenance Wizard

The first step in maintaining hierarchies is to access the Hierarchy Maintenance Wizard. In this wizard, select the SKUs that will be mapped to the various positions of the user-defined dimension. Responses to prompts in the wizard are used to format a new Hierarchy Maintenance workbook.


Hierarchy Maintenance Worksheet

The Hierarchy Maintenance worksheet displays the position assignment fields for the selected custom dimension. Edit the cells associated with the custom dimension as required.

Returning to the example dimension Product Status, you want to classify each selected SKU in your workbook as an A Product, a B Product, or a C Product. This example provides only three positions, or values, in the Product Status dimension; however, you can enter any character string in an individual SKU's Product Status cell. This new string is treated as a separate user-defined grouping. If this is the first time a particular SKU has been mapped to the Product Status dimension, the label assigned to that SKU will not yet be defined. The Product Status field is automatically filled with 'Unassigned.'

Assign labels to each product with regard to the Product Status dimension. In [Figure 6–8](#), products that were previously unassigned are now designated as A, B, or C products.

Figure 6–8 Hierarchy Maintenance Worksheet




Product	Product Status
3001-01 Turtleneck Sweater 20-S	A Product
3001-02 Turtleneck Sweater 20-M	A Product
3001-03 Turtleneck Sweater 20-L	C Product
3041-07 Cashmere Sweater 42-S	B Product
3041-08 Cashmere Sweater 42-M	A Product
3041-09 Cashmere Sweater 42-L	C Product
8031-02 Sports Pique 17-M	C Product
8031-03 Sports Pique 17-L	B Product
8053-05 Sweatshirt 20-M	Unassigned
8053-06 Sweatshirt 20-L	Unassigned

Note: The Oracle Retail Predictive Solutions system is case-sensitive when a new position name (label) is entered in the Hierarchy Maintenance workbook. After the workbook is committed, the typing of the group name is not case-sensitive. For example, "B Product" can later be entered as "b product" after the "B Product" group label has been committed.

After making the A, B, or C Product designations for the selected SKUs, you must commit the workbook for any changes to take effect.

For this example, labels have now been assigned to the various positions within the Product Status dimension, and selected products in the product hierarchy have been classified with regard to the custom dimension. Demand Forecasting treats Product Status, a user-defined dimension, as a normal dimensional level within the product hierarchy.

[Figure 6–9](#) displays the results when you access a quick menu in a wizard and change the rollout to the Product Status dimension. The products shown here are classified according to the position values (A Product, B Product, or C Product) that were assigned while maintaining the Product Status dimension.

Figure 6–9 Product Status Dimension Results

Access Hierarchy Maintenance

1. Select **Open** from the File menu to bypass the Hierarchy Maintenance wizard, and open an existing Hierarchy Maintenance workbook, or select **New** from the File menu.
2. Select the **Administration** tab to display the list of Administration templates.
3. Select **Hierarchy Maintenance** and click **OK**.
4. Select the hierarchy to specify a user-defined dimension (for example, Product or Location). Only the hierarchies that have been set up to contain user-defined dimensions are represented here. Click **Next**.
5. Select the user-defined dimension to be updated. The number and names of available custom dimensions are set at installation. Click **Next**.
6. On the **Available** side of the selection wizard, choose the items to be mapped to positions within the custom dimension.
7. Click the right arrow button to move them to the **Selected** side.
8. After all items to appear in your workbook have been selected, click **Finish**.

Maintain a User-Defined Dimension Within a Hierarchy

Use this procedure to assign product or location items to custom-defined positions within a specialized dimension. Custom-created dimensions are distinct from those in the standard hierarchical roll-ups configured in the system implementation. You can

use these dimensions as you would normal Demand Forecasting levels, aggregating data along these new hierarchical paths.

1. Select **New** from the File menu.
2. Select the **Administration** tab to display the list of Administration templates.
3. Select **Hierarchy Maintenance**. Click **OK**.
4. Select the hierarchy to specify a user-defined dimension (for example, Product or Location). Only the hierarchies that have been set up to contain user-defined dimensions are represented here. Click **Next**.
5. Select the user-defined dimension to be updated. The number and names of available custom dimensions are set at installation. Click **Next**.
6. On the **Available** side of the selection wizard, choose the items to be mapped to positions within the custom dimension.
7. After all items to appear in your workbook have been selected, click **Finish**.
8. The Hierarchy Maintenance workbook is displayed. In the position assignment field for the custom dimension, assign a value to each product or location position in the workbook. Enter any text string in a cell. Each unique string will be treated as a separate user-defined position within the custom dimension.
9. Select **Commit Now** from the File menu to commit the changes to the master database. You can also save the workbook by selecting **Save** from the File menu.
10. To close the workbook, select **Close** from the File menu.

Password Policy Administration Workbook

Using the Password Policy Measures Settings worksheet, administrators can configure password complexity and settings in order to ensure the account security of users and other administrators. With this worksheet, administrators can set the required password complexity, the number of allowable password attempts, the expiration time of a password, and the length of time a user is locked out of the system after failed password attempts. The measures used to control these settings are described in [Figure 6–10](#) and [Table 6–1](#).

Note: The requirements set in the workbook are automatically applied when the user logs in. If a user's password does not meet the requirements, the user is prompted to change it.

Figure 6–10 Password Policy Measures Settings




Table 6–1 Password Policy Measures

Measure	Description	Allowable Values	Default Value
Password Attempts	The number of consecutive failed login attempts (due to an invalid password) before the account is locked. Choose 0 if you do not want to lock the account.	Choose from the set values in the pick list: 0, 1, 2, ..., 10	5
Password Expiration Time	The number of days that a password is valid for an account. After the time passes, the user can to log in, but is prompted to enter a new password.	1 or greater	90 (days)
Password Fail Delay	The number of seconds the server waits before replying to the client that a login failed (due to an invalid password).	0 to 600 seconds	1 second
Hash Algorithm	This is a drop-down list that contains all of the valid hash algorithms for this release. Additionally, the first field in the list will always be Default, which sets the measure to -1. Deprecated algorithms will be marked with a (X), and the default algorithm will be marked with a (*).	The allowable algorithms will be the hash algorithms supported by that release of RPAS. The membership of this list may change between releases.	The default value is the hash algorithm considered most secure at the time of the release of that version of RPAS.
Hash Iterations	This is an integer measure that stores how many iterations the hashing algorithm should be applied.	The minimum setting is 1. There is no maximum.	1000
Password Failure Lockout Time for Administrators	The amount of time that an administrator account is locked after consecutive failed login attempts. After this time passes, the account is unlocked and the administrator can attempt to login again.	Choose from the set values in the pick list: 15 min, 30 min, 1 hour, 1 day	15 minutes
Password Failure Lockout Time for Users	The amount of time that a user account is locked after consecutive failed login attempts. After this time passes, the account is unlocked and the user can attempt to login again.	Choose from the set values in the picklist: 0, 15 min, 30 min, 1 hour, 1 day Choose 0 if you do not want to unlock the account.	15 minutes
Password Minimum Length	The minimum number of characters that a password can contain. (The maximum number of characters is 31.)	1 to 31	6
Password Requires Mixed Case	If this measure is set to true, then the password must contain both a lowercase and an uppercase letter. (Passwords must always contain a letter.)	True, false	True

Table 6–1 (Cont.) Password Policy Measures

Measure	Description	Allowable Values	Default Value
Password Requires Number	If this measure is set to true, then the password must contain a numeric digit.	True, false	True
Password Requires Symbol	If this measure is set to true, then the password must contain a non-alphanumeric character.	True, false	False
Password Reuse Time	The number of days that must pass before a password can be reused for an account.	0 to 10,000 days. Select 0 to place no restrictions on reusing passwords.	365 days

Security Administration Workbook

The security model in RPAS includes workbook templates, workbooks, measures, and positions. The levels of security are defined as measure level, position level, and workbook level.

Security Overview

This section provides the basic information on the security model in RPAS.

User Logon Security

User accounts can be marked as locked by the domain administrator. This prevents the user from logging on with the RPAS Client. The account remains locked until the administrator re-enables the account. Account lockouts can be set or cleared by the domain administrator by using the User Management utility or the Edit User workbook.

Accounts may be configured to automatically lock out after a certain number of failed logon attempts. A domain administrator can configure the number of failed logon attempts and the duration of the lockout using the Password Policy Administration workbook.

Accounts may be marked as requiring the user to change the password. When this is set, users will be prompted to change their password the next time they logon. Users will not be able to proceed using the RPAS client unless they successfully change their password. This is useful for brand new accounts that are created with a stock password. The domain administrator can set or clear this setting using the User Management utility or the Edit User workbook.

Password expiration may be enabled for the domain. The domain administrator may set the number of days after which passwords expire. After this time passes, users will be prompted to change their password the next time the logon. Users will not be able to proceed using the RPAS client unless they successfully change their password.

A password reuse time can be set for the domain. This is often used in combination with password expiration to ensure users do not change their password to a recently used password after one expires. The domain administrator may set the minimum number of days that may pass before users can reuse a previous password using the Password Policy Administration workbook.

Measure Level Security

Measures have access rights; which are read-write, read-only, or denied. Measures that are read-write or read-only may be selected in the extra measures and insert measure dialogs. RPAS ensures that read-only measures are not editable by the user and the presence of read-only measures does not affect the ability to commit a workbook.

Measure security can be specified and changed through the Security Administration workbook. The Measure Rights worksheet allows Read Only, Deny, or Read/Write access to a measure to be specified for each user.

A workbook template can override the security of a measure, but it can only narrow the security of that measure. For example, a measure could have read-write access for a user and a template could specify that all users have read-only access to the measure when a workbook is built. However, if the measure security was read-only, the template could not expand the security of that measure to read-write. Measures that are explicitly made read-only by a workbook template will not be expanded to read-write access by RPAS.

Note: See the *RPAS User Guide for the Classic Client* for more information on the Measure Analysis workbook.

Position Level Security

Position Level Security allows access control for dimensions on a position-by-position basis. This capability is completely optional. If position level security is not explicitly defined and configured, all users in a domain have access to all positions in all hierarchies. After position level security is defined, access to a position can be granted or denied for individual users, users in a group, or for all users.

Position level security can be defined at levels (dimensions) at or above base (such as class in the product hierarchy) in any hierarchy other than calendar. As positions are added at a level/dimension lower in the hierarchy than where the position level security is maintained, access to those positions is automatically granted if a user has access to the parent position. In other words, if security is maintained at the subclass level, users are automatically granted access to all the SKUs in a given subclass if they have access to that subclass. This includes those added after security was established.

Exactly one dimension in each hierarchy can be defined as the security dimension for the hierarchy. If a security dimension is defined for the hierarchy, all dimensions in the hierarchy have position level security enabled, but position security is set at or above the designated dimension. For instance, if the class dimension is designated as the security dimension, an administrator can maintain access to positions in the class dimension or at any level above class. To specify the security dimension for a hierarchy, use the RPAS Configuration Tools or the `hierarchyMgr` utility.

After a security dimension is defined for a hierarchy, all users in the domain default to having access to all positions in any dimension in the hierarchy. Additionally, users automatically have access to newly added positions to a domain. Worksheets in the Security Administration workbook are used to control position access for individual users, user groups, or all users (referred to as world or default access). There are three worksheets in this workbook for each hierarchy with a defined security dimension. The default worksheet controls access to positions for all users (for instance, Prod Security Default); one worksheet controls access to positions by user group (for instance, Prod Security Group); and the last worksheet controls access to positions by individual users (for instance, Prod Security User).

Access must be granted at all levels for a user to have access to a position. This means that a position must have a value of true at the levels default/world, group, and user. [Table 6–2](#) provides details about how access is granted or denied based on all combinations of settings:

In the table, security is set by Position. Denied = False and Granted = True. Based on the settings for User, User Group, and World, the user is either granted or denied access, as shown in the Resulting Access column.

Table 6–2 Grant or Denial of Access by a Combination of Settings

User	User Group	World	Resulting Access
Denied	Denied	Denied	Denied
Denied	Denied	Granted	Denied
Denied	Granted	Denied	Denied
Granted	Denied	Denied	Denied
Denied	Granted	Granted	Denied
Granted	Denied	Granted	Denied
Granted	Granted	Denied	Denied
Granted	Granted	Granted	Granted

Position level security is used when a user selects positions in the wizard process before building a workbook. Only positions to which a user has access are available for selection in the 2-tree, which are then included in the build of the workbook.

Workbook Security

Currently, workbook access is either granted or denied. If users have been granted access to a workbook, they can open, modify, and commit the workbook. No distinction is made between read-write-commit, read-write, and read-only access. Workbook access is automatically granted to the user that built it, and it may be shared with multiple groups or the world.

Note: A user must have access to the workbook template in order to access the workbook, even if the workbook has world or group access rights.

Users with administrator status automatically have access to all workbook templates. By default, administrators have access to all workbooks that are saved with world access. If a workbook is saved with group access, administrators can only access the workbook if they are members of the default user group of the user who saved the workbook.

Another aspect of workbook security is the ability to set limits for the number of workbooks that a user can have saved at any given time. Limits can be set for a user per template, for a user group per template, or for a template for all users. The limits are evaluated in the above order, which means that a limit defined at user-template overrides any values defined at group-template or template. If the above limits are not defined, the default value is one billion.

The limits are checked when the workbook build process is initiated. When the limit is reached, an error message displays informing the user that the workbook build

process cannot complete because the limit has been reached. The message also lets the user know what that limit is. The wizard process then terminates.

Administrative users have full access to all workbook templates regardless of the access rights that other admin users may assign to them in the Security workbook. The administrative user can build the Security workbook to change the access right back, so the nominal assignment does not matter for administrative users.

Non-administrative users do not have access to Security template and User Administration template groups even if the administrator inadvertently assigns them access rights.

Security Administration Workbook

The Security Administration workbook is only available to system administrators. After users and user groups are created, the administrator may set up and maintain access permissions to workbook templates and measures within those workbook templates. This workbook allows the administrator to determine which templates individual users can access, as well as the measures that users can access while manipulating workbooks in the system. The user can also specify and restrict the measures that are available to be added to a given workbook template. Setting access permissions in this way provides a high degree of measure security, because users can be restricted to viewing and editing only certain relevant measures.

All administrative users have full access to all workbook templates regardless of the access rights that they were assigned in the Security workbook by other administrative users. The administrative user can build the Security workbook to change the access right back, so the nominal assignment does not matter for admin users.

The Security Administration workbook has the following worksheets:

- [Workbook Template Rights Worksheet](#)
- [Workbook Template Measure Rights Worksheet](#)
- [Measure Rights Worksheet](#)
- [Dimension Modification Rights Worksheet](#)
- [Position-Level Security Worksheets](#)
- [Workbook Template Limits Worksheets](#)
- [Max Domain Session Limit Worksheet](#)
- [Max User Session Limit Worksheet](#)

Security Template Administration also allows the administrator to modify the label, Admin status, and default workbook template associated with each user. You also access this workbook template to modify the labels associated with user groups, workbook templates, and workbook template groups. Using this workbook, the administrator can:

- Assign and modify access rights of each user to all workbook templates. User/template permissions are set in the Workbook Template Rights worksheet.
- Determine which optional measures are to be accessible through individual workbook templates. Template/measure permissions are set in the Workbook Template Measure Rights worksheet.
- Assign/restrict user access to individual measures. User/measure permissions are established in the Measure Rights worksheet.

Workbook Template Rights Worksheet

The Workbook Template Rights worksheet is for setting and maintaining access permissions of each user to specific workbook templates.

The worksheet contains a drop-down list for each available workbook template and user combination. To grant a user access rights to a workbook template, select one of the following options from the drop-down list for that workbook template:

- Denied
- Read Only
- Full Access

After changing a user's profile, the changes must be committed to the database in order for them to take effect.

The Read Only permission on a template applies only to actual workbooks created by the template. For templates that do not generate a workbook, but only run through a wizard process for other purposes, the Read Only permission for a user on that template will not prevent them from running through the wizard. This applies to standard RPAS templates, such as Add User and Delete User, but it may also apply to various application-specific templates.

Workbook Template Measure Rights Worksheet

The Workbook Template Measure Rights worksheet allows administrators to determine which registered measures will be available for optional inclusion in newly built workbooks.

When a measure is initially registered as a public measure, all templates default to having access to that measure. This means that it is possible for this measure to be added to a workbook template, even if it is not one of the standard measures displayed when a workbook of that type is built. Some new workbook wizards include a dialog that prompts users to select any additional measures to be included in the workbook build. By default, all newly registered measures are included on this list of available additional measures. The other method of inserting new measures into a workbook is with the Insert Measure command.

The Workbook Template Measure Rights worksheet is used to modify template/measure permissions, which allows only certain templates to optionally include specified measures in new workbook builds.

This worksheet contains a check box for each available workbook template and registered measure combination.

Measure Rights Worksheet

The Measure Rights worksheet allows the administrator to restrict user access to individual measures on a user-by-measure basis. User/measure permissions are initially determined by the system by integrating the current user/template and template/measure settings and applying the following rule: "A user cannot have access to any measure that is not available in at least one template to which the user has access."

Permissions can be made even more restrictive on a user by measure basis by using the Measure Rights worksheet to deny users access to measures that they would normally be permitted to edit.

The worksheet contains a drop-down list for each available user and registered measure combination. Three security options are available: Denied, Read-only,

Read/Write. Denied prevents the user from viewing data. Read-only allows the user to view the data. Read/Write allows the user to edit data values. However, a commit rule must be configured for a measure for data to be committed to the RPAS datastore.

A measure will have the security rights it had when it was inserted in the workbook. The change in measure security rights is only reflected in new workbooks when that measure is inserted.

Note:

- If a measure that has dependent measures is inserted into a worksheet, those dependent measures will also be inserted. If the dependent measures have denied measure access, they are still inserted into the workbook but are not visible to users.
 - The Measure Rights worksheet contains only public measures; that is, measures that can be optionally included in a worksheet, depending on choices made in a new workbook wizard. Measures that are registered as private measures will not appear in this worksheet. If there are no public measures available to be displayed in this worksheet, the worksheet will not be built.
-

Dimension Modification Rights Worksheet

The Dimension Modification Rights worksheet allows the administrator to determine which dimensions, if any, a user can modify. The worksheet contains a check box for each available user and dimension combination. A check mark in the cell indicates that the user is permitted to modify the specified dimension.

After changes are made to a user's dimension modification rights, they must be committed before they take effect.

Position-Level Security Worksheets

The position-level security worksheets are used to grant or deny access to positions for individual users, user groups, or all users. Position-level security is set for a specific dimension of a hierarchy (other than calendar). See the *RPAS Configuration Tools User Guide* for more information on setting position-level security dimensions.

For each hierarchy/dimension that has position-level security enabled (normally just a single hierarchy/dimension), there are three worksheets: one each for user, user group, and world/all users.

After changes are made to position-level security, they must be committed before they take effect.

Workbook Template Limits Worksheets

The Workbook Template Limits worksheets are used to limit the number of workbooks that the user can have saved. Limits can be set for a user per template, for a user group per template, or for a template for all users. The limits are evaluated in the above order, which means that a limit defined at user-template will override any values defined at group-template or template. If the above limits are not defined, the default value is 1 billion, but it is not displayed in the workbook.

The limits are checked when the user begins the workbook build process. If the limit has been reached, an error message appears that informs the user that the workbook build process cannot complete because the limit has been reached. The wizard process then terminates.

Max Domain Session Limit Worksheet

The Max Domain Session Limit worksheet is used to limit the number of user sessions that can be attached to a single domain by all users of that domain. The limit is set at the domain level. In a global domain environment, the same limit is applied individually to each local domain and the master domain.

This limit is checked during user login. If the limit has been reached, an error message appears to inform the user that the login has failed due to this limit being reached.

Max User Session Limit Worksheet

The Max User Session Limit worksheet is used to limit the number of concurrent user sessions that can be attached to a single domain by the same user at the same time. The limit is set per user so that admin can control the maximum number of concurrent sessions that are allowed for an individual user. In a global domain environment, the same limit is applied individually to each local domain and the master domain.

This limit is checked during user login. If the limit has been reached, an error message appears to inform the user that the login has failed due to this limit being reached.

Using the Security Administration Workbook

Note: These tasks are performed through the Security Administration workbook. This workbook is available to only system administrators.

Access Security Administration

1. From the main menu, select **File - New**. The New dialog box appears.
2. Select the **Administration** tab to display a list of workbook templates for Administration.
3. Select **Security Administration** and click **OK**.

Set or Modify User Access to Workbook Templates

1. From the File menu, select **New**.
2. Click the **Administration** tab.
3. Select **Security Administration** and click **OK**.
4. On the Workbook Template Rights worksheet, select each template for which a user's access rights require modification. Set to Denied, Read-only or Full Access.
5. Changes must be committed to the master database before they take effect. To commit the changes, select **Commit Now** from the File menu.
6. Save the workbook by selecting **Save** from the File menu, if desired.
7. To close the workbook, select **Close** from the File menu.

Set Measure Availability for Workbook Templates

1. From the File menu, select **New**.
2. Click the **Administration** tab.
3. Select **Security Administration** and click **OK**.

4. On the Workbook Template Measure Rights worksheet, select each registered measure that should be available for inclusion in the associated workbook template. For measures that should not be included in the associated template, make sure there is no check mark.
5. Changes must be committed to the master database before they take effect. To commit the changes, select **Commit Now** from the File menu.
6. Save the workbook by selecting **Save** from the File menu, if desired.
7. To close the workbook, select **Close** from the File menu.

Assign or Restrict User Access to Measures

1. From the File menu, select **New**.
2. Click the **Administration** tab.
3. Select **Security Administration** and click **OK**.
4. On the Measure Rights worksheet, for each measure that a user should have access to, select either **Read Only** or **Read/Write** from the drop-down list. For measures to which the user should not have access, make sure **Denied** is selected.
5. Any changes made must be committed to the master database before they take effect. To commit the changes, select **Commit Now** from the File menu.
6. Save the workbook by selecting **Save** from the File menu, if desired.
7. To close the workbook, select **Close** from the File menu.

Change User Ability to Modify Dimensions

1. From the File menu, select **New**.
2. Click the **Administration** tab.
3. Select **Security Administration** and click **OK**.
4. On the Dimension Modification Rights worksheet, select each dimension for which the user needs modification rights. For dimensions that the user should not be able to modify, make sure there is no check mark.
5. Any changes made must be committed to the master database before they take effect. To commit the changes, select **Commit Now** from the File menu.
6. Save the workbook by selecting **Save** from the File menu, if desired.
7. To close the workbook, select **Close** from the File menu.

Set or Modify Access to Positions

Use this procedure if position level security has been enabled.

1. From the File menu, select **New**.
2. Select the **Administration** tab.
3. Select **Security Administration** and click **OK**.
4. Select the worksheet for which security needs to be set or modified: **User**, **User Group**, or **World**.
5. By default, the dimension (level) at which position level security is enabled will be displayed. To manage security at a level above the designated level (only levels above are possible), right-click and **Select Rollup** to view the available dimensions.

6. To grant access to a position, click the check box of the cell.

Note: A user must have access at the User, User Group, and World levels to have access to a position.

7. Changes must be committed to the domain before exiting in order for them to take effect.

Limit the Number of Workbooks that a User Can Save

1. From the File menu, select **New**.
2. Select the **Administration** tab.
3. Select **Security Administration**.
4. Click **OK**.
5. Select the worksheet for which the limit will be set: User/Template, Group/Template, or Template.
6. Set the values as necessary.
7. Commit the data to the domain before exiting.

Set or Modify the Maximum Domain Session Limit

1. From the File menu, select **New**.
2. Select the **Administration** tab.
3. Select **Security Administration**.
4. Click **OK**.
5. On the Max Domain Session Limit worksheet, modify the scalar measure Maximum Domain Session Limit value with a valid integer value.
6. Changes must be committed to the master database before they take effect. To commit the changes, select **Commit Now** from the File menu.
7. To save the workbook, select **Save** from the File menu.
8. To close the workbook, select **Close** from the File menu.

Set or Modify the Maximum User Session Limit

1. From the File menu, select **New**.
2. Select the **Administration** tab.
3. Select **Security Administration**.
4. Click **OK**.
5. On the Max User Session Limit worksheet, modify the Maximum User Session Limit measure per user.
6. Changes must be committed to the master database before they take effect. To commit the changes, select **Commit Now** from the File menu.
7. To save the workbook, select **Save** from the File menu.
8. To close the workbook, select **Close** from the File menu.

Measure Analysis Workbook

The Measure Analysis workbook allows the user to view data associated with any registered measure in the RPAS applications, such as actual sales data for specified product/location/calendar combinations. The user may also use the Measure Analysis workbook to edit values for writable measures, however commit capability is only allowed to administrative users.

Although a common use of the Measure Analysis workbook is to view actual sales data, the workbook is not restricted to presenting sales data alone. The user can view any data loaded into the RPAS master database, such as selling prices, shipments, and orders. The Measure Analysis Wizard provides a list of all stored measures that have an insertable measure property set to true. (See the *RPAS Configuration Tools User Guide* for more information on measure properties.) The user simply chooses the measures to be displayed in the new workbook.


Note: Formatting settings cannot be saved in the Measure Analysis workbook because of its dynamic nature.

Building a Measure Analysis Workbook

The Measure Analysis Wizard guides the user through the process of creating a new Measure Analysis workbook in which the user can view selected measure data.

1. Access the New Workbook window by clicking **New** in the menu.
2. Select the **Analysis** tab and then select **Measure Analysis**. Click **OK**.

Figure 6–11 Measure Analysis Workbook



3. The Measure Analysis Wizard opens. Select the measures you want to include in the workbook. Click **Next**.
4. The Available Location Positions step appears. Select the locations to be included and click **Next**.
5. The Available Product Positions step appears. Select the products to be included and click **Finish**.

Measure Analysis Worksheet

The Measure Analysis worksheet allows the user to view the chosen measure data for the positions selected from the measure's associated hierarchies. Each Measure Analysis worksheet is displayed at a different dimensional intersection, depending on the measure selections made in the wizard. This dimensional intersection is shown in the worksheet title bar.

Figure 6–12 Example of Measure Analysis Worksheet

	1/4/2008	1/11/2008	1/18/2008	1/25/2008	2/1/2008
10000010Leather Loafer - Black 6 B	777.00	156.00	418.00	547.00	564.00
10000011Leather Loafer - Black 6.5 B	761.00	401.00	620.00	352.00	332.00
10000012Leather Loafer - Black 7 B	765.00	351.00	285.00	325.00	573.00
10000013Leather Loafer - Black 7.5 B	685.00	412.00	432.00	382.00	488.00
10000014Leather Loafer - Black 8 B	382.00	279.00	384.00	422.00	311.00
10000015Leather Loafer - Black 8.5 B	620.00	238.00	535.00	453.00	473.00
10000016Leather Loafer - Black 9 B	515.00	368.00	529.00	398.00	376.00

Figure 6–12 shows a Measure Analysis worksheet that displays Weekly Sales data for several items in a particular store. The location/product/calendar dimensional intersection of this worksheet, as shown in the title bar, is STR (Store), ITEM, WEEK. The Weekly Sales measure, because it is registered as a read/write measure, can be edited in this worksheet. However, only an administrative user can commit overwrites to writable measures in this workbook.

Review and Edit Sales or Other Registered Measure Data

To review and edit sales or other registered measure data, complete the following steps:

1. To open an existing Measure Analysis workbook: select **Open** from the File menu, double-click on the workbook to be opened, and go to step 9.
Or, to open a new workbook, select **New** from the File menu.
2. On the Analysis tab, select **Measure Analysis** and click **OK**.
3. The Measure Analysis Wizard opens and prompts the user to select the measures to be displayed in the new workbook. Use Ctrl-Click or Shift-Click to select multiple measures. Click **Next**.
4. For each hierarchy specified in the base intersection of the measures selected, use the hierarchy wizard to select positions to view. Repeat this step for each hierarchy wizard and click **Next**.
5. Click **Finish** to open the Measure Analysis workbook.
6. On the Measure Analysis Worksheets, view the stored data associated with the measures and hierarchy positions selected in the wizard. Make any changes as required. Administrative users may commit changes.

Hierarchy Management

There are a number of key concepts and processes that are critical to the hierarchy management process:

- Hierarchy structures are loaded into a domain using the loadHier utility.
- RPAS uses integer indexing for simplified hierarchy administration. A set number of hierarchy positions, based on bit size, is allocated to each dimension. The pre-allocation of positions reduces the need for updating the measure data structures.
- The length of position names is 24 characters or less by default. RPAS provides the ability to increase this length using the dimensionMgr utility.
- Position names must consist of *only* the following characters: a-z, A-Z, 0-9, _, &, \$, and %. Position names cannot start with _. Any uppercase letters are converted to lowercase letters by the application. Position names cannot be an empty string.
- RPAS provides the ability to have placeholder positions in the domain that can be used when loading new hierarchy positions.
- RPAS can automatically handle the movement of positions and their corresponding data between local domains when their parent-child relationships change and cause such a scenario. This is only applicable in a global domain environment.
- Positions at the partition level in a global domain environment can be moved between local domains using the reconfigGlobalDomainPartitions utility.
- New local domains can be added to an existing global domain environment using the reconfigGlobalDomainPartitions utility.

Loading Hierarchies Using loadHier

The loadHier utility is used to load and refresh a hierarchy. loadHier supports comma separated value (CSV) or fixed width flat files for loading. The load file should have a .dat file extension when a flat file is loaded and a .csv.dat file extension when a CSV file is loaded. When using a fixed width file, the width of fields (number of characters) is specified in a configuration file before a domain is built. The width of fields can be increased after a domain has been built using the dimensionMgr utility or by changing a property in the Configuration Tools and patching the domain. The utility also allows a simple compression method that can skip duplicated values line by line.

Note: The following information concerns hierarchy loading with intraday:

- Pre-13.2.2: Hierarchy loading is not supported while RPAS users are online unless all users are asked to limit the scope of their activities to workbook calculations only. Other operations such as workbook build, refresh, commit, custom menus, and dynamic position maintenance (DPM) can interfere with the hierarchy loading process. In such scenarios, users and administrators will experience concurrency issues and, in the worst case, hierarchy data corruption can occur.
 - 13.2.2 and later: The Ride framework allows the execution of the loadHier process. Users can stay logged in and continue their activities as long as they do not perform activities such as workbook build, refresh, DPM, and custom menu scripts that touch hierarchies. The system restricts users from performing these activities while the batch process is running. Users can submit commit-ASAP requests, which are queued and executed at the completion of the batch process.
-

To manage the addition, removal, and reclassification of positions in a hierarchy, RPAS uses a methodology called integer indexing. It is used to manage multidimensional data at the storage level. For more information, see "[Integer Indexing](#)".

The loadHier utility supports both the loading of hierarchy positions and purging data in parallel. When RPAS deletes a partition position through purging, RPAS adjusts the cache data in parallel to maintain the correct position or domain mapping.

RPAS allows for multiple input files to be loaded for the same hierarchy. The extra input files should be named with a secondary extension (for example, 'msgs.dat.1'). The extra input files can be loaded only with the main input file. For example, you cannot load 'msgs.dat.1' in a separate loadHier call. Multiple files are often used when the hierarchy load data comes from different sources.

RPAS automatically generates a backup copy of hierarchy files prior to performing a load for a hierarchy. If any type of error occurs during the load process, the hierarchy is restored from the backup copy. Note that in cases where the -loadAll argument is used to load multiple hierarchies at once, any hierarchies already loaded prior to the error will not be reverted; only the hierarchy in which the error was encountered will be restored.

The loadHier utility stops with an error if the loadHierBk directory exists in the data directory of the domain, which indicates that a non-recoverable error may have occurred in the previous run. If this occurs, contact My Oracle Support at <http://www.oracle.com/support/contact.html>. The My Oracle Support team can best determine whether to delete the loadHierBk directory or copy the loadHierBk content back to the domain.

To optimize performance while moving or cleaning data during the hierarchy purging or reclassification processes, the -excludeMeasList or -includeMeasList argument can be specified. Both arguments specify a full path to an XML file, in the following format, which contains a list of measures to either be excluded or included:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes" ?>
</rpas>
  <measures>meas1,meas2,meas3...,measN</measures>
</rpas>
```

Note: It is important to specify the measures with care, especially when using the `-includeMeasList` option. This is because no data moves for those measures not included in the list; such data is lost when loadHier completes.

When working with a domain sharing data through a RPAS Data Mart, loadhier also updates the ITT tables of the RDM to reflect the changes introduced in the hierarchy load. Additionally, loadhier is used on a domain integrated through a RPAS Data Mart, the RDM status tables are updated to reflect the fact that a hierarchy load is in progress. See the "Domain Integration and the RPAS Data Mart" section for more information on the hierarchy load process in an integrated environment.

loadHier Usage

```
loadHier -d domainPath -load hierName -loadAll {-purgeAge purgeage}{-purgeAll
hierarchy1} {-noClean}{-loglevel level} {-defaultDomain ldom#,
ldom#}{-excludeMeasList listName | -includeMeasList listName} {-includeUdd}
```

Table 7-1 provides descriptions of the arguments used by the loadHier utility.

Table 7-1 loadHier Utility Arguments

Argument	Description
-d domainPath	Indicates the domain in which to load the hierarchy data.
-load hierName	Indicates the name of the hierarchy to load and refresh.
-loadAll	Loads all hierarchy input files (with a .dat file extension) that are located in the input directory of the domain. Including this argument disables the reshaping process until all files have been loaded.
-purgeAge <i>purgeage</i>	Specifies the purgeage during loadHier. If not specified, loadHier gets purgeage from domain. In global domains, -purgeAge supports the purge of partition positions when the <i>purgeage</i> is reached.
-purgeAll <i>hierarchy1</i> , <i>hierarchy2</i>	Purges formal, informal, and user-defined positions in the listed hierarchies. It cannot be used on a partition hierarchy or any system hierarchy.
-noClean	Prevents the removal of input files and temporary data files that are generated during the hierarchy load process. Input files remain in the input directory of the domain after the process is completed. This option is often used for debugging or troubleshooting.
-logLoadedPositions	Enables the logging of successfully loaded input file lines into a loaded[HIERNAME].dat file under the processed directory.
-maxProcesses <i>count</i>	If specified, some parts of loadHier will run in parallel, meaning that it will use a maximum of the defined processes, which are specified by <i>count</i> .
-forceInputRollups	Enforces new hierarchy rollup changes. New rollup changes override or dominate existing hierarchy rollups if they conflict with the rollups specified in the input file. This allows you to load a hierarchy file that reclassifies one or more upper level positions while removing one or more discontinued base-level positions that roll up to the reclassified position.
-forceNAConsistency	Forces NA consistency when the current NA value is different from the originally defined NA value for the measure.

Table 7–1 (Cont.) loadHier Utility Arguments

Argument	Description
-includeUdd	Loads user-defined positions back to the domain. The data file must be in CSV format with a headerline. The name of the data file should follow the current standard: <hierarchy name>.csv.dat. All user-defined dimensions must be in the data file. Any missing user-defined dimensions cause an error. All loaded positions will have formal status after running -includeUdd.
-defaultDomain ldom#, ldom#,	Specifies comma-separated default domain paths that are used for accommodating new partitions. The domain paths can point to existing local domains or to new (non-existing) local domain. The local domain names are specified by a fully qualified path. To specify more than one local domain, separate local domain paths with a comma. Example: <pre>loadHier -defaultDomain ldom1, ldom2, ldom3</pre>
-excludeMeasList <i>listName</i> or -includeMeasList <i>listName</i>	Optimizes performance while moving or cleaning data during the hierarchy purging or reclassification processes: ■ Use -excludeMeasList to optimize performance by excluding the list of measures in <i>listName</i>. ■ Use -includeMeasList to optimize performance by including only the list of measures in <i>listName</i>.
-headerLine	Specifies the column order base on the header line. If not specified, the order is based on the start property of the dimension. This option is only applicable to the loading of CSV files; it has no effect on the loading of fixed-width files.

loadHier Notes

When using -defaultDomain, loadHier adds the new partition positions to the specified default domains one by one. The list of default domains is performed in the given order until each new partition position is added.

Example:

For a global domain that consists of two local domains, ldom0 and ldom1, using the following loadHier command:

```
loadHier ... -defaultDomain ldom1,ldom2,ldom3 ...
```

In this call, three new partition positions (part1, part2, and part3) are in the input file. When the loadHier finishes, there will be two new local domains, ldom2 and ldom3, with the following new partition positions included in them:

```
ldom1 --> part1
ldom2 --> part2
ldom3 --> part3
```

In the previous example, if only two new partition positions (part1 and part2) are added, when the loadHier finishes there will only be one new local domain, ldom2. New partition positions will be located as follows:

```
ldom1 --> part1
ldom2 --> part2
```


Using the same example, if five partition positions (part1, part2, part3, part4 and part5) are added, when the loadHier finishes there will be two new local domains, ldom2 and ldom3. New partition positions will be located as follows:

```
ldom1 --> part1,part4
ldom2 --> part2,part5
ldom3 --> part3
```

Position Label Translation

To enable translation of position labels for the desired dimensions using the RPAS Configuration Tools, check the box in the Translate column of the Dimension definition tools. Building or patching the domain with this configuration builds the necessary infrastructure in the domain to manage translations for those dimensions. However, label translations must be separately loaded.

Position label translations are loaded in dimension specific translation measures for every language that is used by the users. If translated labels are not loaded using these measures, workbooks show position names wherever a label has to be shown. Note that for a translatable dimension, RPAS never uses or shows the position labels from the hierarchy load file but always refers to the labels in dimension specific translation measures. This implies that if a domain is patched to make a dimension translatable but the translation measures were not loaded, RPAS users will see position names instead of position labels from the load file.

Dimension specific position translation measures are named as r_<dim name>label, where <dim name> must be replaced with the name of the translated dimension. For example, if the sku dimension is translated, load the r_skulabel measure with translations. These measures must be loaded after loading the hierarchy because RPAS can only load translations for already loaded positions.

The position label translation measure load files have three columns. The first column has the position names, the second column has the language identifier, and the third column has the translation for the language specified in that row.

For example, a translation measure file for the sku dimension is named r_skulabel.csv.ovr and has the content formatted, as shown below. Note that in the following example, the same file contains labels in four languages.

```
10006782,ENGLISH,White Nike Running Shoe size 11
10006782,CHINESE_SIMPLIF,白色耐克?鞋大小 11
10006782,FRENCH,Taille blanche 11 de chaussure de course de Nike
10006782,ITALIAN,Formato bianco 11 del pattino corrente di Nike
10004523,ENGLISH,Black leather shoe size 8
10004523,CHINESE_SIMPLIF,黑皮鞋大小 8
10004523,FRENCH,Taille noire 8 de chaussure en cuir
10004523,ITALIAN,Formato nero 8 del pattino di cuoio
```

Note: For a list of language identifiers, see [Table 11-1](#).

Alternatively, you can manually enter or alter translated labels using the Translations workbook in the Administration tab. In this workbook, a worksheet is available for each dimension that has translations enabled. You can manually enter translated strings for the language of interest. After they are committed, these translations are available for every new workbook.

It is possible that, because of errors that occurred when translation files were prepared, translated labels for some positions may not be loaded. In a situation where RPAS is

unable to look up the label for the locale of the machine on which the RPAS client is being run, RPAS looks for a non-empty label string for the English language. If it fails to find a non-empty label string for the English language, it uses or shows the loaded position name of the position.

Note: For the fixed-width format of translation measure load files, RPAS limits the labels to 80 bytes (RPAS uses UTF-8 encoding). For CSV format files, there is no limit. To avoid the complexity of calculating starting positions for fixed-width format files and the limitation of translation string length, it is recommended that CSV files be used.

Integer Indexing

Integer indexing refers to the RPAS methodology for managing multidimensional data at the storage level. It aims at eliminating some of the inefficiencies around the addition, removal, and reclassification of positions in a hierarchy.

In the earlier versions of RPAS, string identifiers were constructed and used for internal addressing of the multidimensional data. Integer Indexing changes this mechanism in the sense that the data is internally referenced and stored by integer identifiers that are calculated from the dimension size and measure intersection.

Integer index calculation and management is entirely internal to RPAS. From a domain administrator's point of view, loading and exporting hierarchy and measure data is achieved through external names consistent with earlier versions of RPAS.

One major advantage of integer indexing is its impact on reducing the complexity of operations around managing position placeholders (buffers). The number of positions that a dimension can hold is defined by a bit size, which typically provides significant space for including new positions. For instance, a SKU dimension at 24 bits can include in excess of 16 million products. This provides significant growth space for adding new products.

Compared to the earlier versions of RPAS, integer indexing can improve the hierarchy load performance as it eliminates the need for running the frequent rebuffering processes and completely eliminates the reshaping process.

Note the following definitions:

- **Index:** An identifier that allows for random access to the multidimensional data. In RPAS, it is a numerical value that identifies a position in a given dimension. Prior to integer indexing, RPAS used string-type position identifiers and referred to them as internal names. Internal names were internally generated by RPAS. They were often different from the Position IDs provided in the hierarchy input files.
- **Position ID:** The name provided in the input hierarchy files. RPAS maps the position ID to an internally generated integer index.

Consider the following example in which some items are loaded into the Product hierarchy. These items are externally identified by the position names and labels (as shown in the second and third columns in [Table 7-2](#)). The table represents the SKU dimension in the product hierarchy. Integer indexes, names, and position labels are the three major pieces of information associated with a dimension.

Table 7–2 Example 1

Integer Index	Position Name	Position Label
0	SKU1234	blue sweatpants
1	SKU2345	white sweatpants
2	SKU3456	green sweatpants
3	SKU4567	purple sweatpants
4	SKU5678	yellow sweatpants
5	SKU6789	pink sweatpants
6	SKU7890	orange sweatpants
7		
8		
9		

When the green sweatpants item is loaded into the domain, its position ID, SKU3456, is mapped to an available integer index, which is 2. After these items are loaded, RPAS uses the integer indexes to look up those items instead of their item IDs. Therefore, if you performed a query for green sweatpants, RPAS first identifies that the integer index of the green sweatpants is 2. Then, it searches for instances of 2 in the array. The reason that RPAS searches for 2 instead of SKU3456 is because searching for an integer is faster than searching for a string.

Here is what happens if you delete the blue and white sweatpants from your domain because they are discontinued. Rather than delete these position IDs from the dimension, RPAS marks them as inactive. By keeping these position IDs in the dimension, RPAS maintains the position ID to integer index mapping. This way, the green sweatpants are still mapped to the integer index 2.

Table 7–3 Example 2

Integer Index	Position Name	Position Label	Position Status
0	SKU1234	blue sweatpants	Inactive
1	SKU2345	white sweatpants	Inactive
2	SKU3456	green sweatpants	
3	SKU4567	purple sweatpants	
4	SKU5678	yellow sweatpants	
5	SKU6789	pink sweatpants	
6	SKU7890	orange sweatpants	
7			
8			
9			

As you delete items from the domain, the dimension begins to have gaps in it where inactive items are using integer indexes. This is not an issue unless you want to add more position IDs than the dimension has room for. When this happens, you need to reindex the dimension first.

For example, add five more colors of sweatpants to the dimension. This dimension has ten integer indexes, so it can hold ten total position IDs. As the dimension is now, it has only three available integer indexes (7, 8, and 9) even though there are five indexes that are not being used (0, 1, 7, 8 and 9). To use the 0 and 1 integer indexes that the inactive sweatpants are mapped to, use the `reindexDomain` utility to defragment the dimension. This removes the blue and white sweatpants and allows you to use the 0 and 1 integer indexes for the new sweatpants. For more details about `reindexDomain` and defragging, see ["Reindexing Domains Using reindexDomain"](#).

Table 7–4 Example 3

Integer Index	Position Name	Position Label	Position Status
0	SKU3456	green sweatpants	
1	SKU4567	purple sweatpants	
2	SKU5678	yellow sweatpants	
3	SKU6789	pink sweatpants	
4	SKU7890	orange sweatpants	
5			
6			
7			
8			
9			

However, a different result occurs if you add 15 new SKUs instead of 5. Even if you defragment the dimension, there is room for only 5 new SKUs. In this instance, you need to add more room to the dimension by increasing the dimension's bit size. There are two ways to increase the bit size: through the configuration and patching process or through the `dimensionMgr` utility. (For more information about `dimensionMgr`, see ["Setting Properties for Dimensions Using dimensionMgr"](#).)

Changing the Bit Size

To change the bit size after a domain has been built or upgraded to 13.3, complete one of the following sets of steps:

1. Update the configuration and patch the domain. See ["Upgrading and Patching Domains"](#) for instructions. To change the bit size, see the "Defining Dimension Properties" in the *Oracle Retail Predictive Application Server Configuration Tools User Guide*.
2. Run the `reindexDomain` utility. See ["Reindexing Domains Using reindexDomain"](#) for instructions.

Or,

1. Run the dimensionMgr utility to update the bit size. See ["Setting Properties for Dimensions Using dimensionMgr"](#) for instructions.
2. Run the reindexDomain utility. See ["Reindexing Domains Using reindexDomain"](#) for instructions.

Reindexing Domains Using reindexDomain

Use the reindexDomain utility to compress, increase, or decrease the set of physical address space (or indexes) of the multidimensional arrays. The process of compressing and defragmenting the physical IDs makes the domain load and run faster.

Compression is achieved by removing the gaps that develop due to hierarchy operations like purge and delete. Gaps essentially block certain physical IDs within an address space and prevent them from being used to store data in the individual dimension arrays. After the gaps are removed, the dimension space of the multidimensional array is recreated.

Increasing or decreasing the multidimensional array's address space is also achieved by increasing or decreasing the bit size of the dimension arrays and then recreating the dimension space of the array with the reindexDomain utility. In addition to defragmenting, increasing, and decreasing the number of indexes, the reindexDomain utility also updates any measure array affected. The update is similar to the dimension data because they are reindexed and defragmented. An affected measure is one that has a dimension that was reindexed in its base intersection.


Note: While working with a global domain, the reindexDomain utility runs from the master domain and spawns subprocesses across subdomains if more than one process is used.

Domain reindexing should not invalidate existing workbooks.

After the reindexing process, the workbook refresh and build operations may run more slowly. Users who try to open a workbook that was built before the reindexing process should be prompted and informed about the potential slow performance.

A workbook is old if its dimregistry version is older than the domain dimregistry version.

Workbook commit and refresh processes will use external name dimension maps if the workbook is old.

Figure 7–1 Reindexing Domain Process Flow

When operating on a domain that shares data through a RPAS Data Mart, reindexDomain also updates the ITT tables of the RDM to reflect changes made as a result of the reindex process. See the "Domain Integration and the RPAS Data Mart" section for more information on domain operations in an integrated environment

reindexDomain Usage

```
reindexDomain -d <domainPath> {Commands} {Options}
```

There are three unique ways to use the reindexDomain utility: reindex, analyze, and domain properties. These are described in the following sections:

ReindexDomain Option: Reindex

The reindexing options allow you to reindex the entire domain, dimensions within a hierarchy, a specified list of dimensions, or prepend calendar dimension positions.

Table 7–5 reindexDomain: Reindexing

Example	Description
<code>reindexDomain -d <domainPath></code>	Reindexes the entire domain. Checks the reindex condition.
<code>reindexDomain -d <domainPath> -hier <hierName></code>	Reindexes all dimensions of the specified hierarchy.
<code>reindexDomain -d <domainPath> -dimSpec <Dim1,Dim2,...></code>	Reindexes the comma separated list of dimensions.

Table 7–5 (Cont.) reindexDomain: Reindexing

Example	Description
<code>reindexDomain -d <domainPath> -hier CLND -prepend</code>	Creates a new DimRegistry by allocating space to prepend calendar dimensions' positions. The input to this process comes from <code>clndprepend.xml</code> file which is generated by <code>loadHier</code> before it throws an exception trying to prepend. After this process, <code>loadHier</code> must be run again to complete prepending the calendar hierarchy.

Two options can be used with these usages:

- `-force`: Forces reindexing at domain/hierarchy/dimension level. It does not check the reindexing condition.
- `-processes max`: Specifies the number of processes to run in parallel.

ReindexDomain Option: Analysis

The analysis option of the `reindexDomain` utility allow you to analyze the dimensions of a hierarchy or a CSV list of dimensions.

Table 7–6 reindexDomain: Analysis

Example	Description
<code>reindexDomain -d <domainPath> -analyze -hier <hierName></code>	Analyses the dimensions of the hierarchy and prints a report.
<code>reindexDomain -d <domainPath> -analyze -dimSpec <Dim1, ...></code>	Analyses the CSV list of dimensions and prints a report.

The report consists of the following information:

- DimName
- DimRegistry BitSize
- DimInfo BitSize
- Capacity
- Maximum Physical Id
- Maximum Logical Id
- Threshold
- Action Required

ReindexDomain Option: Domain Properties

The `reindexDomain` utility can be used to review the domain properties by appending the `-prop` argument.

Table 7–7 reindexDomain: Domain Properties

Example	Description
<code>reindexDomain -d <domainPath> -prop</code>	Prints the following domain properties: ■ <code>dimregistry_version</code> ■ <code>reindexing_in_progress</code>

[Table 7–8](#) provides descriptions of the arguments used by the reindexDomain utility.

Table 7–8 reindexDomain Arguments

Argument	Description
-d	Specifies the path to the global or non-partitioned domain. This argument is required for all reindex usages.
-hier	Reindexes all the dimensions of the specified hierarchy. When used with the -analyze argument, the -hier argument specifies which hierarchies should be analyzed.
-dimSpec	Reindexes the comma separated list of dimensions. When used with the -analyze argument, the -dimSpec argument specifies which dimensions should be analyzed.
-force	Forces the reindexDomain utility to reindex the entire domain whether it needs it or not. This argument overrides the logic that checks whether reindexing is necessary.
-prepend	Creates a new DimRegistry by allocating space to prepend calendar dimensions' positions. The input to this process comes from clndprepend.xml file. The loadHier utility generates the clndprepend.xml file when you load the calendar hierarchy that has positions been prepended. After loadHier generates this file, it displays a message, stating that reindexDomain needs to be run first. The clndprepend.xml must exist in the master domain root. After you run the reindexDomain utility, you must run the loadHier utility again to complete the prepending the calendar hierarchy process. You can also generate clndprepend.xml manually without using loadHier. Use this argument only when you are prepending the calendar hierarchy.
-processes max	Specifies the number of processes to run in parallel.
-analyze	Checks whether reindexing is necessary for hierarchies or dimensions. This argument returns details about the dimension name, bit size, number of available and used position IDs, the threshold ratio, and whether reindexing is necessary.
-prop	Shows the domain properties. The properties are dimregistry_version and reindexing_in_progress, where the dimension registry number and the progress indicates whether a reindex is currently running. For more information about these domain properties, see Version and ReindexStatus .

When to Reindex

To know whether you should reindex a domain, run the reindexDomain utility with the -analyze option to generate a status report. (Detailed instructions for using the analyze option are described in ["Run Reports"](#).) RPAS calculates how much space is left in the array with the following equation:

$$(C-M) < \text{Thresh}$$

Where:

- C equals the maximum number of position IDs that the array has been allocated to hold. C is the dimension or BitSize cardinality, which is computed with the following equation:

$$C = 2^{\text{BitSize}}$$

BitSize is used to compute the cardinality of a dimension. The BitSize can be represented by up to a 4-byte unsigned long integer, which gives a maximum of

four gigs of address space. Since four gigs is very large, the BitSize is specified at configuration time. The BitSize is determined based on the current size of the dimension and expected increase in future. If the user does not specify BitSize for a dimension at configuration time, then the RPAS Configuration Tools provides a default BitSize of 8 for that dimension.

For domains that were converted to 13.3.0, you must ensure that the hierarchy.xml file contains an entry for all dimensions along with the bit sizes.

- M equals the number of used position IDs in the dimension. This is across all position IDs currently assigned to the dimension positions. The maximum number of utilized physical IDs can differ from the maximum number of utilized logical IDs if the hierarchy is subjected to deletions and additions.
- Thresh is the threshold value you set for a dimension. The threshold is an integer that represents the number of unused position IDs that you want to be available in the dimension array. When the dimension array no longer has this specified amount of free space, RPAS reindexes the dimension array. The threshold is part of the configuration and is configured as a dimension attribute.

The threshold can be specified per dimension and can be configured to any initial value. The default value is 10 percent of the BitSize cardinality for a particular dimension. This value helps determine what approach you need to take:

- Increase the address space of a dimension after reindexing the dimension arrays
- Compress the domain arrays without increasing the BitSize
- Both
- L is the maximum utilized logical ID, which is the logical ID value of the last position in a dimension. The domain sequence generator displays this number, which is incremented as new positions are added to the dimension.

For example, if you have a dimension with the following characteristics:

- C: Dimension or BitSize Cardinality for the dimension: **16** (where Bitsize is 4)
- M: Maximum Utilized Position IDs: **9**
- Thresh: Threshold: **10**
- L: Maximum Utilized Logical ID: **4**

The calculation produces the following:

$$\begin{aligned} (C-M) &< \text{Thresh} \\ (16 - 9) &< 10 \\ 7 &< 10 \end{aligned}$$

Since 7 (the number of available position IDs) is less than 10 (the threshold of required available position IDs), the report returns a message that states the dimension needs to be defragmented.

RPAS also checks the $C - L > \text{or} < \text{Thresh}$ condition. If $C - L > \text{Thresh}$, then the report returns a message that states not to increase the BitSize, but to defragment the dimension to recover unused address space. In the example above, $C - L$ ($16 - 4 = 12$) is less than the threshold (10). Therefore, the available address space is sufficient after the dimension is defragmented.

If you have a dimension with the following characteristics:

- C: Dimension or BitSize Cardinality for the dimension: **16 (where Bitsize is 4)**

- M: Maximum Utilized Position IDs: 9
- Thresh: Threshold: 10
- L: Maximum Utilized Logical ID: 7

The calculation produces the following:

```
(C-M) < Thresh
(16-9) < 10
7 < 10

(C-L) < Thresh
(16-7) < 10
9 < 10
```

Since 9 is less than 10, the report returns a message that states that defragmenting alone will not help, but that a BitSize increase is required to create additional address space.

Before calculating this equation, RPAS validates the inputs. The main inputs to this utility are the domain path, hierarchy names, and dimension names. Validation logic checks the syntax of the command and verifies all inputs

Run Reports

Three reports that you can run with the reindexDomain can help you determine if you should reindex your domain. Depending on the results of these reports, RPAS may recommend that you need to reindex the dimension, that you need to increase its bit size and then reindex, or that you do not need to do either of these because the cost of reindexing outweighs the benefit. To find out which one of these is the case, run the following commands:

Hierarchies and Dimensions This command analyzes one or more dimensions to check whether reindexing is necessary. You can use either the -hier argument or the -dimspec argument.

If the -hier argument is used, the command analyzes all dimensions in the specified hierarchy to check whether reindexing is necessary:

```
reindexDomain -d [domainPath] -analyze -hier [hierName]
```

If the -dimspec argument is used, the command analyzes all the dimensions in the CSV list to check whether reindexing is necessary:

```
reindexDomain -d [domainPath] -analyze -dimSpec [comma seperated dim names]
```

The command returns the reindexing details. The following is an example of the returned details:

```
Dim Name           : splr
DimRegistry BitSize: 8
DimInfo           BitSize: 8
Capacity          (C): 256
MaxPhyId          (M): 3
MaxLogId          (L): 3
Threshold:        (T): 25
```

```
Action Required    : No action required for this dimension at this stage
```

Version and ReindexStatus To see the DimRegistry version and reindex status, run the following argument:

```
reindexDomain -d [domainPath] -prop
```

It returns details about the following two domain properties:

```
dimregistry_version      reindex_in_progress
```

These two domain properties can be set using the domainprop utility. But when reindexDomain is run, these two properties are set automatically. Therefore, there is no need to set them separately unless you want to see their values before using the reindexDomain utility.

- **dimregistry_version:** This domain property stores the current version number of the dim registry. A dim registry defines the version of the indexes. The version number is reset with a new value when reindexing is performed. For instance, if SKU123's index is 15, but then reindexing causes that same SKU to be at index 2, then the dimregistry_version is incremented. The new version is created and stored in the domain at the same level as the previous version. Each version can be identified physically as it is named after its version number. The path up to the version is the same across all versions so that data can be processed from one version to another.
- **reindex_in_progress:** This domain property stores the status of the last run reindex process. It is set to TRUE when the reindexing starts and to FALSE when it successfully completes. When the process aborts in the middle, the status remains TRUE to indicate that reindexing needs to be restarted. As long as the status is TRUE, the domain cannot be used for any purpose other than reindexing.

Note: After the reindexing process starts, the domain must not be used for any other purpose until the process finishes successfully. If the process aborts in the middle, then the reindexing process must be restarted and completed successfully.

Only an administrator should run this utility.

How to Reindex a Dimension

The sections below describe the various scenarios and steps needed for reindexing.

Reindex Entire Domain as Needed

This checks whether the reindex condition is met by at least one dimension in the domain. If so, the reindexDomain utility stops further checks and reindexes the entire domain. If reindex condition is not met by any of the dimensions, the utility exits.

It is recommended that you do not reindex the entire domain. Instead, you should reindex on a hierarchy-by-hierarchy or dimension-by-dimension basis.

```
reindexDomain -d <domainPath>
```

Reindex Entire Domain with Force Option

By using the reindexDomain with the -force option, you can reindex the entire domain. The -force option overrides the logic that checks whether reindexing is necessary and reindexes the entire domain.

```
reindexDomain -d <domainPath> -force
```

Reindex Dimensions

Reindexes one or more dimensions. Use either the `-hier` argument or the `-dimSpec` argument

When used with the `-dimSpec` argument, the utility reindexes only those dimensions in the CSV list that meet the reindex condition. If listing more than one dimension, use commas to separate them.

```
reindexDomain -d <domainPath> -dimSpec <comma separated dim names>
```

When used with `-hier` argument, if at least one dimension in the hierarchy meets the reindex condition, the utility reindexes all the dimensions in the specified hierarchy.

```
reindexDomain -d <domainPath> -hier <hierName>
```

Note: You can use the `-force` option with both the `-hier` and the `-dimSpec` arguments.

Prepending Calendar Dimensions

This is a unique case where you need to add calendar positions before the current start of the calendar. This scenario must be handled differently from scenarios where you add calendar positions at the end of the calendar. This is because the DimRegistry for calendar dimensions must make space at the beginning for the new positions to be added. Here is what happens during the process of prepending positions to calendar dimensions:

1. The loadHier utility detects during its first run that clnd.dat has positions to be prepended. It generates an XML file (clndPrepend.xml) that contains the dimensions to be prepended and the number of new positions to be prepended for each dimension. The loadHier utility then displays an error message, stating that the reindexDomain utility must be run first before prepending the calendar dimensions.

Here is an example of the clndPrepend.xml file:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes" ?>
- <clnd>
- <prepend>
  <dimension>day</dimension>
  <prependsize>1457</prependsize>
</prepend>
- <prepend>
  <dimension>mnth</dimension>
  <prependsize>48</prependsize>
</prepend>
- <prepend>
  <dimension>qtr</dimension>
  <prependsize>16</prependsize>
</prepend>
- <prepend>
  <dimension>ssn</dimension>
  <prependsize>8</prependsize>
</prepend>
- <prepend>
  <dimension>week</dimension>
  <prependsize>208</prependsize>
</prepend>
- <prepend>
  <dimension>year</dimension>
```

```
<prependsize>4</prependsize>
</prepend>
</clnd>
```

2. The reindexDomain utility is run with syntax added specifically for prepending calendar dimensions:

```
reindexDomain -d <domainPath> -hier CLND -prepend
```

3. The reindexDomain domain reads the clndPrepend.xml file and creates a new version of the DimRegistry with space allocated to prepend the calendar dimensions by shifting the physical IDs by an offset equal to the number of new positions. The reindexDomain utility then reindexes all the arrays in the domain that have these calendar dimensions.
4. After running reindexDomain, run the loadHier utility again to complete the process of prepending the calendar dimensions. During both the first and second runs of loadHier, the loadHier utility checks whether calendar dimensions are ready to be prepended or not. The DimRegistry interface is modified to verify this condition.

If you want to avoid using loadHier for generating the clndPrepend.xml file, then you can manually generate the file and place it under the domain root. Ensure that you use the same syntax as the one loadHier generates, as shown in the previous example.

Reindex Domain Examples

This table does not contain a comprehensive list of scenarios that can occur. It only lists some examples of the reindexDomain functionality.

Condition 1: Change of Bit Size

Table 7–9 *Change of Bit Size Scenarios*

Scenario	Utility Option	Result
Increase/Decrease Bit Size for SKU dimension	-dimSpec SKU	Reindexes SKU dimension and create a new dimRegistry version.
Increase/Decrease Bit Size for SKU and Store dimensions	-dimSpec SKU,STR,CLSS	Reindexes SKU and Store dimensions only. The Class dimension does not meet any reindex condition. A new dimRegistry version is created.
Increase/Decrease Bit Size for SKU dimension	-hier PROD	Reindexes all dimensions in the product hierarchy and creates a new dimRegistry version. For SKU dimension, the bit size change is applied to the dimension and measures.
Increase/Decrease Bit Size for SKU	-d	Reindexes all the dimensions in the entire domain and creates a new dimRegistry version. For the SKU dimension, the bit size change is applied to the dimension and measures.

Condition 2: Using the -force Option

Table 7-10 -force Option Scenarios

Scenario	Utility Option	Result
No bit size change or any fragmentation	-dimSpec SKU	Reindexes the SKU dimension and creates a new dimRegistry version.
Bit size change for SKU, fragmentation in CLSS, no fragmentation in STR	-dimSpec SKU,CLSS, STR	Reindexes by changing the bit size for the SKU dimension and defragmenting CLSS and STR. Creates a new version of dimRegistry.
Bit size change for SKU and fragmentation in CLSS and STR	-hier PROD	Reindexes all dimensions in the product hierarchy and creates a new dimRegistry version. SKU dimension has bit size change applied. STR dimension and location hierarchy are not affected.
Bit size change for SKU, fragmentation in CLSS, no fragmentation in STR	-d	Reindexes all the dimensions in the entire domain and creates a new dimRegistry version. For the SKU dimension, the bit size change is applied to the dimension and measures.

Condition 3: Fragmentation Checks

Table 7-11 Fragmentation Check Scenarios

Scenario	Utility Option	Result
Fragmentation in SKU dimension but does not meet reindex condition. In this case, $C-M > TR$	-dimSpec SKU	Nothing happens.
Fragmentation in SKU dim - does not meet the reindex condition (here $C-M > TR$) and fragmentation in STR dimension which meets the reindex condition ($C-M < TR$)	-dimSpec SKU,STR	Only the STR dimension will get reindexed. A new dimRegistry version is created.
Fragmentation in SKU dimension but does not meet reindex condition. In this case, $C-M > TR$	-hier PROD	Nothing happens.
Fragmentation in SKU, CLSS and STR but none meet reindex condition. In this case, all three have $C-M > TR$.	-d	Nothing happens.
Fragmentation in SKU dimension where $C-M < TR$ and $C-L > TR$.	-dimSpec SKU	Reindexes the SKU dimension and creates a new dimRegistry version. This condition means that there is fragmentation in the SKU dimension over the threshold and defragmenting can generate enough address space to resolve the initial reindex condition of $C-M < TR$.

Table 7-11 (Cont.) Fragmentation Check Scenarios

Scenario	Utility Option	Result
Fragmentation in SKU dimension with C-M<TR and C-L>TR and fragmentation in STR with C-M>TR.	-dimSpec SKU,STR	Only the SKU dimension is reindexed. A new dimRegistry version is created.
Fragmentation in CLSS dimension with C-M<TR.	-hier PROD	All dimensions in the product hierarchy are reindexed. A new dimRegistry version is created.
Fragmentation in both the SKU and STR dimensions. Both dimensions have C-M<TR and C-L>TR.	-hier PROD	All dimensions in the product hierarchy are reindexed. Nothing happens in the location hierarchy or to the STR dimension. A new dimRegistry version is created.
Fragmentation in the STR dimension with C-M<TR and C-L>TR.	-d	All dimensions in the entire domain are reindexed. A new dimRegistry version is created.
Fragmentation in the SKU dimension with C-M<TR and C-L<TR.	-dimSpec SKU	No reindexing of SKU takes place as it does not resolve the issue since a bit size increase is needed. A message is provided through the utility indicating a bit size increase is required.
Fragmentation in the SKU dimension with C-M<TR and C-L<TR and STR dimension with no fragmentation	-dimSpec SKU,STR	No reindexing takes place on either SKU or STR. SKU requires a bit size increase. A message is provided through the utility indicating a bit size increase for SKU is required.
Fragmentation of all dimensions in the product hierarchy were C-M<TR and C-L<TR.	-hier PROD	No reindexing, but a message is provided through the utility, stating that all dimensions in the product hierarchy need an increase in bit size.
Fragmentation for the SKU and STR dimension where both have C-M<TR and C-L<TR.	-d	No reindexing, but a message is provided through the utility, stating that the SKU and STR dimensions need an increase in bit size.

Optimizing Domains Using optimizeDomain

Because the RPAS Btree dimension arrays undergo continuous updates and changes as a result of adding and deleting existing positions, over time, measure arrays become full of stale data. This stale data is created when positions are deleted, but the associated index remains. This results in wasted space and inefficient operations. The measure arrays containing any data for these positions should be updated to reflect these deletions. Deleting positions marks the hierarchy data for the corresponding dimensions as changed, but it does not clean up the associated data from the measure arrays. The optimizeDomain utility cleans this stale data from the measure arrays.

When this utility is run in a global domain environment, it is centrally administrated. This means it runs over the master domain and spawns parallel processes over the local domains.

For example, consider a measure called meas1 that exists at intersection sku_str_week. When a user adds a DPM position skudpm1 through an RPAS client, it is assigned to an integer index in the meas1 data array in addition to being assigned within the dimension's array. When this DPM position is deleted, the position is flagged as inactive, but the integer index cannot be reused until the dimension is reindexed. The

measure array also still contains the data associated with the inactive position. This measure array needs to be updated too in order to remove the stale data.

Note: The reindex process cleans data, but there are times when a reindex is not needed but data needs to be cleaned. That is when optimizeDomain should be run.

The same also holds true for formal positions that are purged during a normal hierarchy patch. The measures arrays need to be updated for this set of positions because they may be pointing to stale data too.

Another issue is that, over time, existing arrays in RPAS databases may become fragmented, resulting in wasted space and a possible degradation in efficiency on array operations. This occurs because the RPAS Btree array stores data in chunks of disk memory called pages. Deleting data from the array causes empty spaces or holes to develop in these pages. These holes contain no data, yet they still increase the overall size of the array. Over time, these holes become larger and more frequent, and array operations suffer degraded performance as a result. Also, the pages begin to acquire large amounts of unused space, causing an inefficient disk usage and a larger than optimal domain size.

Notes: While reclassifying a position from one subdomain to another, data associated with the reclassified positions is moved to the new subdomain. But it is also cleaned from the source subdomain. Therefore, the reclass operation in loadHier does not cause stale data.

No matter what option is specified, when an array is picked to be processed by optimizeDomain, that array is both cleaned and defragmented after a successful run.

The optimizeDomain utility is used to correct these two issues. It is recommended that after a full hierarchy purge you run the optimizeDomain utility with the -cleanOnly argument to remove all stale data from the domain.

optimizeDomain Usage

```
optimizeDomain -d [domainpath] [options]
```

Table 7–12 provides descriptions of the arguments used by the optimizeDomain utility.

Table 7–12 *optimizeDomain Usage*

Argument	Description
-d domainpath	Defines the path to the domain that you want to optimize. If no options are specified, by default, optimizeDomain defragmentss and cleans the entire measure set that have arrays that have a BTree density ratio lower than the threshold (80%). A BTree density ratio is the ratio of the populated positions in the array to the total available BTree pages in the array. The density threshold value (80%) is not configured.
-defragOnly	Selectively defragments the domain data based on database fragmentation. Note: -cleanOnly and -defragOnly cannot be specified together. To achieve that behavior, run the utility without any options.

Table 7–12 (Cont.) optimizeDomain Usage

Argument	Description
-cleanOnly	Cleans the stale data in the domain. Note: -cleanOnly and -defragOnly cannot be specified together. To achieve that behavior, run the utility without any options.
-force	Forces the utility to defragment the entire measure arrays set. This must be used in combination with the -defragOnly option. This option makes the defragment process run longer. This option overrides the Btree density threshold (80%) and defragments all measure arrays regardless of their BTree density.
-excludeMeasList	Defines the path to exclude list XML file. Optimize performance by excluding this list of measures while defragmenting or cleaning. This option cannot be used with -includeMeasList.
-includeMeasList:	Defines the path to include list XML file. Optimize performance by including only this list of measures while defragmenting or cleaning. This option cannot be used with -excludeMeasList.
-returnFailure	Returns a non-zero value if any of the databases or arrays fail to defragment or clean. The optimizeDomain process still runs to completion. The default behavior is to log a warning and continue.
-processes max	Defines the maximum number of processes to run in parallel.

Usage Examples

In this example, the measures list and corresponding BTree density ratios. SKU dimension have some positions purged from the domain and marked for cleaning.

Table 7–13 SKU Positions Purged

	Meas1	Meas2	Meas3	Meas4
Base Intersection	SKU	SKU	STR	STR
Density Ratio	60%	90%	60%	90%

Here is the list of measures that are processed based on the specified command.

Table 7–14 Processed Measures

	Meas1	Meas2	Meas3	Meas4
No specified options	X	X	X	
-includeMeas List Meas1, Meas3, Meas4	X		X	
-excludeMeasList Meas1,Meas3		X		
-defragOnly	X		X	
-cleanOnly	X	X		
-defragOnly -force	X	X	X	X
or				
-force				

Table 7–14 (Cont.) Processed Measures

	Meas1	Meas2	Meas3	Meas4
-defragOnly -includeMeasList Meas1,Meas4	X			
-defragOnly -excludeMeasList Meas1			X	
-cleanOnly -includeMeasList Meas1,Meas4	X			
-cleanOnly - excludeMeasList Meas1		X		

Adding New Dimensions to Hierarchies

Using the RPAS Configuration Tools and RPAS utilities, you can add new dimensions to hierarchies in existing domains. This process is described at a high level in this section. Each step is described in greater detail in the following sections.

Caution: Due to its invasive impact, adding new dimensions is different from the normal domain patch process. It is a multi-step process and requires careful user intervention.

1. To add new dimensions to an existing hierarchy, use the Configuration Tools to open and modify the existing configuration so that the new configuration includes the additional dimensions.

Note: This process can be used only for non-partitioned hierarchies. It cannot be used to delete dimensions.

After the configuration has been modified, use the Reports Generator in the Configuration Tools to generate a hierarchy.xml report. For more information about creating this report, see the "Report Generator" section in the *Oracle Retail Predictive Application Server Configuration Tools User Guide*.

2. Export measure data for all measures that have one of the dimensions of the modifying hierarchies. See ["Exporting Measure Data"](#).
3. Export all hierarchy data from the hierarchy, including all DPM positions and user-defined dimension rollup information. Repeat this step for each hierarchy being modified. See ["Exporting Hierarchy Data"](#).
4. Purge all positions from the hierarchies. See ["Purging Hierarchy Data"](#).
5. Use the hierarchyMgr utility with the new hierarchy.xml to add the new dimensions. Repeat this step for each hierarchy being modified. See ["Adding New Dimensions"](#).
6. Create a new data file, or update the old one, to contain positions for the newly added dimensions. Load it back to the domain. If the original hierarchy contains user-defined dimensions, the format of the data file must be in CSV format with a header line indicating all user-defined dimension position and label columns and using the -includeUdd switch. Perform one execution for each modified hierarchy. See ["Reloading Formal Hierarchy Data"](#).

7. Optional: Informalize previous DPM positions. See ["Informalizing DPM Positions"](#).
8. Reload measure data to the domain. See ["Reloading Measure Data"](#).

Exporting Measure Data

Use the `exportMeasure` utility to export all measures for hierarchies. It exports only measures that have storage in the domain (db property).

```
exportMeasure -d [domain path] -hier [hierarchy1, hierarchy2] -outDir
[outputDirectory]
```

To enter multiple hierarchies, separate them with commas, for example, `-hier loc,clnd`. This exports all hierarchy measures that have storage in the domain.

If the output directory that you specify does not exist, it will be created. One output file is created for each HBI (higher based intersection) measure. (In a global domain, the HBI measures are stored in the master domain because the measure's base intersection is above the partition dimension. Forced non-HBI measures are measures that should be HBI measures and stored in the master domain, but the RPAS application has forced these measures to be stored across the local domains.) In addition, one file is created for each non-HBI or FnHBI measure per subdomain. The file names contain an internal subdomain index, for example, `sales.0.csv.rpl`, `sales.1.csv.rpl`, and so on.

For more information about `exportMeasure`, see ["Exporting Measure Data Using exportMeasure"](#).

Exporting Hierarchy Data

Use the `exportHier` utility to include user-defined dimensions. Export each hierarchy individually.

```
exportHier -d [domain path] -hier [hierarchy] -datFile [datFile] -udd
[-listInformal fileName]
```

Use the `-udd` argument to export the user-defined definitions. The `-udd` argument can only be used with the `-onlyFormal` or `-onlyInformal` options. It cannot be used with `-fixedWidth` option. User-defined dimensions can only be exported in CSV format.

If the `-listInformal` argument is used, `exportHier` also creates a file with the `fileName`, which contains a list of informal positions in the domain in a format that can be used by the `informalPositionMgr`. This option cannot be used with the `-onlyFormal` option.

For more information about `exportHier`, see ["Exporting Hierarchy Data Using exportHier"](#).

Purging Hierarchy Data

Use the `loadHier` utility with the `-purgeAll` argument to purge the hierarchy data. The `-purgeAll` argument purges formal, informal, and user-defined positions in the listed hierarchies. It cannot be used on a partition hierarchy or any system hierarchy.

```
loadHier -d [domain path] -purgeAll [hierarchy1,hierarchy2]
```

Use the `-purgeAll` argument to purge hierarchies. Specify multiple hierarchies by using CSV format (`-purgeAll loc,clnd`). For more information about `loadHier`, see ["Loading Hierarchies Using loadHier"](#).

Adding New Dimensions

Use the hierarchyMgr utility to add new dimensions. The hierarchyMgr utility parses the hierarchy.xml file and determines where to add the new dimensions. All start and width properties of all dimensions in the hierarchy are refreshed to be consistent with the new hierarchy.xml.

```
hierarchyMgr -d [domain path] -h [hierName] -addLevels hierarchy.xml
```

One or more dimensions can be added to any existing hierarchy except for partitioned or system hierarchies. Only one hierarchy can be parsed at a time. The original meta, hmain, and language databases are backed up and can be restored if a failure occurs.

Note:

- You cannot add or insert dimensions above or below user-defined dimensions. The hierarchy.xml file does not include user-defined dimensions.
 - Remove and move are not allowed.
 - This process is for inserting formal dimensions into the hierarchy. It is not for adding user-defined dimensions.
-
-

Reloading Formal Hierarchy Data

Use loadHier to load the hierarchy that was purged. The user-defined dimensions included in the input hierarchy file are also loaded.

```
loadHier -d [domain path] -load [hierarchy] [-includeUdd]  
or
```

```
loadHier -d [domain path] -loadAll [-includeUdd]
```

Use the -load argument to load one hierarchy at a time. Use the -loadAll argument to load all hierarchies that have an input file in the input directory of the domain.

Use the -includeUdd argument to load user-defined positions back to the domain. The data file must be in CSV format with a header line. The name of the data file should follow the current standard: <hierarchy name>.csv.dat.

All user-defined dimensions must be in the data file. Any missing user-defined dimensions can cause an error. All loaded positions are in formal status. See the ["Informalizing DPM Positions"](#) section for information on how to convert the status of previous DPM positions to informal.

For more information about loadHier, see ["Loading Hierarchies Using loadHier"](#).

Informalizing DPM Positions

Use informalPositionMgr to informalize DPM positions.

```
informalPositionMgr -d domainPath -hier hierName -operation informalize -file  
inputFile
```

Use the -file argument to enter the input file to be processed. The inputFile is the list file exported by exportHier with -listInformal.

For more information, see ["Informal Position Manager"](#).

Reloading Measure Data

Use `loadmeasure` to load all measure files located in the specified input directory.

```
loadmeasure -d [domain path] -inDir [inputDirectory]
```

Only `.rpl` files can be used with this option, and only the CSV format with a header line is supported. The exported files have one measure per file.

For more information, see ["Loading Measure Data Using loadmeasure"](#).

Adding New Hierarchy Dimensions Sample Script

Here is an example of a script that adds new hierarchy dimensions.

```
#!/bin/ksh
# -----
#
# This is a sample script demonstrates different steps to add levels to an
# existing domain
# -----
DOMPATH=$TEST_GLOBAL_DOMAIN

# Need to be full path.
MEAS_STORAGE_DIR="C:/tak/patchHierScript/measdata"
OLD_HIER_DATA_DIR="c:/tak/patchHierScript/oldhier"
NEW_HIER_DATA_DIR="c:/tak/patchHierScript/newhier"
# can be comma separated list. If multiple is intended, exportHier, hierarchyMgr,
# and loadHier needs to be called in a loop.
# ex. HIERS="loc,clnd"
HIERS="loc"

mkdir -p $MEAS_STORAGE_DIR
mkdir -p $OLD_HIER_DATA_DIR

# Inside NEW_HIER_DATA_DIR, there should be new hierarchy xml file and new hier
# data
# file.

IFS=","

# -----
# Step one, export all positions.
# -----
exportMeasure -d $DOMPATH -hier $HIERS -outDir $MEAS_STORAGE_DIR
if [ $? -ne 0 ]; then
    echo "exportMeasure for exporing $HIERS hierarchy"
    exit 1
fi

# -----
# Step two, export all hierarchy data and informal position list. One hierarchy
# at a time.
# -----
for h in "$HIERS"; do
    exportHier -d $DOMPATH -hier $h -datFile $OLD_HIER_DATA_DIR/$h.csv.dat -udd
    -listInformal $OLD_HIER_DATA_DIR/$h.informal.lst
    if [ $? -ne 0 ]; then
        echo "exportHier failed for $h"
        exit 2
    fi
done
```

```

        fi
    done

    # -----
    # Step 3, purge the hierarchies.
    # -----
    loadHier -d $DOMPATH -purgeAll $HIERS
    if [ $? -ne 0 ]; then
        echo "Unable to purge $HIERS."
        exit 3
    fi

    # BEGIN MODIFICATION and DATA LOADING.

    # -----
    # Step 4, patch the hierarchy (add levels) one hierarchy at a time.
    # -----
    for h in $HIERS; do
        hierarchyMgr -d $DOMPATH -h $h -addLevels $NEW_HIER_DATA_DIR/hierarchy.xml
        if [ $? -ne 0 ]; then
            echo "hierarchyMgr failed for $h"
            exit 4
        fi
    done

    # -----
    # Step 5, Load the hierarchy data back. The structure reflect the
    # new hierarchy structure.
    #
    # NOTE 1:
    # Data should be updated to include positions of the newly added dimensions
    #
    # NOTE 2:
    # Even though reclass is supported by loadHier, measure data is going
    # to be loaded assuming old position is still in the original subdomain.
    # Position reclassification is not recommended and may cause measure data lost.
    # -----
    for h in $HIERS; do
        # copy the data file to input directory of the domain
        cp $NEW_HIER_DATA_DIR/$h.csv.dat $DOMPATH/input/

        loadHier -d $DOMPATH -load $h -includeUdd
        if [ $? -ne 0 ]; then
            echo "loadHier failed while loading $h hierarchy"
            exit 5
        fi
    done

    # -----
    # Step 6, Set the previously DPM position back to DPM
    # -----
    for h in $HIERS; do
        informalPositionMgr -d $DOMPATH -hier $h -file $OLD_HIER_DATA_
        DIR/$h.informal.lst -operation informalize
        if [ $? -ne 0 ]; then
            echo "informalPositionMgr failed while informalizing positions in $h
            hierarchy"
            exit 6
        fi
    done

    #
    # -----

```

```
# Step 7, load measure data that exported with exportMeasure back to the
# domain.
# -----
loadmeasure -d $DOMPATH -inDir $MEAS_STORAGE_DIR
if [ $? -ne 0 ]; then
    echo "Unable to purge load measure data in $MEAS_STORAGE_DIR. Please check the
logs in output directory"
    exit 7
fi

# Completed
```

Transferring Data

The transferData utility allows you to load measure data into one domain by using a second domain as the data source. While this can be achieved with the exportMeasure and loadmeasure utilities, the transferData utility is more efficient since it uses a single process to load a data file into the destination domain.

The transferData utility works by copying the source measure data to the destination measure data along any matching external position names between the source and destination measures. This includes any informal positions that have matching names in the destination. You can choose to exclude informal positions' data from being copied to the destination if you want.

Note: You cannot use the transferData utility for forced non-HBI (higher based intersection) measures. Forced non-HBI measures are measures that should be HBI measures and stored in the master domain, but the RPAS application has forced these measures to be stored across the local domains. For these measures, you must transfer data using export and import utilities as described in ["Adding New Dimensions to Hierarchies"](#).

The measure names do not need to be identical in both domains. However, the source and destination measures must conform in intersection and data type.

The measures to be transferred are identified with the -measMap argument. The format of the argument is a colon-separated list of source measure name, destination measure name, and an optional mask measure name:

```
src1:dest1[:mask1]
```

Multiple measure groups can be specified in pairings if separated by a comma:

```
src1:dest1[:mask1],src2:dest2[:mask2]
```

For a complete example, see ["transferData Usage"](#).

The transferData utility can run in parallel when the destination domain is a global domain. The -processes argument determines how many child processes to spawn to perform the processing. This makes the transferData task faster when used in a global domain destination.

The transferData utility validates several components before transferring data. It validates that the source and destination domains exist, that the source and destination measures exist, and that the source and destination measures have conforming intersections and data types.

After it has validated everything, the `transferData` utility opens the source measure in read mode and opens the destination measure in write mode; this prevents data integrity issues. It iterates over positions in the source measure, pulls the values, and writes them to the destination measure.

The `transferData` utility cannot be used to transfer the data of a measure shared through a RPAS Data Mart. If `transferData` is called on a shared measure, `transferData` will exit with an error message identifying the shared measure.

See the section on `transferFactData` for information on transferring the content of a shared measure between a domain and a RPAS Data Mart.

Masks

You can use a mask measure to restrict the transfer of a range of positions from the source measure. This mask measure is a boolean measure, located in the source domain, that is at the same intersection of the source/destination measure pair or higher than that intersection.

Note: It should be noted that the mask measure should have an NA value of FALSE for optimal performance. The `transferData` utility will stop and inform the user if this is not the case.

In addition, you can use the `-clearDest` flag with or without a mask option. These combinations are described as follows:

No Mask with `-clearDest` Clears the destination arrays. Changes the destination array's NA value to the source measure's NA value at the beginning of the process. Loops over the source arrays and populates the destination.

No Mask without `-clearDest` Does not clear the destination arrays. Does not change the destination array's NA value. Loops over the source arrays and copies to the corresponding destination.

Mask with `-clearDest` Does not clear the destination arrays. Does not change the destination array's NA value. Loops over the mask and copies the source values (populated or unpopulated) to the corresponding destination cells.

Mask without `-clearDest` Does not clear the destination arrays. Does not change the destination array NA value. Loops over the mask and copies the populated source values to the corresponding destination cells. It does not copy unpopulated source values.

Table 7–15 illustrates the results of the mask and flag combinations.

Table 7–15 Mask and Flag Combinations and Outcomes

	No Mask, -clearDest Flag	No Mask, No Flag	Mask, -clearDest Flag	Mask, No Flag
Clears destination arrays	Yes	No	No	No

Table 7–15 (Cont.) Mask and Flag Combinations and Outcomes

	No Mask, -clearDest Flag	No Mask, No Flag	Mask, -clearDest Flag	Mask, No Flag
Changes the destination array's NA value to the source measure's NA value	Yes	No	No	No
Populates the destination arrays	Yes	Yes	Yes, populated and unpopulated source values	Yes, only populated source values

transferData Usage

```
transferdata -d destDomain -src srcDomain -measMap
"src1:dest1[:mask1],src2:dest2[:mask2]" {-processes n} {-clearDest} {-noInformal}
```

Table 7–16 provides descriptions of the arguments used by the transferData utility.

Table 7–16 transferData Utility Arguments

Argument	Description
-d pathToDestDomain	Specifies the path to the domain containing the destination measure. This is required.
-src pathToSrcDomain	Specifies the path to the domain containing the source measure. This is required.
-measMap	Indicates the comma-separated list of source to destination to mask mappings (src:dest[:mask]), where src is the name of the measure to be copied from the source domain, dest is the name of the measure to be written to in the destination domain, and mask is the name of the mask measure to be used, if any. A one-to-one relationship must exist between source measures and destination measures.
-processes max	Defines the maximum number of processes to run in parallel in a global domain setup.
-clearDest	When used with a mask option, ensures that the destination measures array is not cleared and its NA value is not changed. The utility copies the populated and unpopulated data from the source measure positions that are populated in the mask. When used without a mask option, the transferData utility clears the destination arrays and changes the destination array's NA value to the source NA value at the beginning of the process. After that, the populated measure positions from the source array
-noInformal	Ensures that only formal positions that have a matching external name in the destination measure are copied over.

Scenarios

The following sections describe the four scenarios for using the transferData utility.

Simple Domain to Simple Domain When the source and destination domains are simple domains, the data is copied from the source domain to the destination domain.

Simple Domain to Global Domain When the source is a simple domain and the destination is a global domain, and the destination measure is a non-HBI measure, then the data is copied to the local domains of the destination. HBI measure data is copied to the master domain. Since the destination domain is a global domain, the data is transferred in parallel.

Global Domain to Simple Domain When the source is a global domain, for non-HBI measures, the subdomain data array is copied to the destination. HBI measure data is copied from the master domain to the destination domain.

Global Domain to Global Domain Since the destination domain is a global domain, the data is transferred in parallel. However, each subprocess visits each source subdomain when transferring a non-HBI measure. This means that you do not have to run `transferData` over each local domain, but only on the master domain once.

Global-to-global transferring follows this process:

1. The `transferData` utility maps the source's external position names to the destination's external position names.
2. The utility scans the source.
3. The utility writes the values into the destination measure.

The `transferData` utility always runs in parallel when the destination domain is a global domain and the `-processes` option is set to a number greater than 1. When the source domain is a global domain and the source measure is a non-HBI measure, every subprocess running on each local domain of the destination domain visits every local domain on the source, even when the source local domain has no common position with that particular destination local domain. However, if there is no common position, the subprocesses visits should occur quickly.

FilterHier Utility

Sometimes, a retailer has a master file of hierarchy data that needs to be loaded into multiple domains. Some of these domains may be missing one or more levels from the master hierarchy, mostly because the planning levels in these domains are higher than the lowest level in the master and the domains do not need to have all the lower levels. For example, a retailer may have one domain for Merchandise Financial Planning where the lowest level is Category and another for Item Planning where the lowest level is Item. The hierarchies in these two domains have their relevant hierarchy load data in one master file, but using `loadhier`, the retailer cannot load just what is relevant to the domain from the master. System integrators need to write custom scripts to parse out irrelevant columns from the master file to prepare load files suited for individual domains.

The `filterHier` utility does the filtering of columns for the system integrators so that they do not need to write custom scripts. The utility analyzes the target domain and trims down the master file to only those columns that are needed by the target domain. The utility acts on CSV formatted files and requires the input file to contain a header line containing the names of the columns, for example, `SKU,SKU_label,STCO,STCO_label`. The output of the utility is a `.csv.dat` file that can be subsequently used by the `loadHier` utility.

filterHier Usage

```
filterHier -d domainPath -input inputPath [COMMAND] {OPTIONS}
```

Table 7–17 provides descriptions of the arguments used by the filterHier utility.

Table 7–17 filterHier Utility Arguments

Argument	Description
-d domainPath	Indicates the domain in which to load the hierarchy data.
-input inputpath	Indicates the path to the folder where the master files are located.
-filter hiername	Filters the hierarchy named in the parameter to the command.
-filterAll	Filters all hierarchies in the input directory that are relevant to the target domain.
-compress	Creates a compressed .csv.dat output file. RPAS provides a simple, proprietary compression technique to help reduce the file size and file I/O time during loads. This technique simply replaces a column's value with a '?' character to indicate that the column's value for the row matches that of the row above. The compressed file continues to print out '?' characters for a column until a change is encountered. This kind of compression is useful for hierarchy files where the lowest level positions are grouped by the higher level positions. In such cases, the output file will print out '?' characters for higher level positions until a change is encountered, thus significantly reducing the file size. Note that compressed files should not be split up or reprocessed in ways that change the order of rows.

filterHier Notes

This utility combines one or more input files into an output file that can be imported into the target domain using loadHier. The input files must be csv data containing a comma-separated header line listing the name of each column. Column names must be in the format DIM,DIM_label.

Example: SKU,SKU_label,STCO,STCO_label,SIZ1,SIZ1_label

The input files must have the extension .hdr.csv.dat. Optional extensions are allowed at the end of the file.

Example: prod.hdr.csv.dat.foo1

The columns in the output file are arranged to match the hierarchy format of the target domain. Any dimensions from the input files that are not present in the output domain are filtered out.

The output file is a properly formatted .csv.dat file, and is located in the input directory of the target domain.

Error Conditions

The following error conditions may occur during the operation of filterHier:

- Dimension Not Found: If a dimension that exists in the domain is not found in the header, the input file will be skipped.
- No Usable Input Files: If filterHier cannot find a usable input file, it will exit with error.
- Parse Error in Data: If one or more data rows contain an error, filterHier will display an error specifying which lines contain an error and continue processing the rest of the file.

- **Conflicting Base Positions:** If a base position has multiple definitions in the input files, the first definition will be used and all others will be skipped.

Position Repartitioning

Position repartitioning is the automated process of moving positions and all corresponding measure data between local domains. This functionality is only available (and relevant) in global domain environments. Positions must be moved between local domains when they are assigned a new parent that exists in a different local domain. Note that moving positions at the partitioning level is a manual process and requires the use of the `reconfigGlobalDomainPartitions` utility.

For example, imagine `Style1` belongs to `Sub-Class1` in `LocalDomain1`. If `Style1` is reassigned to be a child of `Sub-Class2`, which is located in `LocalDomain2`, RPAS will move the `Style1` position, `Style1`'s children (if any), and all corresponding data to `LocalDomain2`. This process is often referred to as reclassification by RPAS customers. RPAS refers to this functionality as position repartitioning because it technically does not handle the many complex functional requirements of true reclassification as most retailers define the term to mean.

Loading RDF and Curve Parameters after Repartitioning

After repartitioning, default parameters for Curve and RDF are not automatically loaded in new subdomains. To load these parameters, the following scripts, which are located in the `RPAS_HOME/bin` directory, need to be executed:

- `loadCurveParameters.ksh`: Used to load Curve parameters after a repartition.
- `loadRdfParameters.ksh`: Used to load RDF parameters after a repartition.

These scripts are used to load RDF and Curve parameters to a subdomain created as a result of repartitioning. These parameters are usually loaded during a full installation by the plug-ins, but when performing a patchinstall, the parameters are not loaded by default. These parameters include default required method, default source, spreading profile, and others.

You must run these scripts after repartitioning a domain on the new partition.

Syntax

```
scriptname -d <full path to domain> -s <full path to subdomain>
```

Example:

```
loadRdfParameters.ksh -d /vol.nas/forecast/domains/RDF_12 -s /vol.nas/forecast/  
/domains/RDF_12/ldom0/
```

Reconfiguring Partitions of a Global Domain Using `reconfigGlobalDomainPartitions`

It is common for many customers to regularly add, remove, or change the parent-child relationships for positions in hierarchies, most commonly for positions in the product hierarchy. While this movement and reassignment of positions is normally handled automatically within the `loadHier` utility, a special process must be followed for positions at the partition level of a global domain environment.

The RPAS utility `reconfigGlobalDomainPartitions` is used for the following activities in a global domain environment:

- Adding a new position along the partition dimension and allocate it to an existing or new local domain.
- Removing an existing position from the partition dimension.
- Removing local domains (this is automatic if all partition-level positions in a local domain are removed or moved).
- Moving an existing partition position from one local domain to an existing or new local domain.

Runs loadHier to apply the position addition or removal to the hierarchy.

Note: This utility can only be used on a master domain of a global domain set.

The following processes must be followed to add, remove, or move positions at the partition level in a global domain environment:

- The administrator must be notified in advance that positions at the partition level are being added, removed, or moved.
- The administrator must run the utility reconfigGlobalDomainPartitions by specifying the subdomain to which the positions do or will belong.
- This utility calls the loadHier utility at the end of the reconfiguration process to apply the hierarchy changes to the domain. When positions are added (using the -add argument) an updated hierarchy file must be available in the input directory when the reconfigGlobalDomainPartitions utility is called. Otherwise the utility will fail. Updated hierarchy files are not required to remove (using the -remove argument) or move positions (using the -move argument).

Note: The use of this utility is only required for positions at the partition level. Positions below the partition level can be added, removed, or moved between local domains by loading a modified hierarchy input file with these changes.

When operating on a domain integrated with an RPAS Data Mart, reconfigureGlobalDomainPartitions updates the ITT tables of the RDM to reflect the changes made as a result of its execution. Additionally, reconfigureGlobalDomainPartitions updates the RDM status tables to reflect the fact that a reconfigure operation is in progress. See the "Domain Integration and the RPAS Data Mart" section for more information on domain operations in an integrated environment.

reconfigGlobalDomainPartitions Usage

```
reconfigGlobalDomainPartitions -d pathToMasterDomain -remove posName1, posName2,
...
reconfigGlobalDomainPartitions -d pathToMasterDomain -add posName1,posName2, ...
-sub pathToSubDomain
reconfigGlobalDomainPartitions -d pathToMasterDomain -move posName1,posName2, ...
-sub pathToSubDomain
reconfigGlobalDomainPartitions -d pathToMasterDomain -input pathToInputDir
```

Table 7–18 provides descriptions of the arguments used by the reconfigGlobalDomainPartitions utility.

Table 7–18 reconfigGlobalDomainPartitions Utility Arguments

Argument	Description
-d <i>pathToMasterDomain</i>	Specifies the path to the master domain in a global domain environment.
-add <i>posName1, posName2, ...</i>	Adds one or more positions at the partition level to a specified local domain. The path to the local domain must follow the list of positions to add, using the -sub argument. If the specified path is to a local domain that does not yet exist, the system will create a new local domain with the specified positions at the partition level. This argument cannot be used with -remove or -input.
-remove <i>posName1, posName2, ...</i>	Removes the designated positions from the local domain to which the positions belong. The path to the local domain does not need to be specified with this argument. The local domain will be deleted if all the positions at the partition level in a local domain are removed. This argument cannot be used with -add or -input.
-move <i>posName1, posName2, ...</i>	Moves the specified positions at the partition level from the current domain in which the positions are located to the specified local domain. This argument requires specification of the -sub argument. To move positions, all dimensions below the partition level must be enabled for dummy positions.
-sub <i>pathToSubDomain</i>	Specifies the path to the local domain to which positions are being added or the destination local domain for positions being moved. When a new domain path is specified using -sub option, a new local domain is created. This argument is required for the -add argument and -move argument.
-input <i>pathToInputDir</i>	Specifies the path to the input directory that contains an xml configuration file (reconfigpartdim.xml) to specify positions to either add or move. The file must have all the information to run the process including the command name, position names to add or move, and paths to the local domains. This option is useful for adding or moving positions to multiple local domains. This argument does not handle both adding and moving in the same call. This argument cannot be used with -add or -remove.
-maxProcesses <i>count</i>	If specified, some parts of reconfig utility will run in parallel, using up to the given number of processes.
-forceInputRollups	Prevents this utility from failing in instances where there is a rollup conflict in the input file provided to the utility. This argument enforces new hierarchy rollup changes so that they dominate existing hierarchy rollups in case they conflict with the rollups specified in the input file.

Using an Input File

When using the -input argument, the file must be in a particular format and must contain the add or move commands, the path to each local domain to which positions are being added or the destination for positions being moved, and the list of positions for each local domain. The file must be XML and named **reconfigpartdim.xml**.

Note: The -input argument only supports the addition or movement of positions.

Here is the required format of the input file:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes" ?>
<rpas>
  <command name="command_name">
    <subdomain>
      <subpath>path_to_local_domain_1</subpath>
      <subpositions>sample_pos_1</subpositions>
    </subdomain>
    <subdomain>
      <subpath>path_to_local_domain_2</subpath>
      <subpositions>sample_pos_2,sample_pos_3</subpositions>
    </subdomain>
  </command>
</rpas>
```

Note: The entries in bold are the parameters that must be specified in the input file.

Table 7–19 provides descriptions of the arguments used by the input file.

Table 7–19 Arguments Used by the input File

Argument	Description
command_name	Valid values are add or move.
path_to_local_domain	Indicates the path to the local domain to which positions are added or moved.
sample_pos	Indicates one or more positions that are being added or moved to the designated local domains.

Notes, Assumptions, and Limitations

Note the following:

- Position names are separated by commas and must be valid external position names without the prefix of a dimension (for example, 0102,0144,0152,0160).
- The utility backs up the required data and automatically restores the domains to the original state in case of failure.
- In a single call to the utility without using the -input argument, positions can only be added or removed or moved. That is, the -add, -remove, and -move arguments cannot be mixed in the same call.
- Multiple positions can be added or moved to a single local domain in a single call to the utility using the -add or -move option, respectively.
- Multiple positions can be added or moved to multiple local domains in a single call to the utility using the -input option.
- When positions are added at the partition level, an updated hierarchy file must be available in the input directory when running this utility as the loadhier utility is called after adding positions. If the updated hierarchy file is not present in the input directory when attempting to add positions, the utility will fail.

- No updated hierarchy file is required when moving or removing positions. If a hierarchy file is in the input directory, the utility will back up this file.

Renaming Positions Using renamePositions

RPAS provides the ability to change the name of a position using an RPAS utility named `renamePositions`. Positions that are to be renamed should be included in a hierarchy data file. After the hierarchy data files have been updated and placed in the designated location, an administrator must run the `renamePositions` utility.

Note: The `renamePositions` utility cannot be used to rename any RPAS system objects such as dimensions, workbooks, hierarchies, and so on.

When executing on a domain that shares data through an RPAS Data Mart, it may be necessary to run `renamePositions` on all domains interacting with the RPAS Data Mart. If the call to `renamePositions` affects positions used in sharing data between domains, all domains must be updated by `renamePositions` using the same input file to reflect the change in position names.

When renaming:

- Positions of on-shared dimensions, run only on domain in question.
- Informal positions of shared dimensions, run only on domain in question.
- Formal positions of shared dimensions, run on all domains using the dimension in question.

renamePositions Usage

```
renamePositions -d domainPath -hier hierName {-input inputDirectory} {-log  
logFileName} {-dryRun}
```

The `renamePositions` utility renames positions in a specified hierarchy. It looks for a file named `hierName.rn.dat` (for example, `prod.rn.dat`) in the domain's input directory when `-input` is not specified.

The input file is a CSV text file with three columns. Tab-delimited fields are not supported in the input file. This is consistent with other domain utilities in which tab-delimited fields are not supported. One line is required per rename request. Each line consists of a dimension name, the old position name, and a new position name. For example:

```
sku,10000_old,10000_new
```

The utility ignores any line not formatted this way, empty lines, and lines with a specified dimension that does not exist for the specified hierarchy.

The input file is not case sensitive, and all names are converted to lower case before the positions are renamed. Old position names that are specified but that do not actually exist are ignored. New position names cannot be names that are already in use. Lines with invalid position names are ignored and added to the log file. The old position name and new position name should not be prefixed with the name of the dimension.

The width property of the dimension is enforced. The width attribute in the domain must be greater than or equal to the maximum length of the new names in the input

file. If the width of the new name is greater than the width attribute of the corresponding dimension, RPAS prints an error in the log file and ignores the record.

Dimensions specified in the input file should belong to the hierarchy that is specified in arguments. Otherwise, the record is ignored, and RPAS prints an error in the log file.

Note: If you want to swap position names (a->b, b->a), then you must perform these two separate operations: (a->tempa,b->tempb) and (tempa->b, tempb->a).

Table 7–20 provides descriptions of the arguments used by the renamePositions utility.

Table 7–20 Arguments Used by the renamePositions Utility

Argument	Description
-d <i>domainPath</i>	Specifies the full path to the domain.
-hier <i>hierName</i>	Specifies the hierarchy in which positions are being renamed.
-input <i>inputDirectory</i>	Specifies the input directory where the input file with positions to rename is located. Optional. Utility looks for hierarchy data files with "rn" between hierarchy name and .dat extension (for instance, prod.rn.dat).
-log <i>logFileName</i>	Generates a log file to file name other than default. The default file name is hierRename.log. Optional. This argument can be used to name the log file other than the default, which is created as hierRename.log in the current directory.
-dryRun	Does not apply changes, but generates a report that identifies changes that would be applied. This argument generates a log file.

Managing Dimension Attributes Using dattrmgr

The dattrmgr utility is a command-line utility used to manage the dimension attributes in a domain including both static and dynamic attributes. It can be run only on the master domain or simple domains. This utility provides the user the option to register new dimension attributes or unregister dimension attributes that are already registered in the domain. There is also a report option to list out all the static and dynamic dimension attributes that are present in a domain. The report option also prints out the saved split definitions and workbooks where the dimension attributes are used.

dattrmgr Usage

```
dattrmgr -d pathToDomain [COMMAND [OPTIONS]]
```

Table 7–21 provides descriptions of the arguments used by the dattrmgr utility.

Table 7–21 Arguments Used by the dattrmgr Utility

Argument	Description
-d <i>pathToDomain</i>	Specifies the directory in which to run the utility. All arguments except -version require -d pathToDomain.

Table 7–21 (Cont.) Arguments Used by the dattrmgr Utility

Argument	Description
-register <i>measName</i>	Name of a registered one-dimensional measure in the domain. This measure is used to register the dimension attribute.
-dattname <i>attrName</i>	Optional parameter that can be used with -register to specify the name of the dimension attribute. If this is not supplied, the measure name is used for the dimension attribute name.
-dattlabel <i>attrLabel</i>	Optional parameter that can be used with -register to specify the label of the new dimension attribute. If this is not supplied, the measure label is used as the label of the new dimension attribute.
-report	Lists all the static and dynamic dimension attributes that are present in the domain. Also prints out the saved split definitions and workbooks where the dimension attributes are used.
-version	Prints the version of the utility.

dattrmgr -report Examples

Here is a sample output of the -report option. Note that the same split definition may be used in multiple workbooks and there may be split definitions that are not used in any workbooks. The split definition ID is printed under saved splits and also under workbooks.

```

Name          : attr1
Label         : measA
Dim           : sku
Meas          : measa
              :
Saved splits  :
              :
              id      : splt_dm001
              Label   : split1
              User    : adm
              :
              id      : splt_dm3
              Label   : split3
              User    : adm
              :
Workbooks     :
              :
              id      : t0000000000
              Label   : 1
              Split id : splt_dm001
              :
              id      : t0000000002
              Label   : 3
              Split id : splt_dm3
-----
Name          : datt_1
Label         : dynAttr1
Dim           : stco
Meas          : datt_1
              :
Saved splits  :
              :
              id      : splt_dm2
              Label   : split2

```

```

User      : adm
          :
Workbooks :
          :
id        : t0000000001
Label     : 2
Split id  : splt_dm2
-----

```

Setting Properties for Dimensions Using dimensionMgr

The dimension manager utility is used for setting a number of parameters for dimensions and positions. Multiple command arguments are allowed.

dimensionMgr Usage

```
dimensionMgr -d pathToDomain -dim dimensionName [COMMAND] {Options}
```

[Table 7–22](#) provides descriptions of the arguments used by the dimensionMgr utility.

Note: This utility includes arguments that are not documented in this guide as it is recommended that those operations be configured using the Configuration Tools to ensure consistency between the configuration and the domain.

Table 7–22 *dimensionMgr Utility Arguments*

Argument	Description
-d <i>pathToDomain</i>	Specifies the path to the domain.
-dim <i>dimensionName</i>	Specifies the name of the dimension to which the settings will apply.
-label <i>newLabel</i>	Specifies a new dimension label.
-enableDPM	Enables Dynamic Position Maintenance (DPM) for the specified dimension (-dim <i>dimensionName</i>). Using this argument not only enables DPM for the specified dimension, but for all dimensions that roll up to it. DPM cannot be enabled for any dimension in an RPAS-internal hierarchy (DATA, META, RGPS, ADMU, ADMW, MSGS, LNKS).
-enableTranslation <i>width</i>	Enables the specified dimension to use translated labels. When enabling translated labels, the numeric parameter passed in the argument (<i>width</i>) defines the field width for the translated values in the data file, which is loaded using the loadMeasure utility.

Table 7–22 (Cont.) dimensionMgr Utility Arguments

Argument	Description
-specs	Displays the properties of the specified dimension. The dimension properties indicate whether DPM and translation are enabled for the dimension and whether the dimension is image-enabled. Here is the full list of properties that the -specs argument returns: ■ Dimension Label ■ Associated Hierarchy of Dimension ■ Start and Width for File Load ■ Label Start and Width for File Load ■ Number of Used Positions ■ Image Enabled ■ Bit Size ■ Capacity ■ Used ■ Reindex Threshold ■ Translation Enabled ■ Spreads or Aggs (meaning, where the dimension fits in the rollup) ■ If the dimension is at the base level of a hierarchy, that is indicated as well.
-width <i>widthVal</i>	Sets the width of position names for the specified dimension. After you have created a position name, you can only increase its width; you cannot decrease the width.
bitSize	Changes the bit size for the dimension, but does not actually update the dimension arrays. You must run <code>reindexDomain</code> for the dimension data and measure arrays to be updated.
reindexThreshold	Sets a new reindex threshold for the dimension.

Changing the Label of a Dimension

To change the label of a dimension, use the `dimensionMgr` utility. It is not restricted to partitioned hierarchies.

```
dimensionMgr -d domainPath -dim dimName -label newLabel
```

Exporting Hierarchy Data Using exportHier

The `exportHier` utility is used to export all the positions in a hierarchy, including their rollup relations. The `exportHier` utility cannot be run on local domains. It must be run only on the master domain or simple domains. By default, the utility assumes that the file has a CSV flat file format with fixed-width format as an optional argument. The utility exports all hierarchy positions, but the file may be specified to include only formal or informal positions. The resulting file can then be used as a .DAT file with `loadHier`.

exportHier Usage

```
exportHier -d domainPath -hier hier_name -datFile dat_file [-fixedWidth]
[-onlyFormal | -onlyInformal][-upperCase]
```

Table 7–23 provides descriptions of the arguments used by the exportHier utility.

Table 7–23 Arguments Used by the exportHier Utility

Argument	Description
-d <i>pathToDomain</i>	Specifies the path to the domain.
-hier <i>hier_name</i>	Indicates the name of a hierarchy in the domain from which the .DAT file is generated.
-datFile <i>dat_File</i>	Indicates the path/location where the .DAT file is created. This .DAT file can then be used with loadHier to load the hierarchy into a domain.
-fixedWidth	Indicates that the output .DAT file is a fixed-width file instead of a comma-separated value (CSV) file format.
-onlyFormal	Exports only formal positions to the .DAT file. If this option is specified, informal positions will be skipped.
-onlyInformal	Exports only informal positions to the .DAT file. If this option is specified, formal positions will be skipped.
-genheader	Generates a header line for CSV export. The line contains fields to identify the dimension and its label column.
-udd	Exports the user-defined definitions. It can only be used with the -onlyFormal or -onlyInformal options. It cannot be used with the -fixedWidth option. User-defined dimensions can only be exported in CSV format.
-listformal	Creates a file with the fileName, which contains a list of the informal positions in the domain in a format that can be used by the informalPositionMgr. This option cannot be used with the -onlyFormal option.
-upperCase	Converts the position names to all uppercase before writing the output data file. Without this argument, position names are in lowercase since they are stored in lowercase in the domain.

Informal Position Manager

Informal Position Manager is a domain utility that maintains informal positions for any dimension in a domain. Informal positions are those created with the Dynamic Position Manager (DPM). Formal positions are those created with the loadHier utility. This utility can convert positions from formal to informal or from informal to formal. It can also remove informal positions, create informal positions in bulk, and copy data slices between positions in measures.

InformalPositionMgr Usage

The three main areas of use for informalPositionMgr are described in the following sections:

- [Basic Operations: Informalize, Formalize, Remove, and Load](#)
- [Create Informal Positions in Bulk](#)
- [Data Slice Copying](#)

Basic Operations: Informalize, Formalize, Remove, and Load

Informalizing, formalizing, and removing informal positions all have similar command line options.

```
informalPositionMgr -d domainPath -hier hierName -operation [remove | formalize |
informalize | load] [-dir inputDir | -file inputFile] [-retain]
```

Table 7–24 provides descriptions of the arguments used in basic operations of the `informalPositionMgr` utility.

Table 7–24 Arguments Used by the `informalPositionMgr` Utility: Basic Operations

Argument	Description
<code>-d <i>domainPath</i></code>	Indicates the domain to run the utility.
<code>-hier <i>hierName</i></code>	Specifies the hierarchy to operate on.
<code>-operation <i>informalize</i> <i>formalize</i> <i>remove</i> <i>load</i></code>	Specifies one of the three basic operations: <code>informalize</code>, <code>formalize</code>, or <code>remove</code>. You can use only one of these operations at a time. <code>remove</code>: Removes informal positions of the specified hierarchy. <code>formalize</code>: Formalizes positions of the specified hierarchy. <code>informalize</code>: Informalizes positions of the specified hierarchy. <code>load</code>: Load informal positions from one or more input files. All new positions are treated as informal. Can be used to change roll-up for informal positions. Does not change roll-up for formal positions. In a domain sharing data through a RPAS Data Mart, <code>informalPositionMgr</code> cannot informalize formal positions of a shared hierarchy.
<code>-file <i>inputFile</i></code> or <code>-dir <i>inputDir</i></code>	Use one and only one of these arguments to specify the input files to use during the operation specified. <code>-dir <i>inputDirectory</i></code>: Processes all applicable files under <code>inputDirectory</code>. <code>-file <i>inputFile</i></code>: Processes the <code>inputFile</code> only. All files that match the naming patterns below are processed: - Remove files: <code>{hierName}.remove[.extension]</code> - Formalize files: <code>{hierName}.formalize[.extension]</code> - Informalize files: <code>{hierName}.informalize[.extension]</code> - Load files: <code>{hierName}.informal.csv.ovr[.extension]</code> For example: <code>prod.informalize.20100220</code> The format of the above files except load files has one position on each line. The position has a dimension name and a position name that are comma delimited. There is no header line: <pre>dimName1, positionName1 dimName2, positionName2</pre> The format of load files is the same as that of the load files for <code>loadHier</code>.
<code>-retain</code>	Use if you do not want to move the input files to the processed subdirectory after a successful run.

Error Handling

The following errors result the described behavior.

Table 7–25 *informalPositionMgr Error Handling*

Error	Behavior
Input file does not exist	A log error is generated and the operation stops.
Any dimension specified in the input file is not DPM enabled or does not exist	A log error is generated and the operation stops.
One or more positions does not exist	Log warnings are generated, the line is skipped, and the operation continues.
Selecting to convert a position to informal that is already informal OR Selecting to convert a position to formal that is already formal	No action is taken on that position, a warning message is logged, and the operation continues.
Selecting to convert a position to informal in a shared dimension.	No action is taken on that position. A warning message is logged, and the operation continues.

Create Informal Positions in Bulk

This feature allows the user to create a number of informal positions on any DPM-enabled dimension. These positions are automatically named and labeled. However, the user must provide applicable rollup and spread information so that these new positions can be properly placed in the hierarchy.

```
informalPositionMgr -d domainPath -create -dim dimName -n posCount -rollups
dim1:pos1,dim2:pos2,... [-spreads dimA:ratioA,dimB:ratioB,...]
```

For bulk creation of informal positions, the position names and labels are auto-generated. To ensure the names are unique, a prefix and sequence number is used and the width and label width attributes of the dimension should be at least 10. If the width or label width is less than 10, bulk creation is not allowed, and an exception is thrown.

If the dimension specified by option `-dim` is above the root dimension, informal positions are added to the lower levels. The numbers of spread positions for lower level dimensions are defined by the `-spreads` option. Note that the spread ratio determines the number of children for each position at the level immediately above the specified dimension. If it is not specified, the default is 1.

If these new descendent positions roll up to other alternate dimensions, there are a few options to handle the rollups. If those dimensions are DPM-enabled, you can either specify the rollup positions (in the `-rollups` option) or let the program create single rollup position for you. This repeats recursively for those new positions. If any of those affected dimensions are not DPM-enabled, the rollup positions must be specified.

[Table 7–26](#) provides descriptions of the arguments used by the create bulk positions functionality of the `informalPositionMgr` utility.


Table 7–26 *Arguments Used by the informalPositionMgr Utility: Create Bulk Positions*

Argument	Description
<code>-d domainPath</code>	Indicates the domain to run the utility.
<code>-create</code>	Batch creates new positions with auto-generated names.
<code>-dim dimName</code>	Specifies the dimension to operate on.

Table 7–26 (Cont.) Arguments Used by the informalPositionMgr Utility: Create Bulk

Argument	Description
-rollups <i>dim1:pos1,dim2:pos2,...</i>	Specifies the rollup positions for the auto-generated positions. You should include the direct rollups for the new positions and rollups for the descendents of these new positions on alternate branches, if applicable.
-spreads <i>dimA:ratioA,dimB:ratioB,...</i>	Specifies the spread counts for the auto-generated positions. This is optional. It defaults to 1 if not specified.
-n <i>posCount</i>	Use to enter the number of positions to create on the specified dimension.

The following sections provide some examples of creating informal positions in bulk on the product hierarchy. [Figure 7–2](#) displays a product hierarchy with alternative branches and four DPM-enabled dimensions (STYL, STCO, SKU, and SZ12).

Figure 7–2 Sample Product Hierarchy with Alternative Branches

Example 1: Creating Informal Positions on the STYL Dimension

To create informal positions on the STYL dimension, you must specify rollups for SCLS, CLR, and SPLR. You have the following two options for SZ12:

Option 1

Specify the rollup for SZ12.


```
informalPositionMgr -d domainPath -create -dim STYL -n 100 -rollups
SCLS:scls1001,CLR:clrwhite,SPLR:splr100,SZ12:sz12null -spreads STCO:2,SKU:5
```

This command creates the following:

1. The 100 new informal positions on the STYL dimension. These positions all roll up to SCLS:scls1001.
2. Since the spread ratio for STCO dimension is 2, then 200 new informal positions are created at the STCO level, and every two of them roll up to one STYLE position that was created in step 1. These new informal positions all roll up to CLR:clrwhite on the alternate branch.
3. Similarly, 1000 new informal positions are created at the SKU level, and every 5 of them roll up to one STCO position that was created in step 2. They all roll up to SPLR:splr100 and SZ12:sz12null on the two alternate branches.

Option 2

Specify the rollups for SIZ1 and SIZ2. A new position is automatically created on SZ12. You can specify spread ratios for STCO and SKU. If you do not specify the spread ratio, it defaults to 1.

```
informalPositionMgr -d domainPath -create -dim STYL -n 100 -rollups
SCLS:scls1001,CLR:clrwhite,SPLR:splr100,SIZ1:siz1null,SIZ2:siz2lt02 -spreads
STCO:2,SKU:5
```

This command creates the following:

1. The 100 new informal positions on the STYL dimension. These positions all roll up to SCLS:scls1001.
2. One new informal position at the SZ12 level, which rolls up to SIZ1:siz1null and SIZ2:siz2lt02.
3. Since the spread ratio for STCO dimension is 2, then 200 new informal positions are created at the STCO level, and every two of them roll up to one STYL position created in step 1. These new informal positions all roll up to CLR:clrwhite on the alternate branch.
4. Similarly, 1000 new informal positions are created at the SKU level and every 5 of them roll up to one STCO position created in step 3. They all roll up to SPLR:splr100 and the new SZ12 position created in step 2.

Example 2: Creating Informal Positions on the SZ12 Dimension

To create informal positions on SZ12 dimension, you must specify rollups for SIZ1, SIZ2, and SPLR and specify the spread ratio for SKU. For the DPM-enabled dimension STCO, you have the following two options:

Option 1

Specify a rollup position.

```
informalPositionMgr -d domainPath -create -dim SZ12 -n 10 -rollups
SPLR:splr100,SIZ1:siz1null,SIZ2:siz2lt02,STCO:stco2000 -spreads SKU:5
```

This command creates the following:

1. The 10 new informal positions on the SZ12 dimension. These positions all roll up to SIZ1:siz1null and SIZ2:siz2lt02.

2. Since the spread ratio for SKU dimension is 5, then 50 new informal positions are created at the SKU level, and every 5 of them roll up to one position at the SZ12 level, which was created in step 1. These new informal positions all roll up to SPLR:splr100 and STCO:stco2000.

Option 2

Create a rollup position and specify a rollup position on CLR and rollup position on STYL (or create recursively until the non-DPM enabled dimension is reached).

```
informalPositionMgr -d domainPath -create -dim SZ12 -n 10 -rollups  
SPLR:splr100,SIZ1:siz1null,SIZ2:siz2lt02,SCLS:scls1001,CLR:clrwhite, -spreads  
SKU:5
```

This command creates the following:

1. The 10 new informal positions on the SZ12 dimension. These positions all roll up to SIZ1:siz1null and SIZ2:siz2lt02.
2. One new informal position at the STYL level, which rolls up to SCLS:scls1001.
3. One new informal position at the STCO level, which rolls up to CLR:clrwhite, and the new position at STYL level, which was created in step 2.
4. Since the spread ratio for SKU dimension is 5, then 50 new informal positions are created at the SKU level, and every 5 of them roll up to one position at the SZ12 level which was created in step 1. These new informal positions all roll up to SPLR:splr100 and the new STCO position created in step 3.

Position Naming Convention

When the bulk creation feature is used, the names and labels of the new positions are created with the following name conventions:

- **Position name:** {dimName}{dpm}{seq#}

The width is an attribute of the dimension. For auto-generated position names, the minimum width required is 10.

- For width = 10, {4 char dimName}d{5 char seq#}
- For width = 11, {4 char dimName}d{6 char seq#}
- For width = 12, {4 char dimName}d{7 char seq#}
- For width = 13, {4 char dimName}dp{7 char seq#}
- For width >= 14, {4 char dimName}dpm{7 char seq#}

If the dimName is less than four characters, an underscore is appended, for example: sku_. Each dimension uses its own persistent sequence generator in the domain with a unique label DPM_{dimName}. For a global domain, the sequence generators are in the master domain.

- **Position label:** Follows the same convention as that of position name, but uses the label width attribute of the domain.

Data Slice Copying

This feature allows you to copy data slices from one position to those of another. It is useful when you want to merge informal positions with existing formal positions. An inclusive or exclusive measure list can be specified. A Boolean type mask measure can be used to enable selective copying of data slices.

When operating on a domain that shares data through an RPAS Data Mart, if `informalPositionMgr` is used to copy a slice of data involving a dimension shared in the RDM, the slice of data will also be copied within the RDM.

```
informalPositionMgr -d domainPath -copy -file inputXMLFile [-retain]
```

[Table 7-27](#) provides descriptions of the arguments used by the data slice copying functionality of the `informalPositionMgr` utility.

Table 7-27 Arguments Used by the `informalPositionMgr` Utility: Data Slice Copying

Argument	Description
<code>-d domainPath</code>	Indicates the domain to run the utility.
<code>-copy</code>	Copies data slices within measures.
<code>-retain</code>	Use if you do not want to move the input file to the processed subdirectory after a successful run.
<code>-file inputXMLFile</code>	Specifies the XML file that contains the configuration settings for the copying operations.

[Table 7-28](#) shows the available settings for the XML input file. This file defines the configuration settings for the copying operation.

Table 7-28 XML Input File Settings

Section	Attribute	Description	Required	Format
copy		Settings for copying operation.		
copy	removeSource	Removes source data slices after copying.	Yes	"true" or "false"
measures		A list of measures.	No	Comma delimited measure names
measures	option	The measure list is inclusive or exclusive.	Yes, if the measures section exists	"include" or "exclude"
mask		Mask measure.	No	One Boolean type measure name
positionMap		Position mapping for the copying operation.	Yes	
positionMap	dim	The dimension of the mapped positions.	Yes	Dimension name
positions		Two positions. The first one is the position to copy from. The second one is the one to copy to.	Yes	Comma delimited

Here is a sample XML Input File Schema in XSD format file.

```
<?xml version="1.0" encoding="UTF-8"?>
<xsd:schema attributeFormDefault="unqualified" elementFormDefault="qualified"
version="1.0" xmlns:xsd=http://www.w3.org/2001/XMLSchema">
  <xsd:element name="rpas" type="rpasType" />
  <xsd:complexType name="rpasType">
    <xsd:sequence>
      <xsd:element name="copy" type="copyType" />
    </xsd:sequence>
  </xsd:complexType>
```

```

<xsd:complexType name="copyType">
  <xsd:sequence>
    <xsd:element name="measure" type="measuresType" />
    <xsd:element name="mask" type="xsd:string" />
    <xsd:element name="positionMap" type="positionMapType" />
  </xsd:sequence>
  <xsd:attribute name="removesource" type="xsd:boolean" />
</xsd:complexType>
<xsd:complexType name="positionMapType">
  <xsd:sequence>
    <xsd:element maxOccurs="unbounded" name="positions" type="xsd:string" />
  </xsd:sequence>
  <xsd:attribute name="dim" type="xsd:string" />
</xsd:complexType>
<xsd:complexType name="measuresType">
  <xsd:attribute name="option" type="xsd:string" />
</xsd:complexType>
</xsd:schema>

```

Here is a sample XML Input file.

```

<?xml version="1.0" encoding="UTF-8"?>
<rpas>
  <copy removeSource="true">
    <measures option="include">
      r_ex_demoa, r_ex_demob, r_ex_democ
    </measures>
    <mask>maskMeasure</mask>
    <positionMap dim="sku">
      <positions>sku_10000010, sku_10000008</positions>
      <positions>sku_22200001, sku_22200002</positions>
    </positionMap>
  </copy>
</rpas>

```

Merging Informal Positions to Formal Positions

Here are the steps to merge informal positions to formal positions in a domain.

1. Create the XML input file by specifying the position pairs to merge. Optionally, you can specify the measures to update and a mask measure.
2. Create the mask measure if you want to use a filter to protect some data slices from being modified. For example if you want to leave elapsed data alone, you can create a calendar based mask measure and set all the elapsed cells to false.
3. Run informalPositionMgr with -copy option to copy the data slices.
4. Create an input file for the removal of informal positions you no longer need.
5. Run informalPositionMgr with -operation remove to remove those informal positions.

Managing Position Lists as PQDs Using pqdMgr

The pqdMgr utility is a command-line utility used to add and delete Position Query Definitions (PQDs) using XML-formatted position lists. A position list is a multi-level listing of positions along a non-Calendar RPAS hierarchy. For example, along the product hierarchy, a position list could be a single-level listing of SKUs or it could be a multi-level list of classes, styles, and SKUs. A PQD is used to represent a set of selected

positions in a particular hierarchy. Each PQD is identified by a unique name. A user can load the PQD instead of performing manual selections on a two-tree wizard.

This utility can be used to load PQDs in master, local, and non-partitioned domains. PQDs are not shared across local domains. Loading a PQD into a global domain does not affect any local domain. Similarly, loading a PQD into a local domain does not impact the master or other local domains.

The input file must be in the following XML file format. This example shows loading and deleting lists for world and user access levels.

```
<? xml version="1.0" encoding="UTF-8" ?>
<pqdlists>
  <pqdlist name="list_name" hier="hierarchy_name">
    <access level="user">
      <comma_separated_user_names>
    </access>
    <dimension name="dimension_name">
      <pos>
        <position_external_name>
      </pos>
    </dimension>
  </pqdlist>
  <pqdlist name="list_name" hier="hierarchy_name">
    <access level="world">
    </access>
    <dimension name="dimension_name">
      <pos>
        <position_external_name>
      </pos>
    </dimension>
  </pqdlist>
  <pqdlist name="list_name" hier="hierarchy_name" operation="delete">
    <access level="world">
    </access>
  </pqdlist>
  <pqdlist name="list_name" hier="hierarchy_name" operation="delete">
    <access level="user">
      <comma_separated_list_of_users>
    </access>
  </pqdlist>
</pqdlists>
```

pqdMgr Usage

```
pqdMgr -d domainPath -load xmlFile
pqdMgr -d domainPath -delete xmlFile
pqdMgr -d domainPath -deleteAll
pqdMgr -d domainPath -validate xmlFile
pqdMgr -d domainPath -export outFile [-user userName|-world]
```

Table 7–29 provides descriptions of the arguments used by the pqdMgr utility.

Table 7–29 Arguments Used by the pqdMgr Utility

Argument	Description
-d domainPath	Specifies the path to the domain.
-load xmlFile	Loads position lists from an input XML file. Position lists with an operation attribute of delete are ignored.

Table 7–29 (Cont.) Arguments Used by the pqdMgr Utility

Argument	Description
-delete <i>xmlFile</i>	Deletes PQDs as specified in the input XML file with an operation attribute of delete. Position lists with an operation attribute of load are ignored.
-deleteAll	Deletes all PQD lists from the domain.
-validate <i>xmlFile</i>	Validates the XML file and report any errors. No impact on the existing PQD files in the domain.
-export <i>outFile</i> [-user <i>userName</i> -world]	Exports existing PQDs in the domain for a user or world access level. The file specified by <i>outFile</i> is overwritten. Requires one of the following options: ▪ -user <i>userName</i> exports PQDs for the provided <i>userName</i> in the same XML format as used for a load. ▪ -world exports PQDs with world permission in the same XML format as used for a load.

pqdMgr Notes

Note the following:

- Input files are validated before loading. All dimensions, hierarchies, and user names provided in the input file must be consistent with the existing hierarchies and registered users in the domain. The utility fails (return a non-zero error code) if it finds such inconsistencies in the input file. The errors are reported to the standard output.
- Multiple list operations are allowed in the XML input file.
- The supported operations are load and delete. If no operation is specified for a list, the default is load.
- The name *list_name* must be unique within an access level and, if the access level is user, for the user name.
- Each list definition consists of the list name, hierarchy, and access level. One or more dimension definitions are allowed. One or more position definitions are allowed for each dimension. Only external names are allowed to describe positions.
- When specifying an access level of user, a single user name or a comma-separated list of user names is required. A PQD file is created for each user name in the list.
- For the access level of world, the PQD file that is created is saved in the following path:
`<domain_root>/wizardPQD/<hierarchy_name>/_world/<list_name>`
- For the access level of user, the PQD file that is created is saved in the following path:
`<domain_root>/wizardPQD/<hierarchy_name>/<user_name>/<list_name>`

Data Management

This chapter explains the processes involved in RPAS data management. It includes the following sections:

- [Loading Measure Data Using loadmeasure](#)
- [Exporting Measure Data Using exportMeasure](#)
- [Exporting Measure Data Using exportData](#)
- [Mapping Data Between Domains Using mapData](#)
- [Moving Data Between Arrays Using updateArray](#)
- [Scan Domain Data Using scanDomain](#)
- [Repair Domain Metadata Using fixDomain](#)
- [Sync Measure NA Value with Array NA Value Using syncNAValue](#)

Loading Measure Data Using loadmeasure

The loadmeasure utility is used to load measure data from text files into the domain. The administrator must specify the measure names and the path to the domain that contains the measures.


The loadmeasure utility supports the use of fixed width and CSV (comma-separated variable) files for loading measure data. RPAS recommends the use of CSV files to reduce the size of the load file and to reduce disk I/O time.

To load measure data, system administrators must copy or create one or more load files in the input folder of the domain directory. The administrator can then call loadmeasure to load data.

Example:

Measure A has a base intersection of SKU and a clearint of dept.

If there is data for only a few SKUs (3, 5, and 11) in the incoming file, and SKUs 3 and 5 roll up to dept1 while sku11 rolls up to dept3, the data in all of the SKUs that rolls into dept1 and dept3 will be cleared.

Figure 8–1 Loading Measure Data

- The base intersection is SKU.
- The clearint is dept.
- Data is present for SKUs 3, 5, and 11; these fall under dept 1 and dept3.
- Data are cleared for dept1 (SKUs 1, 2, 3, 4, 5, 6) and dept3 (SKUs 11, and 12).
- Data for SKUs in dept2 (7, 8, 9, and 10) are untouched.

Load File Names and Load Behavior

System integrators must pay close attention to file naming. If a file name has not specifically been configured in the domain configuration, the file must be named the same as the measure name, with the appropriate extension depending on the type of the load.

For example, if the measure is named "rsal" and does not have a filename configured in the domain configuration, then the basic filename will also be "rsal". This name should be appended with one of the following extensions to indicate the type of load. If the load is an overlay, then the filename should be rsal.ovr; if it is an increment the file name should be rsal.inc, and so on. If a CSV file is being used, then the load type extension should be prefixed with the .csv extension; for example, rsal.csv.ovr and rsal.csv.inc.

RPAS supports the following types of loads (identified by file name extension):

- .ovr (Overlay): Existing values in the measure are overlaid with the values in the input file. Any values not included in the input file are not changed in the measure.

Note: For string type measures, an empty cell in the ovr file is treated as a valid string; as a result, the loadmeasure utility overwrites the previously loaded string with an empty string.

For other measure types, an empty cell in the ovr file is treated as invalid data. It is discarded and the previously loaded value is retained.

- .rpl (Replace): The existing measure is cleared and the values in the input file are taken as the new values for the measure. Existing values for cells that do not exist in the load file are switched to NA. In other words, all data at the base intersection for the measure are removed before cells are populated with the data from the incoming file.
- .inc (Increment): Increment mode should only be used with numeric measures in which the load file contains incremental values. Therefore, if a cell had a value of 2 and the .inc file provided a value of 3 for the cell, the new value for the cell will be 5 (2 incremented with 3).

- **.clr (Clear):** Clear mode is a variation of replace mode. It is used when measure data is loaded in parts or staggered in time, so that data for all positions grouped by an aggregate level position is replaced if one or more positions for that group of positions are being loaded.

In other words, data at the base intersection of a measure is partially cleared based on incoming data and the clearint attribute for the measure. The clearint attribute defines an intersection above the base intersection. All cells at the base intersection that are descended from a given position at the clearint level will be removed if data exists in the incoming file for at least one of those descending positions.

For example, assume that there are four regions, each with several stores, and the data is loaded region by region or for a subset of regions at a time. When loading data, ensure that data for a region is completely replaced with the new load if the load file has data for one or more stores from that region; however, other regions should be left untouched. This is made possible by clear loads where the clear intersection (clearint) property of a measure specifies the aggregate level at which to group positions for completely replacing the data. In this example, the clear intersection is at the region level. Clear intersection does not have to be performed along one hierarchy, but can be performed at the intersection of multiple hierarchies.

However, if you load multiple .clr files with region as the clear intersection, and data for one of those regions is in multiple files, then the last loaded .clr file for that region will replace any information that the previous .clr files loaded for that particular region.

The loadmeasure utility allows more than one load file to be present in the input folder at the same time for the same measure. If more than one load file is present in the input folder at the same time, each will be loaded. Since RPAS has a strict naming convention for measure file names, in order to add more than one load file at the same time, integrators must append the filenames as described above with file-distinguishing extensions.

For example, with the file names rsal.csv.ovr.1 and rsal.csv.ovr.2, RPAS does not care about the form of the multi-file extension. The extensions can be anything, number or text, and RPAS will still load them.

Note:

- Backup files should not be named as rsal.csv.ovr.bak or they will be loaded as well.
 - loadmeasure does not guarantee any specific ordering of loads based on the appended extensions.
-

The loadmeasure utility also allows multiple types of load files to be present in the input directory at the same time. RPAS loads .rpl files first, then .clr, .ovr, and .inc files. Since .rpl files completely erase existing measure data and then load the given data, you should not have multiple .rpl files at the same time.

Loading Multiple Measures from One File

The loadmeasure utility allows multiple measures to be loaded from a single file. You can load measures from ["CSV Files"](#) or ["Fixed Width Files"](#).

Note: See the “Data Interface Tool” section of the *RPAS Configuration Tools User Guide* for more information.

CSV Files

If a CSV file is used for loading measure data, loadmeasure will use the order that measures were specified on the command line to determine the order of columns in the CSV format. For example, if a file named multiple is used to load measures A, B, and C, where the call to loadmeasure listed the -measure argument as A,C,B, then when using the CSV file multiple.csv.ovr, loadmeasure will assume that after the dimension columns, the first column is A, then C, and then B, because that is the order they were passed in the call to loadmeasure.

It is not necessary to load all measures in a multiple measures file. If a file contains more columns than the call to loadmeasure requires, the trailing columns are ignored. If a line in the file does not contain enough columns to hold values for all measures specified, the line will be skipped. For example, if a file containing three measures (after the dimension positions), but only one measure is specified in the command line, only the first measure field will be used and the rest of the line is ignored.

Note: Even though it is not required to specify all measures contained in the multi-measure CSV file in a single loadmeasure command, there is no way to skip data columns in the CSV file.

Fixed Width Files

With a fixed width file, a single measure's data can be loaded from a file containing multiple measures.

Loading Data from Below the Base Intersection of the Measures

The loadmeasure utility supports loading measure data from an intersection lower than the base intersection of the measure. The load intersection has to be pre-specified in the configuration (loadint property) and the load time aggregation (loadagg property) method must also be specified. See the *RPAS Configuration Tools User Guide* for information on setting up measure properties.

When loadmeasure loads data from below the base intersection, all low-level data corresponding to a cell at the base intersection must be available in the load file for RPAS to be able to correctly aggregate the low-level data to the base level. A mistake in the values of a subset of cells that aggregate up to one cell at the base level can only be corrected by reloading the data for all low-level cells that correspond to the cell at the base level. If any low-level cells are missing, RPAS replaces their value with NA.

To perform a lower level load, RPAS first aggregates the data and then applies the appropriate load type to update the measure value, overwriting the existing value with the aggregate of the input cells if .ovr files were used, or incrementing the existing value with the aggregate of the input cells if .inc files were used.

Staging Measure Loads

RPAS supports the notion of stage-only measures. For stage-only measures, loadmeasure queues the loaded data in an intermediate staging area, but does not load it into the measure until it is called with the -applyloads parameter. For stage-only measures, loadmeasure should be called twice, once to stage the measures and then

with the `-applyloads` parameter to subsequently load the staged data in the measure arrays. The loadmeasure utility cannot simultaneously stage loads and apply the staged loads.

Measure staging should be performed when measure data can arrive from different sources, in different load formats, and staggered in time, when system administrators want to queue all these loads up and apply them at once while honoring the data arrival queue. Measure staging can be performed while the system is online as it does not cause measure data-related contention (it has the potential to cause metadata-related contention). When staging measure data, loadmeasure splits the data and purges the data files if data purging is enabled; it does not purge measure data until the loads are applied. This staging time preprocessing significantly reduces the load time when the loads are actually applied.

Note: The replace (.rpl) format cannot be used for staging. Furthermore, data loads from below the base intersection of the measure cannot be staged.

Running Pre-Load or Post-Load Scripts

The loadmeasure utility provides the ability to automatically run scripts before and after the utility is executed. These are referred to as preprocessing and post-processing scripts.

When loadmeasure is called, the utility checks for the existence of scripts named `pre<measurename>.sh` in the `./scripts` directory of the domain. If scripts exist, they will be run prior to the execution of the utility. Similarly, after the utility has completed running, the utility checks for the existence of scripts named `post<measurename>.sh` and executes them.

When multiple measures are loaded in a single call, only the preprocessing script for the first listed measure has any effect on the data.

Purging Old Measure Data

System administrators can purge old measure data during a load. When the base intersection of a measure involves the Calendar hierarchy, the setting for the purgeage measure property defines how and when existing data gets purged to an NA value. If `purgeAge` has not been set, the data never gets purged. If a purge age of zero or more has been set, data is purged for all dates before `RPAS_TODAY` - purgeage days. That is, if `purgeage` is 5, then at data load time, all data that is older than 5 days before `RPAS_TODAY` will be purged.

Behavior in an Integrated Environment

The loadmeasure utility cannot be used to load measure data for a measure that is shared through an RPAS Data Mart. If one of the measures passed to loadmeasure is a shared measure, loadmeasure will exit with an error message identifying the shared measure.

See the section on loadFactData for information on loading data into shared measures.

Behavior in a Global Domain Environment

In a global domain environment, loadmeasure is centralized and can only be called in the master domain. The loadmeasure utility loads one or more input files that can contain data from one or all of the local domains within the given global domain

environment. The utility then splits the input files and loads them into the required domain (which is the local domain to which the position belongs), or the master domain if the measure has a base intersection above the partition level. The split only occurs once in the case of multiple measures. Local domains will be checked for files even if there is no file in the global domain. The utility can be run in parallel in a global domain environment.

loadmeasure Usage

```
loadmeasure -d pathToDomain -measure measureName{,measureName,...}
{-applyloads}{-processes max} {-noClean} {-forcePurge}{-splitOnly | -noSplit}
{-defrag} {-loglevel level} {-recordLogLevel level} {-inDir inputDirectory}
```

Table 8–1 provides descriptions of the arguments used by the loadmeasure utility.

Table 8–1 Arguments Used by the loadmeasure Utility

Argument	Description
-d <i>pathToDomain</i>	Specifies the domain in which to load the measure.
-measure <i>measureNames</i>	Specifies the name of the measures to load. Measure names must be lowercase (for example, measurename1, measurename2, measurename3). If more than one measure is specified, all the measures must be in the same input file.
-applyloads	Applies any staged loads for the named measure. If the measure is registered to be a stage-only measure, loadmeasure will put the load in a staging area but will not update the measure until loadmeasure is called again with this argument. Upon the use of this argument, loadmeasure applies all loads that have been queued up in the staging area. It clears out the staged loads unless the measure's loadsToKeep property has been set to a non-zero number. In that case, it does not clear out the latest loadsToKeep loads. Note that only .ovr, .inc, and .clr loads can be staged. .rpl loads cannot be staged. Additionally, staging is only allowed for base intersection loads. RPAS cannot stage loads where load intersection is below the base intersection of the measure. This argument must not be used for measures that are not stage-only.
-processes <i>max</i>	Specifies the maximum number of child processes for parallel splitting of files and loading of measures across local domains in a global domain environment. For instance, if you specify five as the maximum number of processes, then up to five child processes can run concurrently in the split or load operations. If this argument is omitted or if only one process is specified, the application will perform all processing in a single process and no child processes will be created. This only specifies the number of child processes. The controlling process is not included (<i>max</i> + 1 is the actual number of processes).

Table 8–1 (Cont.) Arguments Used by the loadmeasure Utility

Argument	Description
-noClean	Prevents the input files from being moved to the processed directory. This option is used when a single file is used to load multiple measures, but not all measures from the file are loaded at once. The use of this option instructs loadmeasure to leave the load file behind for subsequent loading of unloaded measures. The user might want to use this option to perform intermediate processing between loads of measures available from the same file.
-forcePurge	Forces the purge routine to run even if no new data is loaded. This purges old measure data. This option can be applied to stage-only measures without having to apply loads. When a measure has the Calendar hierarchy in its base intersection, the setting for the purgeAge measure property defines how and when existing data gets purged to a NA value. If purgeAge has not been set, the data never gets purged. If a purge age of zero or more has been set, data is purged for all dates that are before RPAS_TODAY -purgeAge days. That is, if purgeAge is five, at data load time all data that is more than five days before RPAS_TODAY will be purged. This option does not require you to load any new data.
-loglevel <i>level</i>	Sets the logger verbosity level. Possible values: all, profile, audit, information, warning, error, or none.
-noSplit	Loads the pre-split input files (created by -splitOnly) into the local domains. This option should only be used in global domain environments.
-splitonly	Causes the input files in the global domain to be split across the local domains, but does not do any further processing of the input files. Subsequently, loadmeasure can be used with the -noSplit argument to load these pre-split input files into the local domains. File-splitting is a fairly time consuming activity and can consume up to 80 percent of the load time. System integrators may be able to improve batch performance by breaking away file-splitting from actual measure loading. This is useful if a multi-measure file is being used in such a way that subsets of measures are loaded at different steps in a batch process. The file can be split with multiple processes by specifying the -processes argument. This option should only be used in global domain environments.
-defrag	Defragments the domain at the end of the measure loading process to reduce the physical size of the domain. This space-saving is achieved by replacing the existing fragmented pages with copied, fully populated BTree database pages.

Table 8–1 (Cont.) Arguments Used by the loadmeasure Utility

Argument	Description
-recordLogLevel <i>level</i>	Sets a logging level for record loading issues. Issues such as parsing errors, missing positions, and data conversion errors are evaluated for every record in the measure load file. By default, these are logged as errors in the log file of the loadmeasure utility. However, customers might want to downgrade the logging level for such record loading issues. They can do that by using the -recordLogLevel <i>level</i> argument. The standard log levels, error, warning, information, and profile, can be used as parameters to this argument. When logging, loadmeasure compares this logging level to the utility's logging level (set using -loglevel). If the utility's logging level is less verbose than the record logging level, then record issues will not be logged. If utility's logging level is at same or higher verbosity as the record logging level, the record issues will be logged with the log indicator as set using this argument.
-inDir <i>InputDirectory</i>	Only .rpl files can be used with this option, and only the CSV format with header line is supported. The header line is used to map the columns to dimensions and measures (for example: SKU,STR,DAY,Sales). Enter one measure per input file. The name of the measure is extracted from the file name; for example, sales.csv.rpl corresponds to measure sales. The input data must be at the base intersection of the measure. If the measure is normally partitioned (non-HBI), a sub-domain index may be used for further performance optimization by avoiding the data-splitting step. For measures that may contain duplicate positions in different sub-domains (FnHBI measures), the sub-domain index is required. In either case, the name of the file is used to figure the sub-domain index (for example, sales.0.csv.rpl corresponds to the first local domain; sales.1.csv.rpl corresponds to the second local domain, and so on). Note: The sub-domain index is designed to be used in conjunction with exportMeasure -hier only. Manual name-indexing of the files is not recommended. The filename property of the measure is not considered with the -inDir option.

Loading Image Paths for Positions

A configuration and backend process may also be used to support the load of image paths for one or more positions of a dimension at a time. The paths of the images must be stored in a measure called `r_images_<dimension name>` where `<dimension name>` must be replaced with the RPAS Name of the image-enabled dimension (for example, `r_images_sku` if loading image paths for the sku dimension). This measure is single-dimensional, defined on the image enabled dimension. An .ovr file is required with position names and the image paths for those positions formatted according to the RPAS measure load formats. The loadMeasure utility is then used to load this data into the domain.

Note: See the *RPAS Configuration Tools User Guide* and the "Position Images" section in the *RPAS User Guide for the Classic Client* for more information on Image Display.

Example

```
loadmeasure -d <domain path> -measure r_images_sku
```

where <domain path> is the path to the domain.

Exporting Measure Data Using exportMeasure

The exportMeasure utility is a command-line utility that may be used to export domain or workbook measure data from RPAS in either a CSV or a fixed-width file format. A single measure, or multiple measures, may be exported based a specified intersection. If the measure's base intersection is not the same as the export intersection, the measure's default aggregation method will be used to aggregate data to an intersection higher than base, or replication will be used for spreading measure data if the data is required at an intersection lower than base. This utility:

- Supports export of data in a user-specified range, which can be a single mask measure, a range specified on Calendar dimension, or a combination of the two.
- Supports multiple processes for better performance in a global domain environment.

Note: The exportMeasure utility allows multiple measures to be exported into the same file when multiple measure names are provided. The same measure name cannot be specified more than once (using comma separation) in a single call; therefore, a measure can be exported only once per file.

The exportMeasure utility cannot be used to export the data contained within a shared measure in a domain integrated with an RPAS Data Mart. If one or more of the measures passed to exportMeasure is shared through an RPAS Data Mart, exportMeasure will exit with an error message identifying the shared measure.

exportMeasure Usage

```
exportMeasure -d pathToDomain -out outFile [COMMAND] [OPTIONS]
```

Table 8–2 provides descriptions of the arguments used by the exportMeasure utility.

Table 8–2 Arguments Used by the exportMeasure Utility

Argument	Description
-d <i>pathToDomain</i>	Specifies the path to the domain.
-out <i>outFile</i>	Specifies the output file name. It is required and must be a valid file name including the path.
-wb <i>wbname</i>	If specified, exportMeasure exports data from the specified workbook (wbname). A valid workbook name must be used.

Table 8–2 (Cont.) Arguments Used by the exportMeasure Utility

Argument	Description
-intx <i>intxString</i>	Specifies the intersection at which to export measures. If the measure's base intersection is higher than the export intersection, replication is used to spread the measure down to the export intersection. If the measure's base intersection is lower than the export intersection, the measure's default method (defagg) is used for aggregation. The export intersection must be either at, above, or below the base intersection of the measure. The export intersection cannot have some dimensions above the dimension in the base intersection of the measure and some below. The RPAS dimension names in an intersection should be four characters in length. If a RPAS dimension name is less than four character long, then an underscore character (" _ ") must be used as a filler at the end of a dimension name.
-mask <i>measureName</i>	Specifies a mask measure, which must be a valid Boolean measure registered. In the current measure store, its baseintx must be at same export intx.
-range <i>start:end</i>	Specifies a range of positions along the innermost dimension. Only values in the range are considered for export.
-processes <i>max</i>	Defines the maximum number of processes to run in parallel.
-append	Appends new output to the current output file. If not specified, the current output file will be erased and replaced with new data.
-nomerge	If run in a global domain environment and exporting intersection below partition dimension, and have processes set greater than 1, specifying nomerge will stop exportMeasure from merging multiple output files created from each local domain to the master output file. Output files created from local domain are stored at masterdomain/output/exportMeasure[TS] folder, where [TS] represents a timestamp. Files are named as out000X.txt, where 000X is the index of the local domain.
-compress	Specifies that the output file should be in the compressed CSV format.
-hier <i>hierarchy1, hierarchy2</i>	Exports all measures for hierarchies. It exports only measures that have storage in the domain. Multiple hierarchies can be specified in a comma-separated list.
-outDir <i>outputDirectory</i>	Updates the output directory. If the output directory does not exist, the utility creates one. The measure names are used to generate the output file names. A CSV file with a header line can identify the dimensions of the base intersection and the name of the measure that is generated for each file. The files always have a csv.rpl extension (e.g. sales.csv.rpl). Old files are overwritten. One output directory is created for each HBI measure. In addition, one file is created for each non-HBI or FnHBI measure per sub-domain. The file names contain an internal sub-domain index, for example sales.0.csv.rpl, sales.1.csv.rpl, and so on.
-upperCase	Converts the position names to all uppercase before writing the output data file. Without this argument, position names are in lowercase since they are stored in lowercase in the domain.

Table 8–2 (Cont.) Arguments Used by the exportMeasure Utility

Argument	Description
-meas " <i>measSpec</i> , <i>measSpec</i> ..."	Must specify one. <i>measSpec</i> is <i>measName.modifier</i>. The -meas argument may be repeated to export multiple measure arrays to the same output file. modifier include the following: .precision<double>, specifies the precision for numeric measure .format<formatString>, specifies the user defined export format The examples below provide valid measure specifications given <i>MeasNameA</i> is a valid real type measure. Examples: <pre>-meas MeasNameA -meas MeasNameA.precision(0.0001) -meas MeasNameA.format("%13.2f").precision(0.01) -meas MeasNameA.precision(0.01).format("%13.2f")</pre> For specifying date and time, the following formats are supported: %Y: four-digit year %y: two-digit year %m: month %d: day %B: full name of the month %b: three character abbreviation for the month %H: hour %M: minute %S: second %s: milli-second The examples below provide valid measure specifications given <i>MeasNameB</i> is a valid date/time type measure. Examples: <pre>-meas MeasNameB -meas MeasNameB.format("%Y%m%d") -meas MeasNameB.format("%d%B%Y%H%M%S")</pre>

Exporting Measure Data Using exportData

Use exportData to export measure data from RPAS into text files. Each line that is exported contains the position name for the exported dimension followed by the value in the cell for each array being exported.

Note: More than one array may be exported and more than one dimension in each array can be exported.

The utility may be invoked by specifying all parameters on the command line or by specifying an array that contains a list of the parameters.

When running this utility in a global domain environment, the utility must only be called to export data from the master domain. The utility extracts the data from either the local domains or the master domain depending on where the data resides, which in turn depends on the level at which the global domain environment is partitioned.

The exportData utility cannot be used to export the data contained within a shared measure in a domain integrated with an RPAS Data Mart. If one or more of the measures passed to exportData is shared through an RPAS Data Mart, exportData will exit with an error message identifying the shared measure.

The parameters specify what arrays and dimensions are exported and how to format the data. It is best to specify the arrays first. An array specification begins with -array followed by the array information. This includes the array name, formatting string, NA cell value, and NA cell value formatting string. The formatting string for both the cell value and NA value is based on the C language printf function formats. See the documentation on the printf for more information on the possible values. The -array parameter can be repeated as needed to export more than one array into the same export file. Remember that the order in which the arrays appear in the -array parameter is the order that they appear in the export file.

After the arrays have been specified, the administration must specify the dimensions to be exported within the arrays. The -dim parameter is used to specify a dimension in an array. The -dim parameter is followed by the dimension name, a convert option, the formatting string (just like an array), and the order the dimension appears in the export file. Because arrays are not required to contain identical dimensions, it is important to list all dimensions in all arrays with the -dim parameter. This makes it possible to track dimensions across arrays and line the data up correctly. If a dimension in an array is not to be in the export file, set the last value of this parameter to 0. The conversion option specifies either the number of characters to be removed from the position name or it specifies an array that contains the real position name. If an array name is given, this array must be a vector. The function will go to this array and use the original position name to jump to the cell of the same position name. It will then get the cell value and use that as the position name in the export.

It is possible to specify the number of decimal places when exporting numeric measures of data type real. This setting is defined in the specifications for measures, arrays, and dimensions (measSpec, arraySpec, and dimSpec). The format is %[precision]type where [.precision] is the number of decimal places and type is the letter f. For example, the setting %.2f exports numbers with two decimal places. Other settings are provided below.

If all parameters are contained in an array, after the export file name and source database name, the -params parameter is used to specify the database name and array name that contains all of the parameters needed for the export.

Note: Either the -array, -meas, or -params parameters must be specified when using this utility.

exportData Usage

```
exportData -d domainPath -out outputFile -params db array
exportData -d domainPath -out outputFile -array "arraySpec" {options}
exportData -d domainPath -out outputFile -meas <measspec> | -array <arrayspec> |
-params <paramspec> -wb <wbName> {options}
```

Table 8–3 provides descriptions of the arguments used by the exportData utility.

Table 8–3 Arguments Used by the exportData Utility

Argument	Description
-d domainPath	Specifies the domain that contains the data that to export.

Table 8–3 (Cont.) Arguments Used by the exportData Utility

Argument	Description
-out <i>outputFile</i>	Specifies the file that will contain the exported data. The <i>outputFile</i> is relative to the domain unless the full path is specified.
-meas \" <i>measSpec</i> \"	Specifies the measures to export. <i>measSpec</i> must be quoted, and the format is \" <i>measName</i> <i>cellFormat</i> <i>naValue</i> <i>naFormat</i> \" The -meas argument can be repeated to export multiple measure arrays to the same output file. Measures are exported at the base intersection.
-array \" <i>arraySpec</i> \"	Specifies the array to export. <i>arraySpec</i> must be quoted, and the format is \" <i>dbName</i> <i>arrayName</i> <i>cellFormat</i> <i>naValue</i> <i>naFormat</i> \" ■ <i>dbName</i> can be a path to the database (relative paths are relative to the domain root). ■ Both <i>cellFormat</i> and <i>naFormat</i> use printf format commands. See the printf function for more information on the possible values. The -array argument can be repeated to export multiple arrays to the same output file. The order in which arrays are listed is the order in which they will be exported. Note: This argument cannot be used in a global domain environment and can only be used in simple domains. This argument cannot be used with -useLoadFormat.
-params <i>db array</i>	Instead of specifying all parameters on the command line, this parameter allows the parameters to be read from an array. ■ <i>db</i> specifies the name of a .ary file where the array of parameters is stored. ■ <i>array</i> specifies the name of an array in the specified database that has the above parameters.
-wb <i>WbName</i>	Used after specifying any of the command arguments: -array, -meas, or -params.
-append	Specifies that output is appended at end of output file. The default is to overwrite output file.

Table 8–3 (Cont.) Arguments Used by the exportData Utility

Argument	Description
-dim <i>"dimSpec"</i>	Specifies the dimension to be exported. ■ <i>dimSpec</i> must be quoted, and the format is <i>"dimName conversion format order"</i> ■ <i>conversion</i> is either a count of the number of characters to strip from the start of the position name or the name of an array to be used to translate the position name before writing to the output file. ■ <i>format</i> is a printf-style format for the position names. See the printf function for more information on the possible values. ■ <i>order</i> indicates the order the dimension is listed in the output file. If the value is 0, then the dimension is not exported. The -dim parameter can be repeated. The -dim parameter is not allowed with the -useLoadFormat. When using with the -wide parameter, the -dim parameter should not be used for the innermost dimension.
-skipNA <i>always allna anyna arrayna</i>	Controls whether a line of data is exported based on having NAs in a cell. ■ <i>always</i> exports data regardless of whether or not it contains NAs. ■ <i>allna</i> does not export a row of data if all columns are NA (default). ■ <i>anyna</i> does not export a row of data if any cell contains a NA value. ■ <i>arrayna</i> does not export a row of data if the value in the given array name is NA (requires -naArray).
-naArray <i>arrayName</i>	When <i>arrayna</i> is specified using the -skipNA parameter, this option specifies the export array that is checked to determine if data is exported.
-wide	Causes the data to be exported wide, which means the innermost dimension will go across the row instead of each cell on a separate line. This is most useful when the innermost dimension is time. The -range parameter can be used in conjunction with wide format (-wide) to specify a range along the innermost dimension. The -dim parameter should not be used for the innermost dimension when -wide is being used.
-range <i>start:end</i>	Used to limit the export to positions in the range. The range can only be specified for the innermost dimension. May be used in conjunction with the -wide parameter.
-time	Specifies the YYYYMMDD format for dates.

Table 8–3 (Cont.) Arguments Used by the exportData Utility

Argument	Description
-precision <i>precisionValue</i>	Causes the utility to avoid exporting values that differ from the NA value by the specified value. Any values smaller than the precision value are not exported. For example, consider a measure with the NA value of zero and a precision value of 0.01. A value of 0.0034 would not be exported while a value of 0.34 would be exported. The precision value must be less than one. If a value greater than one is provided the utility returns a warning.
-processes <i>max</i>	Defines the maximum number of processes to run in parallel.
-useArrayNaValue	Enables the use of the NA value of the array instead of the NA value specified in measSpec or arraySpec.
-upperCase	Converts the position names to all uppercase before writing the output data file. Without this argument, position names are in lowercase since they are stored in lowercase in the domain.
-displayArrayNaValue	Controls the display of measures used as a mask in -naArray option. The default is to not display the mask NA measure value. However, if this option is specified, then the NA Array measure values are also exported.
-useLoadFormat	Enables the use of the format as specified by the measure property. The level at which the data is stored in the domain is used. The -dim parameter is not allowed with the -useLoadFormat.

The -useLoadFormat Parameter

Use the format specified by the measure's loading format to export the measure. This loading format includes Start and Width, which defines the column that corresponds to this measure's data in the measure load file. The measure is exported into the same column in the output file. If the full measure export specifications are not provided, including the cellFormat, naValue and naFormat, the default format will be used. The default export formats for each type of measure are as follows:

- Integer: %<width>.0f
- Real: %<width>f
- String: %<width>s
- Date: %Y%m%d
- Boolean: TRUE or FALSE as string

All values are exported right aligned, as in the measure loading file.

If users provide full measure specifications, then user-specified cellFormat, naValue, and naFormat will be used rather than the default format.

Users can either use the default format by specifying the measure name only or give the full specifications. Partial measure specifications are not permitted.

If users specify multiple measures to be exported into the same file, each of these measures will occupy a column in the file defined by its start and width attributes. If two measures occupy the same column, exportData will throw an exception with an error message saying "overlapping measures in the output file" and exit. If a measure's column is overlapping with the columns occupied by the position names, exportData will throw an exception with an error message saying "measure column is overlapping

with position columns" then exit. Basically, if the measure cannot be exported correctly, exportData will not try to export it but simply exit and alert user with a proper exception.

The -dim and -array parameters are not allowed if -useLoadFormat is used. All dimensions in the measure's base intersection are exported by default. The external position name is exported to the export file, in the order specified by the hierarchy's order attribute, usually in the order of CLND, PROD, and LOC. The position names are left aligned in the export file.

Mapping Data Between Domains Using mapData

The mapData utility is used to move data from one domain to another. Specifically, it copies data from an existing domain, database, or array to a new domain, database or array.

Before this utility is run, the new hierarchy must be loaded in the destination domain. After mapData has copied data, administrators can purge the source domain by calling loadHier with a purge age of 0. Tasks such as hierarchy loading, hierarchy purging, and the validation of source and destination domains are performed outside of this utility.

Note: This utility does not update buffer positions.

mapdata Usage

```
mapdata -d SrcPath -dest destPath [-db dbName [-array arrayName]]
{-db dbName {-array arrayName}} {-loglevel}
```

Table 8–4 provides descriptions of the arguments used by the mapData utility.

Table 8–4 Arguments Used by the mapdata Utility

Argument	Description
-d <i>SrcPath</i>	Specifies the path to the source domain.
-dest <i>DestPath</i>	Specifies the path to the destination domain.
-db <i>dbName</i>	Applies mapdata only on the given database. Must be a valid file. If this argument is not specified, the entire domain will be included in the operation.
-array <i>arrayName</i>	Applies mapdata only on the given array. The database in which the array resides must be specified with the -db argument.

Moving Data Between Arrays Using updateArray

The updateArray utility moves data from a source array to a destination array. The destination array must contain the superset of dimensions in both source arrays. The source array's dimensions may be at the same or higher level, as mapped by the dimension dictionary. If a dimension in the source array is at a higher level, the results are spread across the lower level dimension in the destination. If there are extra dimensions in the destination array, the results are replicated across these extra dimensions. The NA value of the destination array remains unchanged.

To limit the scope of the update, a mask array and an innermost range may be specified. If a mask array is given, the update is limited to cells in the source array for which the corresponding mask cell is on. If an innermost range is given for source or

destination array, the update is limited to cells that are within the start and end of this range on the innermost dimension. If the source and destination arrays are not in the same domain, the measure store associated with the source domain is used to find hierarchy information.

Note: This utility does not update buffer positions.

updateArray Usage

```
updateArray -destArray dbPath.arrayName {-srcArray dbPath.arrayName} {-destDomain
domainPath {-srcDomain domainPath} {-maskDomain domainPath} {-maskArray
dbPath.arrayName} {-updateMethod method} {-srcRange first:last} {-destRange
first:last} {-srcScalar scalarCell} {-version} {-loglevel level}
updateArray -argFile filename {-version} {-loglevel level}
```

Table 8–5 provides descriptions of the arguments used by the updateArray utility.

Table 8–5 Arguments Used by the updateArray Utility

Argument	Description
-destArray <i>dbPath.arrayName</i>	Specifies the destination array where the data is copied. Required. <i>dbPath</i> is relative to destDomain.
-srcArray <i>dbPath.arrayName</i>	Optional argument. Default is no source array. Note: This parameter cannot be used with -srcScalar scalarCell.
-destDomain <i>domainPath</i>	Optional argument. Default is current working directory.
-srcDomain <i>domainPath</i>	Optional argument. Default is current working directory.
-maskDomain <i>domainPath</i>	Optional argument. Default is current working directory.
-updateMethod <i>method</i>	Optional argument. Default is OVERLAY. The following update methods are available: ■ SKIPNA - Omit NA cells in source. ■ SKIPPOP - Omit populated cells in source. OVERLAYNA - Update NA cells in destination. ■ OVERLAYPOP - Update populated cells in destination. ■ OVERLAY - Update all cells in destination with source.
-srcRange <i>first:last</i>	Optional argument. Default is no range. Defines range along innermost dimension of source array.

Table 8–5 (Cont.) Arguments Used by the updateArray Utility

Argument	Description
-destRange <i>first:last</i>	Optional argument. Default is no range. Defines range along innermost dimension of destination array. The position names of the innermost dimension are the range value. For example, if the range values is one week, the range should be specified as -srcRange WEEK200811011:WEEK200811022 -destRange WEEK200811011:WEEK200811022
-srcScalar "TYPE:VALUE"	Optional argument. Default is NA cell. Format for scalar cell is one of: ■ NUMERIC: numeric value ■ STRING: literal value ■ BOOL: Boolean value ■ NA Note: This parameter cannot be used with -srcArray dbPath.arrayName.

Scan Domain Data Using scanDomain

The scanDomain utility is a domain utility used for detecting data loss and repairing data corruption in an RPAS database.

Data loss occurs when an RPAS process is abnormally terminated. This can happen when an external mechanism, such as a power failure, causes a sudden termination of an RPAS process. Data loss can also occur due to unexpected program breakdown.

Data corruption can occur if an external program modifies the RPAS database files or an unforeseen defect occurs in the processes using the RPAS database (an extremely rare event).

The scanDomain utility can detect both corruption and data loss, but it can only fix corruption. This utility can operate on global, non-partitioned, and local domains. It supports parallelization when repairing databases in a domain.

While the utility is attempting to perform a repair of the databases, it can use the command line option (-backup) to enable backing up the original databases. While running in detection mode (-detectDataLoss or -detectCorruption option), the utility does not change any of the RPAS databases, and therefore, it does not create such backups.

In detection mode, the utility prints a list of databases with data loss or data corruption to the screen. The output can be directed to a file.

scanDomain Usage

```
scanDomain -version
scanDomain -d domainPath [-detectDataLoss ] [-detectCorruption] [-loglevel level]
[-noheader]
scanDomain -d domainPath -repairCorruption [-backup][--processes
maximumNumberOfProcesses] [-loglevel level] [-noheader]
scanDomain -d domainPath -listUnusedData outputFile [-processes
maximumNumberOfProcesses] [-loglevel level]
```



```
scanDomain [-?|-help|-usage]
```

If the user intends to detect both corruption and data loss, it is more efficient to run the utility once with both the `-detectDataLoss` and `-detectCorruption` options. The user can run two consecutive commands for detecting corruption and data loss, although this is less efficient.

When running `scanDomain` to detect unused data, the user sees a list of databases that may not be needed by the domain. This information includes:

- User directories for users who are not registered in the domain
- Measure databases whose corresponding measures have been removed from the domain
- Other databases not referenced by the domain (for example, measure load databases, backup databases, and temporary databases)

The data contained in the specified databases may not be needed by the domain. If the user can confirm that the data is actually unnecessary, then the user can remove those databases before proceeding with the upgrade.

Table 8–6 provides descriptions of the arguments used by the `scanDomain` utility.

Table 8–6 Arguments Used by the `scanDomain` Utility

Argument	Description
<code>-version</code>	Prints the version of the utility.
<code>-d domainPath</code>	Specifies the path to a global, non-partitioned, or local domain. Required.
<code>-detectDataLoss</code>	Checks for data loss in the specified domain.
<code>-detectCorruption</code>	Checks for database corruption in the specified domain.
<code>-repairCorruption</code>	Repairs the database corruption in the specified domain. Note: This argument cannot be used with <code>-detectDataLoss</code> or <code>-detectCorruption</code> .
<code>-backup</code>	Backs up database files before attempting to repair them. Optional. Note: This argument can only be used with <code>-repairCorruption</code> .
<code>-listUnusedData</code>	Lists potentially unused data in the domain. Information about this unused data is output to a file whose location is passed as the <code>outputFile</code> argument.
<code>-processes maximumNumberOfProcesses</code>	Specifies the maximum number of processes to be started to repair Btree database corruptions. Optional. Note: This argument can only be used with <code>-repairCorruption</code> .
<code>-loglevel level</code>	Sets the logger verbosity level. Possible values: all, profile, debug, audit, information, warning, error, or none. Optional.
<code>-noheader</code>	Disables the timestamp header. Optional.
<code>-? -help -usage</code>	Obtains the usage text. Optional.

Repair Domain Metadata Using `fixDomain`

Discrepancies in the metadata of a domain can cause serious problems during domain operations, and such discrepancies can be difficult to correct. The `fixDomain` utility analyzes a domain in order to detect problems in the metadata of a domain and can be used to attempt to fix any detected problems.

By default, fixDomain runs in report-only mode. Users should run fixDomain in this mode first in order to detect discrepancies in the metadata. Should any problems be discovered, fixDomain provides functionality to fix certain problems in the metadata of the domain.

Sync Measure NA Value with Array NA Value Using syncNAValue

The measure array present in the measure's database can have an NA value which is different from the NA value of the measure in the measure's property. The array NA value can change depending on the state of the array. The NA value of the array is calculated based on the current value of populated cells in the array and hence this is not constant. The syncNAValue utility is used to set a measure's array NA value to its measure NA value.

The synchNAValue cannot be used to modify the array NA value of a measure shared through an RPAS Data Mart. If synchNAValue is called with the optional -m argument and the specified measure is a shared measure, synchNAValue will exit with an error message stating that the measure is shared. If it is called without the optional -m argument, synchNAValue will skip any shared measures within the domain.

syncNAValue Usage

```
syncNAValue -d [domainpath] [options]
```

Table 8–7 provides descriptions of the arguments used by the syncNAValue utility.

Table 8–7 Arguments Used by the syncNAValue Utility

Argument	Description
-d <i>domainPath</i>	Specifies the path to the domain where the utility needs to be run. This is a required argument. SyncNAValue can also be run from sub-domains. Caution: If syncNaValue is run with the -d option and without any additional arguments, it changes the array NA value to the measure NA value for every measure registered in the domain.
-m <i>measName</i>	Optional argument to run syncNaValue on a single measure only instead of all measures in the domain.
-report	Optional argument to print out all measures whose array NA value is different from its measure NA value. This option does not update any measure arrays. If used along with the -m option, it prints the measure array information even if the array and measure NA values are the same.

Usage Examples

In the following examples, it is assumed that syncNAValue is run from the domain path:

```
syncNAValue -d . -report
syncNAValue -d .
syncNAValue -d . -m measA
syncNAValue -d . -m measA -report
```

Operational Utilities

This chapter describes the following operational utilities of RPAS:

- [Finding Alerts Using alertmgr](#)
- [Copying Domains Using copyDomain](#)
- [Moving a Domain Using moveDomain](#)
- [Setting Miscellaneous Domain Properties Using domainprop](#)
- [Using the mace Calculation Engine](#)
- [Managing the Workbook Batch Queue Using wbbatch](#)
- [Managing Workbooks Using wbmgr](#)
- [Registering Measures Using regmeasure](#)
- [Registering Token Measures Using regTokenMeasure](#)
- [Batch Plug-In Tasks: execPluginTask.sh](#)

Finding Alerts Using alertmgr

Alerts are an exception management tool for users. An alert is a measure that evaluates a business rule (returning a value of true or false). RPAS then notifies users of the true conditions and allows users to build workbooks to resolve the scenario that drove the alert.

Alert measures are first defined in the domain using the RPAS Configuration Tools. These measures are of type Boolean, which means they have a value of true or false. Next, rules (expressions) are registered in the domain for the alert measures to define the business rules used to evaluate the alert.

After the registration process is complete, the alert utility is run to find the alerts in the domain. After the alert finder has been run, the identified alerts can be viewed in the Alert Manager window in the RPAS Classic Client.

The following is a summary of the process for defining and finding an alert:

1. Create an alert measure. This must be a Boolean measure (values are true-false or yes-no) and must be defined in the RPAS Configuration Tools. Its aggregation state and base state must be read-only.
2. Create the alert (the expression) for which the alert should be evaluated using the Configuration Tools. Using alertmgr, register the alert with a category, categoryLabel, and the above expression. This flags the registered measure as an alert so that it is recognized when the alert finder is run.

3. Repeat steps 1 and 2 for any additional alerts to be registered in the domain.
4. Run the alert finder on the domain to evaluate the number of instances in which one or more alert expressions are true. This operation is completed using the RPAS utility alertmgr.

alertmgr Usage

```
alertmgr -d domainPath [COMMAND [parameters]]
alertmgr -d pathToDomain -findAlerts {-navigationThresholdhits} {-alerts "a1 a2 ..."
| -categories "cat1 cat2 ..."}

```

Note: This utility includes arguments that are not documented in this guide as it is recommended that those operations be configured using the Configuration Tools to ensure consistency between the configuration and the domain.

Table 9–1 provides descriptions of the arguments used by the alertmgr utility.

Table 9–1 Arguments Used by the alertmgr Utility

Argument	Description
-d <i>pathToDomain</i>	Specifies the directory in which to run the utility. All commands except -version require -d domainPath.
-findAlerts	Finds alerts in the specified domain. The utility will find all alerts in the domain if neither the -alerts or -categories arguments are specified. If the -alerts or -categories list are not specified, findAlerts is run on all alerts. findAlerts can be run from either Master or Local Domains.
-alerts <i>a1 a2...</i>	Evaluates specific alerts in the domain. <i>a1 a2 ...</i> must be valid names of alerts that are defined in the domain.
-categories <i>cat1 cat2 ...</i>	Evaluates all alerts in the domain that are associated with specific categories of alerts. <i>cat1 cat2 ...</i> must be valid names of alert categories that are defined in the domain.
-sumAlerts	Sums up the hit counts of alerts across local domains. It can be run based on a list of alerts or alert categories. If none are provided, then the respective hit count of each alert across all local domains is summed. -sumAlerts can be used only from Master Domain. Note: -findAlerts must be run first to generate hit counts of alerts.
-navigationThreshold <i>hits</i>	Defines the maximum number of alert hits for the Find Next/Previous Alert functionality to remain operational in a workbook. If over that threshold, the Find Alert functionality will only work up to that number. The default value is 5000.

Note: The alertmgr utility can be run on the local domains individually. The administrator may spawn several processes in parallel and, when needed, run alertmgr -sumAlerts again to aggregate the results to the global domain. If parallelization is desired, the administrator should create a script to spawn the parallel processes.

Copying Domains Using copyDomain

The copyDomain utility is used to copy a simple domain, all domains included in a global domain environment, or a subset of domains in a global domain environment. Domains are often copied before the domains are upgraded after receiving a patch to RPAS.

For a standard, simple domain (in other words, not a global domain environment), copyDomain copies the domain directory recursively from one location to another.

For a global domain environment, copyDomain copies the master domain to the specified destination and then copies each local domain into the corresponding subdirectories of the new location. As part of this particular replication process, the utility also updates all relevant data structures so that the domains are properly connected together.

Relative paths are supported with this utility and are used when creating the new copies of all the underlying data structures (arrays). Relative paths are based on the full pathname of the domain's root directory.

copyDomain Usage

```
copyDomain -xmlConfigFile filename {OPTIONS}
copyDomain -d pathToSrc {OPTIONS}
copyDomain -version
```

Table 9–2 provides descriptions of the arguments used by the copyDomain utility.

Table 9–2 Arguments Used by the copyDomain Utility

Argument	Description
-d <i>pathToSrcDomain</i>	Specifies the path of the domain to be copied. This argument and -dest should not be used with -xmlConfigFile.
-xmlConfigFile <i>pathToXmlConfigFile.xml</i>	Allows copyDomain to copy each subdomain into user-instructed specific locations. This argument should not be used with -d OR -dest. See " copyDomain: Format of the XML Configuration File " for the file format.
The following arguments are valid for -xmlConfigFile and -d:	
-force	Deletes the existing domain at the specified destination path before copying the source domain.
-clone <i>dimposlist</i>	Copies a subset of a domain environment. Copies only positions specified in a format as <i>dim1,pos1,...,posn:dim2,pos1,...,posn</i> where the sequence <i>dim1,pos1,...,posn</i> specifies the selected positions along <i>dim1</i> . Multiple dimensions may be specified, but only one dimension per each hierarchy is allowed.
-partitionPositions <i>positions</i>	Deprecated. Use -clone instead.
-copyWorkbooks <i>workbookList</i>	Copies only the specified workbooks to the destination location. <i>workbookList</i> is either a comma-separated list of the workbooks to copy, or the value none such that no workbooks are copied. If this argument is not specified all workbooks in the environment are copied.
-skipInput	Do not copy the input directory located in the source domain.
-skipConfig	Do not copy the configuration directory located in the source domains.
-skipEmptyDir	Do not copy the empty directory located in the source domain.

Table 9–2 (Cont.) Arguments Used by the copyDomain Utility

Argument	Description
-maxProcesses <i>count</i>	If this argument is specified, some parts of copyDomain run in parallel, utilizing up to the given number of processes.
-noSubDomains	Do not copy any local domains in the source domain.
The following arguments are valid only with -d:	
-dest <i>pathToDestDomain</i>	Specifies the path to where the domain is to be copied. The copied domain can also be renamed in this step by providing a name different than the source domain. This argument must be provided when using any other option (other than -xmlConfigFile or -relativizePaths) of the utility. If this argument is not provided, the domain is updated to have relative paths.
-export	Exports each database from the source domain into a format that can be used on a UNIX platform. This argument cannot be used when specifying an -xmlConfigFile.
-gzip	Compresses the copied domain into a gzip format. This argument cannot be used when specifying an -xmlConfigFile.
-dimDictOnly	Copies only the source domain structure, metadata, and hierarchy data. Running copyDomain with this option result in a non-functional domain. Therefore, this argument should be used for diagnostic purposes only.
-relativizePaths	Updates the existing master and subdomain path references to relative paths. If the current absolute path references are invalid paths, subdomains are searched for in the same location as the master and within the master domain directory. When this argument is used, no domain copy is made. Note: When using this argument, do not provide a destination with the -dest argument. For example, if you build a domain, this is what it looks like at first: <pre>(PGRP100 INFO): "C:\Oracle\Domains\1323\mfprtl\ldom3" (PGRP200 INFO): "C:\Oracle\Domains\1323\mfprtl\ldom3" (PGRP21 INFO) : "C:\Oracle\Domains\1323\mfprtl\ldom0" (PGRP22 INFO) : "C:\Oracle\Domains\1323\mfprtl\ldom1" (PGRP300 INFO): "C:\Oracle\Domains\1323\mfprtl\ldom3" (PGRP31 INFO) : "C:\Oracle\Domains\1323\mfprtl\ldom2" (PGRP32 INFO) : "C:\Oracle\Domains\1323\mfprtl\ldom2"</pre> This tells you, by partition dimension, the path to the local domain. You can see that it is a full path and that the domains are subdirectories of the master domain. If you then run copyDomain with the -relativizePaths option, the paths look like the following: <pre>(PGRP100 INFO) : "ldom3" (PGRP200 INFO) : "ldom3" (PGRP21 INFO) : "ldom0" (PGRP22 INFO) : "ldom1" (PGRP31 INFO) : "ldom2" (PGRP300 INFO) : "ldom3" (PGRP32 INFO) : "ldom2"</pre>

Notes for copyDomain

Note the following:

- If the `-dest` or `-pathToDestDomain` arguments are not provided, the utility no longer makes the paths to the subdomains relative paths. Instead, use the `-relativizePaths` argument.
- When used with `-clone` or `-noSubDomains`, `copyDomain` does not affect workbook metadata or hierarchies.
- Workbooks that are not included in the list used with the `-copyWorkbooks` option are not included in the new domain.
- Any existing workbooks in a domain copied with the `-clone` or `-noSubDomains` options may not be able to be committed back to the new domain.
- When used with `-dimDictOnly`, the `-clone` or `-noSubDomains` options cannot be specified.
- The `-dimDictOnly` switch implies `-copyWorkbooks none`.
- Use `-xmlConfigFile` to specify destination locations for individual subdomains.
- To get the usage text, use `-, -help`, or `-usage`.
- To get the version of this utility, use `-version`.
- To set the logger verbosity level, use `-loglevel` with the following values: `all`, `profile`, `debug`, `audit`, `information`, `warning`, `error`, or `none`.
- To disable timestamp header, use `-noheader`.

copyDomain: Format of the XML Configuration File

The XML configuration file contains source and destination fields for the location of the master domain and each of the subdomains. Here is a basic example:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes" ?>
<rpas>
  <globaldomain>
    <srcPath>C:\usr\Rpas\Domains\GlobalDomain</srcPath>
    <dstPath>C:\usr\Rpas\Domains\GlobalDomain2</dstPath>
    <subdomain>
      <srcPath>C:\usr\Rpas\Domains\GlobalDomain\ldom0</srcPath>
      <dstPath>C:\usr\Rpas\Domains\GlobalDomain2\ldom0</dstPath>
    </subdomain>
    <subdomain>
      <srcPath>C:\usr\Rpas\Domains\GlobalDomain\ldom1</srcPath>
      <dstPath>C:\usr\Rpas\Domains\GlobalDomain2\ldom1</dstPath>
    </subdomain>
    <subdomain>
      <srcPath>C:\usr\Rpas\Domains\GlobalDomain\ldom2</srcPath>
      <dstPath>C:\usr\Rpas\Domains\ldom2</dstPath>
    </subdomain>
    <subdomain>
      <srcPath>C:\usr\Rpas\Domains\GlobalDomain\ldom3</srcPath>
      <dstPath>C:\usr\Rpas\Domains\ldom3</dstPath>
    </subdomain>
  </globaldomain>
</rpas>
```

The `globaldomain` tag should contain one `srcPath` tag, one `stPath` tag, and a `subdomain` tag for each subdomain. Each `subdomain` tag should contain one `srcPath` tag and one `dstpath` tag. Each `srcPath` tag should be a path to either the master or subdomain begins copied. Each matching `dstPath` tag should be a path to where to copy that part of the domain.

The copyDomain utility validates the configuration xml file first before any files are copied. If any of the subdomain source paths do not match a subdomain path of the global domain being copied, a "can't find source subdomain 'subdomain' " error will be report. If the global domain being copied contains any subdomain that does not have a matching srcPath tag, a "subdomain 'subdomain' doesn't have a subdomain xml tag" error will be reported. If the global domain srcPath tag does not contain the path of a valid global domain then an "invalid source path 'srcPath' to global domain" will be reported.

The destination paths in all cases is validated when that part of the global domain is being copied. Unless the switch -force is provided, the destination must not exist and must be writable.

There are two options that control the number of subdomains to be copied. These options still limit the number of subdomains that are copied; however, the configuration file must still contain entries for all domains.

Moving a Domain Using moveDomain

The moveDomain utility provides the flexibility to move elements of global domains such as individual local domains and the master domain to pre-specified locations based on a given XML configuration file. The utility automatically updates RPAS metadata to reflect the modified directory paths in local and master domains. This utility also ensures that the globalDomainConfig.xml file is updated as domains are moved.

The XML configuration being used is simple and designed to fit the required task. It contains fields for the locations of the source master domain and destination master domain as well as source and destination fields for each of the subdomains that need to be moved.

moveDomain Usage

```
moveDomain -version
moveDomain -xmlConfigFile filename
moveDomain -d master -srcSubDomain src -dstSubDomain dst
```

[Table 9–3](#) provides descriptions of the arguments used by the moveDomain utility.

Table 9–3 Arguments Used by the moveDomain Utility

Argument	Description
-xmlConfigFile <i>pathToXmlConfigFile.xml</i>	Allows moveDomain to move a subdomain into user-instructed specific locations based paths specified in an XML file. This argument should not be used with the -d, the -srcSubDomain, and the -dstSubDomain parameters. See " moveDomain: Format of the XML Configuration File " for the file format.
-d <i>pathToMaster</i>	Allows moveDomain to move each subdomain based on the user-specified paths. Enter the path to the master domain.
-srcSubDomain <i>src</i>	Indicates the path of the subdomain to be moved.
-destSubDomain <i>src</i>	Indicates the path where the subdomain is to be moved.

moveDomain: Format of the XML Configuration File

The XML configuration being used is simple and designed to fit the required task. It contains fields for the locations of the source master domain and destination master

domain as well as source and destination fields for each of the subdomains that need to be moved. Here is a basic example of the XML configuration file.

```
<?xml version="1.0" encoding="UTF-8" standalone="yes" ?>
<rpas>
  <globaldomain>
    <srcPath>C:\usr\Rpas\Domains\GlobalDomain</srcPath>
    <dstPath>C:\usr\Rpas\Domains\GlobalDomain2</dstPath>
    <subdomain>
      <srcPath>C:\usr\Rpas\Domains\GlobalDomain\ldom2</srcPath>
      <dstPath>C:\usr\Rpas\Domains\ldom2</dstPath>
    </subdomain>
    <subdomain>
      <srcPath>C:\usr\Rpas\Domains\GlobalDomain\ldom3</srcPath>
      <dstPath>C:\usr\Rpas\Domains\ldom3</dstPath>
    </subdomain>
  </globaldomain>
</rpas>
```

The globaldomain tag must contain a srcPath tag and dstPath tag for the master domain. The master domain is not moved if srcPath and dstPath are the same. It is essential to specify srcPath and dstPath for the master domain even if the master is not intended to be moved; otherwise an error condition will be incurred.

The srcPath and dstPath tags for local domains are required only if the local domain is intended to be moved; otherwise, the lack of tags for a specific local domain indicates that the local domain will not be relocated. If srcPath and dstPath are identical for a given local domain, it will not be moved.

When global domain srcPath and dstPath are different, that is, when moving global domains, all local domains that reside under the global domain folder and are not included in the XML file, will be moved to the destination global domain folder. Other local domains with a specified destination location will be moved according to the configuration.

Assumptions and Requirements

The following rules apply to the XML configuration settings:

- All source and destination paths must be absolute.
- All source paths must correspond to existing directories.
- All destination paths must be valid, in the sense that:
 - The parent of the destination directory must exist.
 - The parent directory must be writable by the user.
 - The destination directory itself must not exist.
 - The source and destination master domain paths are required.
- The source and destination subdomain paths are required only for the domains that must be moved.
- The subdomains that must be moved can be specified or those subdomains that will remain under the master domain can be left out. If a subdomain is not specified, it will be moved along with the master domain.
- If the xmlConfigFile contents do not abide by the above mentioned rules, the utility does not clear the validation phase and terminates with the appropriate error message.

Minimum Space Requirement

Minimum space requirement for moving a global domain is the size of (only) the master domain plus the size of the largest local domain.

Setting Miscellaneous Domain Properties Using domainprop

Use the domainprop utility to manipulate the properties of a domain. Specify password properties, lock user accounts, and determine whether or not a daemon is currently managing a domain. The domainprop utility can be run on a global domain master to set values in all subdomains.

domainprop Usage

```
domainprop -d pathToDomain {-property propertyname=value}
```

Table 9–4 provides descriptions of the arguments used by the domainprop utility.

Table 9–4 Arguments Used by the domainprop Utility

Argument	Description
-d <i>pathToDomain</i>	Specifies the domain path.
-property <i>propertyname=value</i>	Specifies the property to be changed. See Table 9–5 for a list of properties that can be set with this utility. To view the current property setting, use the property command with no value.
-reportSubDomains	Shows property values for subdomains (should all match).

Available Properties

Table 9–5 Available Properties for the domainprop Utility

Domain Property Name	Type	Description
disable_commit_asap	Boolean	If this property is set to TRUE, the ability to use "Commit ASAP" in the File menu of the RPAS Client is disabled. Furthermore, the "Commit ASAP" radio button on the workbook close dialog is disabled. By default, this property is set to FALSE.
disable_commit_later	Boolean	If this property is set to TRUE, the ability to use "Commit Later" in the File menu of the RPAS Client is disabled. Furthermore, the "Commit Later" radio button on the workbook close dialog is disabled. By default, this property is set to FALSE.
domain_name	String	If this property is set, the domain name is displayed on the "About" page of the RPAS Client (menu Help > About).
insert_measure_disabled	Boolean	If this property is set to TRUE, the "insert measure..." item under the Edit menu is disabled. By default, this property is set to FALSE.

Table 9–5 (Cont.) Available Properties for the domainprop Utility

Domain Property Name	Type	Description
help_path	String	This property defines the path to the help files. It is set when using custom help files for the RPAS Client instead of the default help files that are provided with the RPAS Client installation. These files must be in WebHelp format. This path is normally a path to a network folder where the common help files are located. The folder must have a file named "Start.htm". This file is the starting point for the custom help. The default RPAS help has several pages for context-sensitive help. Help searches for these folders in the preset relative location from the root folder of WebHelp. If an implementation uses custom help, it must locate the context help files in the expected default locations relative to that folder. Otherwise, context-sensitive help will not be available and an error will occur every time the user tries to invoke context-sensitive help for a particular function. If this property is not set, the path specified in the foundation.ini file of the Client machine is used. If a path is not specified in that file, help looks in the default location in the RPAS Client installation for help files. If the default files have been removed, help is not available. For the location of the foundation.ini file, see "Translation Administration".
meas_fillclr_precedence	Boolean	By default, when deciding which color to fill a particular cell with, the RPAS grid uses the following order of formatting settings: Read-only, Measure, Hierarchical, and then Read/Write. That is, if the cell is in a read-only state, it will use the read-only formatting setting. However, if that is not the case, the grid will check if there is any Measure level formatting. Failing to find it will fall through to checking for the hierarchical setting and then the read-write setting. However, some customers want RPAS to follow a different priority order for fill color formatting decision making. They may want it to try Measure, then Read-only, then Hierarchical, and finally Read/Write. This change from default can be made by setting this domain property to TRUE. To reset behavior, this domain property can be reset to FALSE.
measure_locking_disabled	Boolean	If this property is set to FALSE, the user can lock a measure on a work sheet. By default, this property is set to FALSE. To disable measure locking, set this property to TRUE.

Table 9–5 (Cont.) Available Properties for the domainprop Utility

Domain Property Name	Type	Description
ovr_def_admin_privileges	Boolean	Using the Security Administration workbook, administrators can set workbook template access for every user in the system. Non-administrative users cannot access the workbook templates to which they have not explicitly been given access. However, if a user is an administrator, by default that user can see all the workbooks in the system. Some retailers want to prevent this from happening. Reasons for this include reducing clutter and having different kind of administrators who manage different administrative tasks in their RPAS systems. Ability to control template access for administrators from the Security Administration workbook is made possible by setting this domain property to TRUE. By default, this property is FALSE.
r_upgrade_hash_alg	Integer	Sets the value for r_crypt_hash_alg in the upgraded domain. Because the list of algorithms can change with each release, it is not feasible to provide a list here. To get the list of values that can be used here, go to the rpasHome of the version to which you are upgrading and run: <pre>usermgr -listHashAlgs</pre>
r_upgrade_has_iter	Integer	This property sets the value for r_crypto_hash_iter in the upgrade process. It is the number of iterations of hashing to perform on the password.

Using the mace Calculation Engine

The mace utility (Multi-dimensional Array Calculation Engine) allows the administrator to evaluate rule groups or expressions in order to manipulate measures. The mace utility supports the use of the RPAS calculation engine in batch.

The mace utility is most commonly used to run a rule group or an expression, but can also be used to:

- create rules and rule groups
- add rules to rule groups
- add expressions to rules
- delete rules not contained in a rule group
- remove any or all rule groups
- validate expressions
- print a list of rules or rule groups

Parallelization

The mace utility can execute in parallel under the following circumstances:

- The utility must be invoked on a master domain.
- Parallelization is only applicable to single-expression evaluation (-run -expression argument). Parallelization does not apply to rule group evaluation.
- The evaluated expression cannot be a SpecialExpression.

- All of the measures appearing on the left-hand side of the expression must be non-HBI; that is, the base intersection of the measures must be below the partition level.

The mace utility creates multiple child processes based on the `-processes` argument, and each child mace process evaluates the expression in one local domain. This functionality enables mace to achieve higher levels of CPU utilization using parallelization on systems with multiple CPUs. It also simplifies the user script when the same expression must be evaluated in all local domains.

Centralization

When mace is run on a master domain, the following command line options apply to the master as well as all local domains. For example, running `mace -d domain -newRule ...` creates a new rule in the master and all local domains.

- `-newRule`: create a new rule in the domain
- `-delRule`: delete an existing rule from the domain
- `-addRule`: add a new rule to a specific rule group
- `-removeRule`: remove an existing rule from a specific rule group
- `-newGroup`: create a new rule group in the domain
- `-remove Group`: remove an existing rule group from the domain
- `-addExpression`: add an expression to a specific rule
- `-purgeRules`: remove all rules not contained in any rule group from the domain
- `-removeAllRuleData`: remove all rule and rule group data from the domain

The behavior and usage of the following commands is unchanged:

- `-find`: search all expressions for the specific measure and print all rules and rule groups that use it
- `-check`: validate the specific expression
- `-resolve`: order but do not evaluate expressions within the rule group
- `-transit`: rule calc engine by transitting over a list of rule groups
- `-print`: print all specific rules and groups
- `-validate`: validate rule groups

mace Usage

```
mace -version
mace -d domainPath -find string
mace -d domainPath -newRule {-ruleName ruleName} {-label ruleLabel}{-processes numProcesses}
mace -d domainPath -delRule ruleName {-processes numProcesses}
mace -d domainPath -addRule      groupName:ruleName {-processes numProcesses}
mace -d domainPath -removeRule  groupName:ruleName {-processes numProcesses}
mace -d domainPath -newGroup   groupName {-label groupLabel}{-processes numProcesses}
mace -d domainPath -removeGroup groupName {-processes numProcesses}
mace -d domainPath -addExpression ruleName -expression exprString{-processes numProcesses}
mace -d domainPath -check -expression exprString
mace -d domainPath -run -group groupName      {-debugRuleEngine}
mace -d domainPath -run -expression expString {-processes
```

```

numProcesses}{-debugRuleEngine}
mace -d domainPath -resolve groupName -measures measList {-debugRuleEngine}
mace -d domainPath -transit workbookName -group groupList {-debugRuleEngine}
mace -d domainPath -print -rule ruleList
mace -d domainPath -print -group groupList
mace -d domainPath -print -allGroups
mace -d domainPath -purgeRules {-processes numProcesses}
mace -d domainPath -removeAllRuleData {-processes numProcesses}
mace -d domainPath -validate calc -ruleGroup groupName
mace -d domainPath -validate general -ruleGroup groupName
mace -d domainPath -validate refresh -ruleGroup groupName -calcRuleGroup calc

```

Table 9–6 provides descriptions of the arguments used by the mace utility.

Table 9–6 Arguments Used by the mace Utility

Argument	Description
-d <i>domainPath</i>	Specifies the domain in which to load the measure.
-find <i>string</i>	Searches all expressions for the specified string and prints all the rules and rule groups that have these expressions.
-newRule {-ruleName <i>ruleName</i> }	Creates a new empty rule. If desired, use the -ruleName argument to specify a name for the rule.
-label {ruleLabel groupLabel}	Specifies the label of the rule with the -newRule argument or label of the group with the -newGroup argument.
-processes <i>numProcesses</i>	Specifies the number of child processes to be run in parallel.
-delRule <i>ruleName</i>	Removes the specified rule.
-addRule <i>groupName:ruleName</i>	Adds the specified rule to the group specified by <i>groupName</i> .
-removeRule <i>groupName:ruleName</i>	Removes the specified <i>ruleName</i> from the group specified by <i>groupName</i> .
-newGroup <i>groupName</i>	Creates a new rule group with the specified name.
-removeGroup <i>groupName</i>	Removes the specified group and non-shared rules in it.
-addExpression <i>ruleName</i>	Adds an expression to the specified rule. The -expression should be used with this argument.
-check	Validates the specified expression. The -expression should be used with this argument.
-run	Evaluates the specified expression or rule group. The -expression should be used with this argument.
-resolve <i>groupName</i>	Orders (does not evaluate) expressions within rule group. Requires a comma-separated list of edited measures.
-transit <i>workbookName</i>	Runs a calc engine by transitioning over a list of rule groups. Requires the name of an existing workbook and a comma-separated list of rule-group names.
-print {ruleList groupList true}	Prints all the specified rules and rule groups. The ruleList is a comma-separated list of rule names. The groupList is a comma-separated list of group names. If "true" is supplied for either ruleList or groupList, all rules or rule groups are printed.
-purgeRules	Removes all rules not contained in any rule groups.
-removeAllRuleData	Removes all rule groups and all rules.

Table 9–6 (Cont.) Arguments Used by the mace Utility

Argument	Description
-validate {calc general refresh}	Validates rule groups. Use calc to validate a calc rule group. To validate a refresh rule group, use the refresh parameter along with -calcRuleGroup to specify the corresponding calc rule group. For all other types of rule groups, use general.
-debugRuleEngine	Generates a file mace.log in the working directory for logging RuleEngine specific debug information.
-expression <i>exprString</i>	Specifies the expression. This argument is used in conjunction with the -addExpression, -check, and -run arguments.
-group <i>groupName</i>	Specifies the rule group to evaluate using the -run argument.
-measures <i>measureList</i>	Specifies the measures to resolve.
-group <i>groupList</i>	Specifies a list of group names, separated by commas. Use this argument in conjunction with the -transit and -print arguments.
-rule <i>ruleList</i>	Specifies a list of rule names, separated by commas. Use this argument in conjunction with the -print argument.
-allGroups	Use this argument in conjunction with the -print argument to print all rule groups.
-addGroup	Creates a new rule group with the specified name
-beginBatch	Instructs mace to set the batch status flag for domains integrated with an RPAS Data Mart.
-endBatch	Instructs mace to perform the end-of-batch data synchronization and clean-up operations associated with the conclusion of a batch process for a domain integration with an RPAS Data Mart.
-abortBatch	Instructs mace to perform the clean-up operations associated with the early termination of a batch process in a domain integrated with an RPAS Data Mart.

mace Restrictions in Global Domains

Problems were found when running simple rule groups using the mace utility. Mace did not allow executing a rule group consisting of non-HBI measures from the master domain for the following reasons:

1. Even though the rule group contains only non-HBI measures, mace was not designed to run it from the master domain and spawn it across local domain as it required too many checks to make it work differently for special expressions as sometimes the special expressions could take more than one left-hand side measure and some of them could be HBI and some non-HBI.
2. It is not a properly configured rule group if the expressions in the rule group have both HBI and non-HBI measures which need to be evaluated. RPAS may not be able to detect this at the time of configuring the solution as the partition dimension can be changed in the globalDomainConfig.xml.
3. The rule group may contain special expressions and regular RPAS expressions (as in 1 and 2) and that adds to the complexity to the rule engine to properly evaluate these expressions when run from the master domain.

4. The above restrictions do not apply when mace is run on a single expression where the LHS can be a HBI or non-HBI measure.

There is only one situation where a rule group can be run from the master domain and that is when all the LHS measures of the expressions in the rule group are HBI measures. This type of rule group cannot be run from local domains and needs to be run from the master domain.

Proper validation is added to the mace utility to identify these conditions and flag them to the user and exit gracefully. The Oracle Retail Predictive Application Server Administration Guide will be updated with these additional validation points in future release. The following validations are added:

1. If mace is executed against a master domain and the executing expressions within the rule group contain non-HBI measures on the left-hand side of expressions, excluding special expressions, mace exits with an exception.
2. If mace is executed against a local domain and the executing expressions within the rule group contain an HBI measure on the left-hand side, excluding special expressions, mace outputs a warning message but the execution continues.

The end result is that mace cannot execute a rule group that contains non-HBI measures on the LHS against the master domain and provides a warning when running a rule group with HBI measures on the LHS against a local domain.

Executing mace in an Integrated Environment

When operating on a domain that shares data through an RPAS Data Mart, mace may encounter situations in which one or more of the measures involved in a calculation are shared measures whose data is stored within the RDM. To support this situation, users can make use of the `-beginBatch` and `-endBatch` operations to cause mace to interact with the RPAS Data Mart associated with a domain during calculation.

The use of `-beginBatch` sets a status flag within the domain. When this status is active, mace evaluates all rule groups for the presence of shared measures prior to execution. In the case where shared measures are present in the calculation, mace makes use of the `transferFactData` utility to fetch the measure data for the shared fact from the RPAS Data Mart.

For the duration of the batch operation, shared measures whose data are updated by mace hold their data within the domain. At the conclusion of the batch, if mace is called with the `-endBatch` operation, mace determines the set of shared measures whose data has been modified and writes those measures back to the RPAS Data Mart.

Upon successful completion of the write-back operation, mace deletes the temporary measure copies within the domain and clears the status flag set by `-beginBatch`. If an error results in the need for a batch process to end prior to the point where `-endBatch` would normally complete the batch process, the `-abortBatch` operation will clear the status flag set by `-beginBatch` without performing the write-back of data to the RPAS Data Mart. This prevents partially or incorrectly updated data from entering the RDM.

The following error and warning conditions exist in association with `-beginBatch`, `-endBatch` and `-abortBatch`:

- A second call to `-beginBatch` without either `-endBatch` or `-abortBatch` having been called.
- A call to `-endBatch` without `-beginBatch` having first been called.

- Any mace call that evaluates an expression (-run -group or -run -expression), if any included measure is an RDM-shared measure, and -beginBatch has not first been called.
- Any call to -beginBatch, -endBatch, or -abortBatch with the -d parameter referencing a subdomain.

The following conditions are reported by mace as a warning, with a successful (zero) return code:

- A call to -abortBatch or -printBatchStatus when -beginBatch was not first called.
- A call to -beginBatch, -endBatch, or -printBatchStatus on a domain that is not HSA-integrated.

Managing the Workbook Batch Queue Using wbbatch

The wbbatch utility is used to manage workbook batch categories and workbooks in the workbook batch queue. The workbook batch queue is updated by using the standard RPAS wizard Auto Workbook Build or using various options of the wbbatch utility.

The most common use of this utility is to build workbooks that have been scheduled to be automatically built using the Auto Workbook Build wizard in the RPAS clients. It is also used to add, update, and delete batch categories, update assignments of workbook build entries to workbook batch categories, provide workbook batch categories when adding workbooks to the refresh queue, and update the assignments to workbooks already in the refresh queue.

When a user defers a workbook commit (using Commit Later), that workbook commit process is added to the Commit Later queue, which is committed using this utility. An administrator can also add a workbook to the commit later queue with this utility.

RPAS provides the ability to update workbook data with domain data without having to rebuild the workbook; this refreshing process is completed using a workbook's default refresh rule group. Workbooks are added to the queue to be refreshed and refreshed using this utility.

The build and refresh operations can be executed in multiple, parallel processes using the -processes argument.

wbbatch Usage

```
wbbatch -version
wbbatch -d pathToDomain -build queueIndex
wbbatch -d pathToDomain -refresh workbookName
wbbatch -d pathToDomain -commit workbookName
wbbatch -d pathToDomain -scheduleRefresh workbookName [-category categoryName]
wbbatch -d pathToDomain -unscheduleRefresh workbookName
wbbatch -d pathToDomain -scheduleCommit workbookName
wbbatch -d pathToDomain -unscheduleCommit workbookName
wbbatch -d pathToDomain -startQueue [all|build|refresh|commit] [-processes max]
[-categories catName1,catName2]
wbbatch -d pathToDomain -printQueue [all|build|refresh|commit]
wbbatch -d pathToDomain -listCategories
wbbatch -d pathToDomain -addCategories catName1:Label1,catName2:Label2
wbbatch -d pathToDomain -deleteCategories catName1,catName2
wbbatch -d pathToDomain -changeCategoryLabels catName1:NewLabel1,catName2:NewLabel2
wbbatch -d pathToDomain -queue build -updateCategories queueIndex1:newCatName1,
```

```
queueIndex2:newCatName2
wbbatch -d pathToDomain -queue refresh -updateCategories
workbookName1:newCatName1, workbookName2:newCatName2
```

Table 9–7 describes the arguments used by the wbbatch utility.

Table 9–7 Arguments Used by the wbbatch Utility

Argument	Description
-d pathToDomain	Specifies the domain containing the workbooks.
-build <i>queueIndex</i>	Runs workbook build for provided <i>queueIndex</i> .
-refresh <i>workbookName</i>	Refreshes workbooks scheduled to be refreshed using this utility. To refresh a single workbook in the queue, specify the name of the workbook. If no name is provided, all workbooks scheduled to be refreshed will be completed.
Commits workbooks with deferred commits. To commit a single workbook in the commit later queue, specify the name of a workbook. If no name is provided, all workbooks in the commit later queue will be committed.	-commit <i>workbookName</i>
-processes <i>count</i>	Used with either -build or -refresh to build or refresh workbooks in the auto-workbook queue in parallel using the specified number of parallel processes.
-scheduleRefresh	Schedules a workbook to be refreshed later by adding it to the workbook refresh batch queue. If the -category option is specified, the scheduled workbook will be in that category. Otherwise, it will be in the default category.
-unscheduleRefresh <i>workbookName</i>	Removes a workbook from the workbook refresh batch queue.
-scheduleCommit <i>workbookName</i>	Schedules a workbook to be committed later by adding it to the workbook commit batch queue.
-unscheduleCommit <i>workbookName</i>	Removes a workbook from the workbook commit batch queue.
-startQueue	Runs all workbooks in provided queue. The queue options are build, refresh, and commit. If the -category option is used and one or more categories are specified, only the workbooks in those categories are built or refreshed. Categories do not apply to committing.
-printQueue	Prints the contents of the queue argument. The queue indexes for auto workbooks in the build queue are shown when printing the build queue. If “all” is specified, all three queues (build, refresh, and commit) are displayed.
-listCategories	Lists both the name and label for all categories.

Table 9–7 (Cont.) Arguments Used by the wbbatch Utility

Argument	Description
-addCategories <i>Name1:Label1, Name2:Label2, Name3:Label3</i>	Adds a new category by providing a name and label, separated by a colon. Multiple categories can be specified on the same command line if separated by a comma. If the users use a different language other than the one typed in the command line, the administrator should use the Workbook Batch Category Management wizard to create new categories.
-deleteCategories <i>catName1, catName2</i>	Deletes a category by specifying the name of that category. Multiple categories can be deleted if separated by a comma.
-changeCategoryLabels <i>Name1:NewLabel1, Name2:NewLabel2</i>	Changes the label of an existing category by specifying the category name and providing a new category label. If the users use a different language other than the one typed in the command line, the administrator should use the Workbook Batch Category Management wizard to change category labels.
-updateCategories <i>queueIndex1:newCatName1, queueIndex2:newCatName2</i> Or: -updateCategories <i>wbName1:newCatName1, wbName2:newCatName2</i>	Updates the category for an entry in the build queue or to change the workbook category of an existing entry in the refresh queue. Multiple category assignment for workbook auto build queue entries and refresh entries can be updated. If using the -build option, list the queue index. If using the -refresh option, list the workbook name. Build example: wbbatch -d pathToDomain -queue build -updateCategories <i>queueIndex1:newCatName1, queueIndex2:newCatName2</i> Refresh example: wbbatch -d pathToDomain -queue refresh -updateCategories <i>workbookName1:newCatName1, workbookName2:newCatName2</i>

Managing Workbooks Using wbmgr

Use the Workbook Manager utility to inspect or remove the existing workbooks. It is recommended that administrators use this utility to remove workbooks rather than doing so manually.

wbmgr Usage

```
wbmgr -version
wbmgr -d pathToDomain -list -all
wbmgr -d pathToDomain -list -user userName
wbmgr -d pathToDomain -print -wbList wb1,wb2,...
wbmgr -d pathToDomain -remove -all
wbmgr -d pathToDomain -remove -user userName
wbmgr -d pathToDomain -remove -user userName -wbList wb1,wb2,...
```

Table 9–8 provides descriptions of the arguments used by the wbmgr utility.

Table 9–8 Arguments Used by the wbmgr Utility

Argument	Description
-d <i>pathToDomain</i>	Specifies the domain that contains the workbooks.
-list -all	Lists all workbooks in the domain.
-list -user <i>userName</i>	Lists all workbooks belonging to the user.
-print -wbList <i>wb1,wb2,...</i>	Prints detailed information about workbooks in the list.

Table 9–8 (Cont.) Arguments Used by the wbmgr Utility

Argument	Description
-remove -all	Removes all workbooks from the domain.
-remove -user <i>userName</i>	Removes all workbooks from the domain belonging to the specified user.
-remove -user <i>userName</i> -wbList <i>wb1,wb2</i>	Removes all the workbooks in the specified list for the specified user.

Registering Measures Using regmeasure

The regmeasure utility is used for batch measure registration. The following functionality is included:

- Register a new measure in the user-specified domain with the user-specified measure properties. If the domain specified by the user is a global domain, this measure will be registered in the master domain and all its local domains. The user must provide a minimum set of measure properties, type, and base intersection. Other measure properties are optional, such as default aggregation and spreading method. If the user omits an optional measure property, the measure will be registered with default value of that property.
- Unregister an existing measure identified by its name from the user-specified domain. If the specified domain is a global domain, this measure will be removed from the master domain and all local domains. Unregistering a measure from a domain causes the measure definition and all the related measure data arrays and supporting arrays to be removed from the domain.
- Modify measure properties of an existing measure. Not all measure values can be modified, such as type, base intersection, NV value, and database name. These properties cannot be changed once the measure is registered. Measure properties such as default aggregation method, default spread method, base state, agg state, and so on can be modified after the measure is registered.

regmeasure Usage

```
regmeasure -version
regmeasure -d pathToDomain -add measureName -type typeName(-baseint
baseIntersection|-scalar) {-label labelString} {-db dataDbPath} {-navalue naValue}
{-defagg aggType} {-defspread spreadType} {-allowedaggs "aggType1
aggType2"} {-refreshable (true/false)} {-insertable (true/false)} {-basestate
(read/write)} {-aggstate (read/write)} {-stageonly (true/false)} {-filename
fileName} {-loadint loadIntersectionString} {-clearint
clearIntersectionString} {-loadstokeep loadsToKeep} {-start fieldStart} {-width
fieldWidth} {-loadagg loadAgg} {-range range} {-purgeage purgeAge} {-viewtype
viewType} {-syncwith syncWith} {-description descriptionString}
{-picklist} {-materialized (persistent/display)} {-lowerbound measurename}
{-upperbound measurename} {-attr attrName -attrpos attrPosName} {-scriptname
scriptName} {-specialval action:specval:behavior,action:specval:behavior,...}
{-fnhbi} {-hybridaggspec hiername:aggop,hiername:aggop,...} {-periodstartvalue
(true/false)}

regmeasure -d pathToDomain -modify measureName {-label labelString} {-defagg
aggType} {-defspread spreadType} {-allowedaggs "aggType1
aggType2..."} {-refreshable (true/false)} {-insertable (true/false)} {-basestate
(read/write)} {-aggstate (read/write)} {-stageonly (true/false)} {-filename
fileName} {-clearint clearIntersectionString} {-loadstokeep loadsToKeep} {-start
fieldStart} {-width fieldWidth} {-loadagg loadAgg} {-range rangeString} {-purgeage
```

```

purgeAge|-clearPurgeAge}{-viewtype viewType} {-syncwith syncWith} {-description
descriptionString}{-picklist|-nopicklist} {-materialized
(persistent|display)}{-lowerbound measureName} {-upperbound measureName}{-attr
attrName -attrpos attrPosName} {-scriptname scriptName}{-specialval
action:specval:behavior,action:specval:behavior,...}{-hybridaggspec
hiername:aggOp,hiername:aggOp,...}{-periodstartvalue (true|false)}
regmeasure -d pathToDomain -remove measureName

```

Table 9–9 provides descriptions of the arguments used by the regmeasure utility.

Table 9–9 Arguments Used by the regmeasure Utility

Argument	Description
-d <i>pathToDomain</i>	Specifies the path to the domain. A valid domain path must be specified.
-add <i>measureName</i>	Adds a measure with the specified name. Set the values for the measure by using the required arguments and any of the optional arguments. Measure names can be up to 30 characters long.
-type <i>typeName</i>	Specifies the measure data type. It can be set to int, real, string, date, or boolean. Required with the -add option. Not available with the -modify option.
-baseint <i>baseIntersection</i> -scalar	Specifies the base intersection of the measure. Non-scalar measures must use the -baseint option. Scalar measures must use the -scalar option. Required with the -add option. Not available with the -modify option.
-label <i>labelString</i>	Specifies the measure label. If not specified, it defaults to the measure name specified for the -add option.
-db <i>dataDbPath</i>	Specifies the database path for the measure's data arrays. A valid database path name must be specified. If not specified, the measure will be registered without a database. As a result, the measure will not be able to store any data in the domain. However, if the measure is not a Display only type, it will still be assigned a database in the workbook. Not available with the -modify option.
-navalue <i>naValue</i>	Specifies the na value for the measure's base level data array. The navalue must be the same type as the measure. For date, the navalue must be formatted as 'YYYYmmddHHMMSSsss'. If not specified, it defaults to the type's default value: 0 for numeric type, false for boolean type, an empty string for string type, and 0001/01/01 for date type. Not available with the -modify option.
-defagg <i>aggType</i>	Specifies the default aggregation method for the measure. It must be an aggregation name valid for the type of measure. For a list of valid aggregation type names, see the <i>RPAS Configuration Tools User Guide</i> . If not specified, it defaults to the measure type's default aggregation method: Total for int and real, Ambig for string and date, and OR for boolean.
-defspread <i>spreadType</i>	Specifies the default spread method for the measure. It must be a spread method valid for the type of measure. For a list of valid spread methods, see the <i>RPAS Configuration Tools User Guide</i> . If not specified, it defaults to the measure type's default spread method: Ratio for int and real, and Replicate for string, date, and boolean.

Table 9–9 (Cont.) Arguments Used by the regmeasure Utility

Argument	Description
-allowedaggs "aggType1 aggType2..."	Specifies a list of aggregation methods that are allowed for this measure. The aggregation methods must be valid for the type of measure. If not specified, it defaults to the default allowed aggs for the type of measure. For numeric (int or real type) measures: total, total_pop, first, first_pop, last, last_pop, min, min_pop, max, max_pop, average, average_pop, popcount, nobcount, ambig, ambig_pop, none, period_start_total, period_end_total, period_start_average, period_end_average, median, median_pop, recalc, hybrid. For string type measures: ambig, ambig_pop, none, popcount, nobcount, first, first_pop, last, last_pop, recalc, hybrid. For date type measure: ambig, ambig_pop, pop_count, nob_count, first, first_pop, last, last_pop, min, min_pop, max, max_pop, non, recalc, hybrid. For boolean measure: boolean_and, boolean_or, pop_count, nob_count, ambig, ambig_pop, none, first, first_pop, last, last_pop, recalc, hybrid.
-refreshable (true false)	Note: This option is no longer supported but is kept for compatibility.
-insertable (true false)	Specifies whether the measure can be dynamically inserted into the workbook. If not specified, it defaults to true.
-basestate (read write)	Specifies the workbook access right for the base array of the measure. If not specified, it defaults to read. The access rights of this measure will be further restricted by the RPAS security features. As a result, write access specified by this option does not guarantee write access of this measure in a specific workbook.
-aggstate (read write)	Specifies the workbook access right for the aggregated level of the measure. If not specified, it defaults to read. The access rights of this measure are further restricted by the RPAS security features. As a result, write access specified by this option does not guarantee write access of this measure in a specific workbook.
-stageonly (true false)	Specifies whether the measure is a stage only measure. If not specified, it defaults to false. Measure data loaded by loadmeasure for stage only measures are not automatically applied to the measure's base data array. User intervention is usually required to manually approve the loaded measure data and apply the approved loads to the measure's base data array.
-filename fileName	Specifies the file name of this measure's loading file. It should not include any extensions. If not specified, it defaults to the measure name in lower case.
-loadint loadIntersectionString	Specifies the intersection to load data for this measure. It must be a valid intersection string that is either the same or lower than the base intersection of this measure. If loadint is lower than the base intersection of the measure, the aggregation method specified by the -loadagg option will be used to aggregate the loaded data to the base array of the measure. Not available with the -modify option.
-clearint clearIntersectionString	Specifies the clear intersection for the clear load of this measure. For more information on the various loading methods including clear load, refer to the Loading Measure Data - loadmeasure section in this guide.

Table 9–9 (Cont.) Arguments Used by the regmeasure Utility

Argument	Description
-loadstokeep loadsToKeep	Specifies the number of temporary measure load arrays to be kept in the staging database. If not specified, it defaults to 1.
-start fieldStart	Specifies the starting column of this measure's data in the measure loading file. If not specified, it is calculated based on the loadint of the measure.
-width fieldWidth	Specifies the number of characters this measure's data occupies in the measure loading file. If not specified, it defaults to the default width of the measure type: 8 for integer, real, and date, 24 for string, and 1 for boolean.
-loadagg loadAgg	Specifies the aggregation method used to aggregate the temporary load array to the measure's base array if the measure's loadint is lower than its baseint. If not specified, it defaults to the measure type's default aggregation method: Total for int and real, Ambig for string and date, and OR for Boolean.
-range rangeString	Specifies the valid range for the measure. The value of the range parameter depends on the measure type. For int or real types, the format is min:max, where min is the lowest possible value of the measure and max is the highest possible value of the measure. For picklist measures, to give the allowed options, the format of the string argument is 'a(Label A),b(Label B),c,d', where a, b, c, and d are allowed measure values and Label A and Label B are optional labels for the values. In addition, the list of allowed options can be changed dynamically with the cell the user is clicking in. For this functionality, the measure's range is specified as 'measurerange=measurename' where measure name is the name of the measure that contains strings in the above format of value/label pairs. For date types, the range must be in the format MMddyyyyhhmmss:MMddyyyyhhmmss, where the first string is the starting date and time of the range and the second date is the ending date and time of the range. If you omit the time portion of the string (hhmmss), the default time is used. The default lower bound is 000000 (12:00:00AM), which is used for the beginning of the day. The default upper bound is 235959 (11:59:59PM), which is used for the end of the day. If the range begins with a negative number (which may confuse the command-line argument parser), enclose the entire range string in square brackets, such as -range [-10:10].
-purgeage (purgeAge)	Specifies the number of days (or whatever the base dimension of the calendar hierarchy is) of measure data that should be kept in the measure's base data array after the measure load. This is used to keep the measure's data size small. If not specified, it defaults to -1, in which case the measure data will never be purged. When using the -modify option, -purgeage or -clearPurgeAge can be specified.
-clearPurgeAge	Resets the number of days (or whatever the base dimension of the calendar hierarchy is) of measure data that should be kept in the measure's base data array after measure load to -1. This means that the measure data is never purged. clearPurgeAge is only available with the -modify option. When using the -modify option, -purgeage or -clearPurgeAge can be specified.

Table 9–9 (Cont.) Arguments Used by the regmeasure Utility

Argument	Description
-viewtype viewtype	Specifies the view type of this measure on the RPAS Client. The valid values are: 0 for none, 1 for view_only, 2 for sync_first_lag, 3 for sync_lead_last, 4 for sync_first, and 5 for sync_last. If not specified, it defaults to none. If the view type starts with "sync", the measure is called a 'Virtual Measure'. A measure of sync_first_lag type must have two sync measure names specified by the -syncwith option. The first syncwith measure name is a 'Period Start Value' type of measure, like opening stock. Measure data at the beginning period of the calendar is synchronized with this period start value kind of measure. The subsequent measure data is synchronized with the other measure data but lagged one period. A measure of sync_lead_last type must have two sync measure names specified by the -syncwith option. The first measure is a 'Period End Value' type of measure. Measure data at the last period of the calendar is synchronized with this period end value. Measure data of previous periods is synchronized with the other measure lead one period data. A measure of sync_first type must have one measure name specified by the -syncwith option. The data of the beginning period is synchronized with this syncwith measure. A measure of sync_last type must have one measure name specified by the -syncwith option. The data of the ending period is synchronized with this syncwith measure. Measures of view_only type are non-persistent. View only measures can only be used in workbooks. Their measure data is calculated during the Fetch process using a calc expression usually specified in the workbook's calc rule group.
-syncwith syncWith	Specifies the measures that the measure must be synchronized with. This option must be specified if the measure is not a virtual measure. For sync_first_lag and sync_lead_last measures, the syncwith option must have two measure names separated by a comma. The first measure is used to synchronize the data at the first or the last calendar period. The second measure is used to synchronize data at other periods. For sync_first and sync_last measures, the syncwith option must be specified with a single measure name that will be used to synchronize the first or last calendar period.
-description descriptionString	Specifies the description of the measure.
-picklist -nopicklist	Specifies whether the measure is displayed as a picklist in the Client. The actual value of the picklist is specified by the -range option of the measure. -nopicklist is only available with the -modify option. It means the measure should not be displayed as a picklist measure in the RPAS Client.

Table 9–9 (Cont.) Arguments Used by the regmeasure Utility

Argument	Description
-materialized (persistent display)	Specifies whether the measure is persistent or display only on the RPAS Server side. Persistent measures must have a valid database and arrays to store the measure data. Display only measures do not have permanent data arrays associated with it. The data for a display only measure must be calculated on the fly. As a result, display only measures can not be used on the RHS of any expression. Display only measures can still be used on the LHS of a calc expression used in a workbook, in which case a temporary array will be created in the workbook to hold the temporary data for the display measure.
-lowerbound measureName	Specifies a measure name that defines the lower bound for each cell of the measure. The difference between the -lowerbound and -range options is that the -range option specifies a single scalar as the lower bound for all cells of the measure, but the lower bound value specified by the -lowerbound option can be different from cell to cell.
-upperbound measureName	Specifies a measure name that defines the upper bound for each cell of the measure. The difference between the -upperbound and -range options is that the -range option specifies a single scalar as the upper bound for all cells of the measure, but the upper bound value specified by the -upperbound option can be different cell to cell.
-attr attrName	Specifies the measure attribute name. If not specified, it defaults to no attribute is assigned to the measure. Note: If this option is specified, the -attrpos option must also be specified.
-attrpos attrPosName	Specifies the measure attribute position name. Combined with the -attr option, the measure attribute provides a way to group measures together based on measure attributes. Note: If this option is specified, the -attr option must also be specified.
-scriptname scriptname	Specifies a shell script that must be executed as part of a specific event. Currently, the only script that is handled is to give the option of selecting a hierarchy position name as the content of a string measure. In other words, when a user clicks in a cell, the user is presented with a hierarchy dimension single-tree pop-up. The format for this is 'SingleSelect(HIER="<HIER>", DIM="<DIM>")' where <HIER> and <DIM> should be replaced with the actual names of the hierarchy and dimension for which the single-tree pop-up should be created.
-specialval action:specval:behavior, action:specval:behavior,...	Specifies a list of measure special values in the form of "Action:SpecialValue:Behavior,...". The special values are stored in the domain's meta data database. For Action, the only action supported is: "DISPLAY". The only SpecialValue supported is "NAVAL". For Behavior, "NULL" means translate any na cell to a blank cell for display. "CELLVALUE" means no translation, just display the navalue as a regular value.

Table 9–9 (Cont.) Arguments Used by the regmeasure Utility

Argument	Description
-fnhbi	Specifies that this measure is a "Forced non=HBI" measure, which means that although the base intersection of this measure is above the partition dimension, the measure data must still be stored in each local domain. Not available with the -modify option.
-hybridaggspec hiername:aggOp,hiername:aggOp,...	Specifies the aggregation method to be used for each hierarchy in the base intersection. This option is only valid when the default aggregation method for the measure is hybrid.
-periodstatevalue (true false)	Specifies that this measure stores a "Period Start" type of data, like beginning inventory. PeriodStart measures usually use Period Start Total or Period Start Average for the default aggregation method. It also has different behavior in elapsed lock. At the aggregated calendar level, if the starting period is elapse locked, then the whole aggregated period is locked.
-modify measureName	Modifies the measure with the specified name. Set the updated values for the measure by using any of the optional arguments.
-remove measureName	Removes the measure with the specified name.

Registering Token Measures Using regTokenMeasure

The regTokenMeasure utility is used to register, list, and remove RPAS Token Measures.

RPAS Token Measure provides placeholder functionality for measure names in RPAS expressions. An RPAS Token Measure is a special RPAS measure.

An RPAS Token Measure is always registered as a scalar measure of string type, with the measure property called tokenmeas set to true. Its measure data holds a valid value measure name as a single string. The data arrays for all token measures are stored in one database called token under the data directory in the RPAS domain.

Token measure can be used in RPAS expressions by prefixing @ in front of the token measure name, either on the LHS or RHS of the expression. Before evaluation, @TokenMeasName in the expression is replaced with the value measure name that is associated with the token measure. As a result, the expression is evaluated against the value measure. A token measure name cannot be used in expression without the prefixing @.

In the following example, TM1 is a token measure registered with the value measure name VM1.

The following expression:

```
@TM1 = a + b
```

Is evaluated as:

```
VM1 = a + b
```

The following expression is not valid, because TM1 is used without prefixing it with @:

```
TM1 = "sth"
```

If mace is used to evaluate, it will throw a ParseException with the message that the token measure "TM1" is used without prefixing @. This functionality prevents the modification of the token measure's data, which is actually the value measure's measure name.

The regTokenMeasure utility cannot be used to create a token measure which is a measure shared through an RPAS Data Mart. If regTokenMeasure is called to register a shared measure as a token measure, regTokenMeasure will exit with an error identifying the shared measure.

regTokenMeasure Usage

```
regTokenMeasure -version
regTokenMeasure -d pathToDomain -add tokenMeasure=valueMeasure {-fnhbi}
regTokenMeasure -d pathToDomain -list
regTokenMeasure -d pathToDomain -remove tokenMeasure=valueMeasure
```

Table 9–10 provides descriptions of the arguments used by the regTokenMeasure utility.

Table 9–10 Arguments Used by the regTokenMeasure Utility

Argument	Description
-d <i>pathToDomain</i>	Specifies the path to the domain. A valid domain path must be specified.
-add <i>tokenMeasure=valueMeasure</i>	Adds a token measure with the specified token measure name and value measure that the token measure points to.
-fnhbi	If specified, the token measure will be registered as an fnhbi measure in the global domain. Its data is stored in each local domain rather than the global domain, although by definition the token measure should be always be HBI measure since it is scalar type.
-list	Prints all token measure names and the value measure names associated with the token measure, which are registered in the domain specified by the -d option.
-remove <i>tokenMeasure=valueMeasure</i>	Removes the token measure with the specified token measure name and value measure. The token measure is unregistered from the domain specified by -d option. Unregistering the token measure has no side effect to the value measure that the token measure is associated with.

Batch Plug-In Tasks: execPluginTask.sh

Many RPAS applications can be used to extend the functionality of the RPAS Configuration Tools and the rpaInstall process. These extensions are provided by plug-ins included as part of the solution. Most plug-in extensions execute either within the RPAS Configuration Tools or within the rpaInstall process. However, some applications may include plug-in extensions that can be executed from a command prompt.

The execPluginTask.sh utility is the method by which these command line plug-in tasks can be performed. When execPluginTask.sh is executed, it loads a configuration, locates the desired plug-in extension, and executes that extension on the loaded configuration. Then, depending on the task being executed, the configuration may or may not be saved. Extensions that do not modify the configuration, for example, ones that generate resources based on the contents of the configuration, do not require saving the configuration at the conclusion of the automation.

Individual applications document the command line extensions they support, if any, along with information on how those extensions may be used.

execPluginTask.sh usage

In addition to the standard arguments described below, it is possible to pass any number of additional arguments to `execPluginTask.sh`. These additional arguments are passed to the extension being executed to allow control over its operation. Details on the arguments any given command line extension accept and their function are located in the application-specific documentation of the extension.

```
execPluginTask.sh plug-inName:taskName pathToConfiguration {plug-inTaskArguments}  
execPluginTask.sh -version  
execPluginTask.sh -help
```

Table 9–11 Arguments Used by the `execPluginTask.sh` Utility

Argument	Description
plug-inName	Indicates the name of the plug-in containing the extension to be executed.
taskName	Indicates the name of the extension to be executed.
pathToConfiguration	The path to the configuration <code>execPluginTask.sh</code> should load to execute the extension. As with the <code>rpasInstall</code> process, paths passed as arguments should be absolute paths and not relative paths.
plug-inArguments	Indicates additional arguments that are passed to the extension.
-version	Prints the version information for <code>execPluginTask.sh</code> .
-help	Prints the usage statement for <code>execPluginTask.sh</code> .

Informational Utilities

Many RPAS utilities can be used for finding information about many of the different components of a domain or domain data. The following utilities are solely for retrieving information and not for making any changes to a domain or data in a domain.

- [Retrieving Domain Information Using domaininfo](#)
- [Checking the Validity of a Domain Using checkDomain](#)
- [Determining RPAS Server Version Using rpassversion](#)
- [Listing Contents of a Database Using listDb](#)
- [Printing Data from Arrays Using printArray](#)
- [Printing Data from Measures Using printmeasure](#)

Retrieving Domain Information Using domaininfo

The domaininfo utility is used to provide miscellaneous details about a domain, such as the type of domain (simple, master, subdomain, or local), and the upgrade and version history of the domain.

domaininfo Usage

```
domaininfo -d pathToDomain [Command]
domaininfo -expectedversion
```

Note: Domain path (-d) is required for all the options except -expectedversion.

[Table 10-1](#) provides descriptions of the arguments used by the domaininfo utility.

Table 10-1 Arguments Used by the domaininfo Utility

Argument	Description
-d	Indicates the path to the domain. Required for all options except -expectedversion.
-domainversion	Displays the RPAS version of the specified domain.
-expectedversion	Displays the expected RPAS version of the domain that the utility expects to find.

Table 10–1 (Cont.) Arguments Used by the domaininfo Utility

Argument	Description
-type	Displays the type of the domain. Possible values are Simple, Global, and Sub. A Simple domain is a traditional, non-partitioned (non-global) domain. A Global domain is the central or master domain of a global domain environment. A Subdomain is one local domain in a global domain environment that can contain one or more partitions.
-apptag	Displays the application associated with domain.
-history	Displays the version history of the domain, specifically when the domain was upgraded to new versions of RPAS (patches or releases).
-listsubdomains	Displays a list of all the local domains in a global domain environment and indicates which positions at the partition level are in each local domain. This argument is only valid when run on a global domain.
-masterdomaininfo	Lists the master domain path and partition dimensions for subdomains.
-integrationStatus	States if the domain is integrated with an RPAS Data Mart. If the domain is integrated, it also prints out the domain identifier for the domain and the path to the RDM repository.
-all	Displays the domain version, expected domain version, domain type, associated application, history, subdomains, and master domain information (where applicable).
-domainsize	Displays the file size information for the domain.
-stringstats	Displays the number of strings of all given lengths that occur in the domain. The output returns a list by string length, for example: <pre> 0 758 1 21 2 8 3 69 ...</pre>
-stringvalues	Displays how many occurrences of each unique string exist in the domain. The output returns a list of each unique string, for example: <pre> 785 '' 1 '#copied <> #pasted' 2 '%' 1 '%1 <= passwd <= %2' 2 '%1 Window' ...</pre>
-arraydensity	Displays array btree density.
-arrayschemas	Displays how many arrays in the domain are formatted for each schema. The output returns a list like the following: <pre> Count of Array Schemas for Domain 'domain' Array schema 10: 371 Unknown schema:</pre>
-version	Displays the version of this utility.

Table 10–1 (Cont.) Arguments Used by the domaininfo Utility

Argument	Description
-subdomain <i>dim,pos</i>	Indicates to which local domain the specified position belongs. The position can be at or below the partition level.
-registrypaths	Displays all array paths to dimRegistry.
-showrelativepaths	When listing subdomains, indicates if paths are relative. Only relevant in combination with -listsubdomains or -all.
-terse	Removes the header information from the output, therefore returning only the requested information.

Checking the Validity of a Domain Using checkDomain

This utility is used to check the validity of an existing domain. Its primary purpose is to verify that a master domain matches its respective local domains and report all discrepancies to the administrator.

checkDomain Usage

```
checkDomain -d pathToDomain -type expectedType {-q}
```

Table 10–2 provides descriptions of the arguments used by the checkDomain utility.

Table 10–2 Arguments Used by the checkDomain Utility

Argument	Description
-d <i>pathToDomain</i>	Path to the domain that needs to be validated.
-type <i>expectedType</i>	Expected type of domain: simple, master, or subdomain.
-q	Quiet mode. Do not display progress messages.

When checkDomain is run on a simple domain the following two items get validated:

- The domain directory exists
- It is of type simple

If checkDomain is run on a global domain, it verifies the following:

- The global domain exists
- The global domain is of type master
- The global domain checks all of the subdomains for:
 - The subdomain directory exists and is of type sub
 - If the master domain and the subdomain have a repos directory
 - The measures, rules, rule groups, templates, and functions are the same in the global and subdomain

If it is run on a subdomain, it checks all of the items listed above for the global domain, but the validation is only performed between the global domain and the specified subdomain.

Determining RPAS Server Version Using rpasversion

Use the rpasversion utility to determine which version of the RPAS Server is running in a particular location.

rpasversion Usage

```
rpasversion -l pathToLibrary
```

Listing Contents of a Database Using listDb

Use listDb to list the basic information of all arrays contained in the databases provided.

listDb Usage

```
listDb pathToDb*
listDb -row -db pathToDb*
listDb -row -pageUsage -db pathToDb*
listDb -row -standardOptions -db pathToDb*
listDb -standardOptions -db pathToDb*
listDb -version
```

[Table 10–3](#) provides descriptions of the arguments used by the listDb utility.

Table 10–3 Arguments Used by the listDb Utility

Argument	Description
-db <i>pathToDb</i>	Specifies the database to list the contents.
-row	Lists array information in a row format.
-pageUsage	Shows btree page usage. Requires -row switch to be active.
-standardOptions	Lists only standard options.

Printing Data from Arrays Using printArray

Use printArray to print the contents of an array.

printArray Usage

```
printArray -array db.array -specs {-maxpos num}
printArray -array db.array {-cell "dim1:pos1,dim2:pos2,..."}
{-format "formatString"}
printArray -array db.array -slice "dim1:pos1,dim2:pos2,..."{-format
"formatString"} {-cellsperrow num} {-nposnames}
printArray -array db.array -allpopulatedcells {-format
"formatString"}{-cellsperrow num} {-nposnames}
```

[Table 10–4](#) provides descriptions of the arguments used by the printArray utility.

Table 10–4 Arguments Used by the printArray Utility

Argument	Description
-array <i>db.array</i>	Specifies the array to print. Specifies the full path to the database containing the array. Required for all commands except -version. <i>db</i> is a full or relative path to a database. Do not specify the .ary suffix. If no other commands are included, the array defaults to -allpopulatedcells with cells per row 1. The -allpopulatedcells command is still available, but now functions as a useful default action. The -nuposnames, -cellsprow, and -format parameters may still be specified when relying on the implicit -allpopulatedcells behavior.
-specs	Prints the specifications of the array and positions along each dimension.
-popcount	Outputs only the popcount of the specified array. Useful to shell script writers to get the popcount value into a shell script variable. For example, export POPCOUNT=`printArray -array hmaint.dim_year -popcount`
-cell <i>CELLSPEC</i>	Prints a specific cell value from the array. Must not contain spaces. Must identify a single of 1-D slice. Specify using the format "dim1:pos1,dim2:pos2,..."
-cellplain <i>CELLSPEC</i>	Outputs a specific cell value with no space padding. Useful for scripts when capturing cell values into shell variables. Must not contain spaces. Must identify a single of 1-D slice. Specify using the format "dim1:pos1,dim2:pos2,..."
-slice <i>CELLSPEC</i>	Prints a one-dimensional slice from the array. Must not contain spaces. Must identify a single of 1-D slice. Specify using the format "dim1:pos1,dim2:pos2,..."
-allpopulatedcells	Print all populated cells including the nvalue of the array.
-format " <i>fmtstr</i> "	If -format is specified, any cells with numeric values are interpreted as dates. <i>fmtstr</i> (formatString) determines how dates are interpreted and can include: ■ %Y: 4 digit year ■ %m: month number (01 to 12) ■ %d: numeric day of month (01 to 31) ■ %H: 24 hour clock (00 to 23) ■ %M: minute (00 to 59) ■ %S: seconds (00 to 61) ■ %s: milliseconds
-cellsprow <i>num</i>	For multi-cell output commands (-slice and --allpopulatedcells), indicates how many cells should be printed on each line.
-nuposnames	Suppresses the output of position names; only cell values are shown.

Printing Data from Measures Using printmeasure

Use the printMeasure utility to print measure information. The printmeasure utility cannot be used to print the data of a shared measure. If printmeasure is called with a

shared measure as an argument, it will exit with an error messages stating that the measure is shared.

printmeasure Usage

```
printmeasure -d domainPath {-wb wbName} {-m measure} [COMMAND]
```

Table 10–5 provides descriptions of the arguments used by the printmeasure utility.

Table 10–5 Arguments Used by the printmeasure Utility

Argument	Description
-d <i>pathToDomain</i>	Specifies the domain that contains the measure to print. Requires the -m parameter.
-m <i>measure</i>	Specifies the measure to print.
-wb <i>workbookName</i>	Specifies the workbook associated with the measure to print. If -wb is not used, the domain measure information is printed. Requires the -m parameter.
-list	Returns a list of all registered measures in the domain. This argument does not require -d domainPath.
-listHBIMeasures	In a global domain, returns a list of all measures registered at or above the partition dimension.
-specs	Returns the list of measure properties. Requires the -m parameter. If the measure in question is shared through an RPAS Data Mart, the fact and fact group of the shared measure will also be returned.
-listDataIntersections	Returns the base intersection of the measure.
-slice CELLSPEC	Prints a one-dimensional slice from the array. Must not contain spaces. Must identify a single of 1-D slice. Specify using the format "dim1:pos1,dim2:pos2,..." Cannot be used for a shared measure in the RDM.
-cell CELLSPEC	Prints a specific cell value from the array. Must not contain spaces. Must identify a single of 1-D slice. Specify using the format "dim1:pos1,dim2:pos2,..." Cannot be used for a shared measure in the RDM.
-allpopulatedcells	Print all populated cells including the nvalue of the array. Cannot be used for a shared measure in the RDM.

Internationalization

Internationalization is the process of creating software that can be translated more easily. Changes to the code are not specific to any particular market. This section describes configuration settings and features of the software that ensure that the base application can handle multiple languages.

Oracle Retail applications have been internationalized to support multiple languages.

Translation

Translation is the process of interpreting and adapting text from one language into another. Although the code itself is not translated, components of the application that are translated include the following:

- Graphical user interface (GUI)
- Error messages

The following components are not translated:

- Documentation (online help, release notes, installation guide, user guide, operations guide)
- Batch programs and messages
- Log files
- Configuration tools
- Reports
- Demonstration data
- Training materials

The user interface has been translated into the following languages:

Note: In [Table 11-1](#), the language identifier is used for position labels. For more information, see ["Position Label Translation"](#). The Windows Language ID is in the foundation.ini file. For more information, see ["Translation Administration"](#).

Table 11-1 *Supported Languages with Language Identifiers*

Language	Language Identifier	Windows Language ID
Chinese (Simplified)	CHINESE_SIMPLIFIED	2052

Table 11–1 (Cont.) Supported Languages with Language Identifiers

Language	Language Identifier	Windows Language ID
Chinese (Traditional)	CHINESE_TRADITIONAL	1028
Croatian	CROATIAN	26
Dutch	DUTCH	19
English	ENGLISH	9
French	FRENCH	12
German	GERMAN	7
Greek	GREEK	8
Hungarian	HUNGARIAN	14
Italian	ITALIAN	16
Japanese	JAPANESE	17
Korean	KOREAN	18
Polish	POLISH	21
Portuguese (Brazilian)	PORTUGUESE	22
Russian	RUSSIAN	25
Spanish	SPANISH	10
Swedish	SWEDISH	29
Turkish	TURKISH	31

Translation Administration

Note: For information on the translation of position labels, see [Position Label Translation](#).

Every product, location, and calendar position can be presented in multiple languages, as can messages presented through the client. However, before translated strings can be viewed in the client, the following processes must be followed to set up the environment to support multiple languages.

1. Build the domains with the Multi-Language setting enabled in the Configuration Tool properties.
2. By default, the language of the RPAS Client is determined by the language of the client-side operating system. You can override this behavior by setting the Language entry in the Options section of the foundation.ini file:

```
[Options]
Language=LANGID
```

LANGID represents a value in the Windows Language ID column of [Table 11–1](#).

Or, you can set the following option:

```
[Options]
ResourceDLL=LANG.dll
```

LANG.dll represents a value in the Language Identifier column of [Table 11–1](#).

If the RPAS Client cannot find the library for the language you specify, it defaults to English.

3. Locate the foundation.ini file. It is found in one of the following places. When RPAS searches for the file, it searches in the following order:

- a. The user's Application Data folder:

- For Microsoft Windows XP:

C:\Documents and Settings*<user>*\Application Data\Oracle Predictive Solutions

- For Microsoft Windows 7:

C:\Users*<user>*\AppData\Local\Oracle Predictive Solutions

- b. The same folder as the currently loaded foundation.exe file.

- c. The current working folder, which can be one of the following:

- The debugger's setting, if started by debugger
- The same folder as the foundation.exe file, if started manually
- ./winnt/system32, if started from automation

- d. C:\Windows

If the file is not found, it is created in the Application Data folder described in Step a.

4. Place the solution and RPAS .ovr files for the language selected into the 'input' directory of your domain. The .ovr files are found in RPAS_HOME/translations. The input directory should also contain the lngs.dat file that will be loaded prior to loading the .ovr files. If not packaged with the language .ovr files, the current lngs.dat file will be in the "input/processed" folder of the domain. Move it back to the input folder and remove the timestamp extension. Then reload the hierarchy using loadHier -d /path -load lngs. Make sure that the language you are loading exists in the lngs.dat file before loading it. Load each .ovr file within the input directory by using the loadmeasure utility (for example, loadmeasure -d /path -measure r_msglabel).

Note:

- The files are removed from your input folder when they are loaded. Therefore, any file that is left in this directory still needs to be loaded.
- For Japanese, Korean, Chinese Simplified or Traditional, and Russian, you must reboot your machine in each of these languages to properly see text if you changed the Windows regional options to one of these languages, as described in the next step.
- The translated .ovr files are released in UTF-8 encoding. Any changes to these files must be made with a UTF-8 capable editor and saved without BOM (Byte Order Mark).

The .ovr files consist of three data columns, as shown below.

```
Gub_mean_      SPANISH      Media máxima
```

The data is populated on a specific column and has definite width. When making any change to this data, you must ensure that the new data abides by the predefined width and position for that data for the file.

5. Set the default language to match.
 - On Windows XP, in Regional Options, set the **Standards and formats** list to the desired language. Set **Location** to the appropriate country. Click the **Advanced** tab and set **Language for non-Unicode programs** to the appropriate language.
 - On Windows 7, access the Control Panel and click **Clock, Language, and Region**, then click **Region and Language**. In the Region and Language dialog box, set the **Format** field to the desired language. Click the **Location** tab and set the **Current Location** field to the desired location. Click the **Administrative** tab and click **Change System Locale** button to set the desired language.

Translation Administration Workbook

The Translation Administration workbook contains worksheets for translating text used in measure labels, workbook template names, template group names, user group labels, and general areas (for instance, wizard instructions, and error messages).

Note: RPAS and solution-specific messages to the user should not be modified. If changes are made to these messages they may be overwritten when patching occurs.

Hierarchy Labels Worksheet

The Hierarchy Labels worksheet allows the user to view and edit the translations of hierarchy labels. Translations are supported for each of the system's allowable alternative languages.

Dimension Labels Worksheet

The Dimension Labels worksheet allows the user to view and edit the translations of dimension labels. Translations are supported for each of the system's allowable alternative languages.

Workbook Template Group Labels Worksheet

The Template Group Translations worksheet allows the user to view and edit the translations of template group names. Translations are supported for each of the system's allowable alternative languages. Translations in this worksheet affect the labels on the tabs that appear in the File - New dialog (for example (in English), Administration, Analysis, and Predict).

Workbook Template Labels Worksheet

The Template Translation worksheet allows the user to view and edit the translations of workbook template names. Translations are supported for each of the system's allowable alternative languages.

Measure Labels Worksheet

The Measure Translations worksheet allows the user to view and edit the translations of measure labels. Translations are supported for each of the system's allowable alternative languages.

Measure Descriptions Worksheet

The Measure Descriptions worksheet allows the user to view and edit the translations of measure descriptions. Translations are supported for each of the system's allowable alternative languages.

User Group Labels Worksheet

The User Group Translations worksheet allows the user to view and edit the translations of user group labels. Translations are supported for each of the system's allowable alternative languages. The list of user groups includes the Administration, Default, and Internal user groups, plus any other user group names set up by the system administrator. For products in the Oracle Retail Predictive Planning Suite, the list of user groups also includes the various planning roles.

Message Labels Worksheet

The Message Labels worksheet allows the user to view and edit the translations of messages displayed to users in the RPAS Client. Translations are supported for each of the system's allowable alternative languages.

RGRP Labels Worksheet

The RGRP Labels worksheet allows the user to view and edit the translations of rule group labels displayed to users in the RPAS Client. Translations are supported for each of the system's allowable alternative languages.

Commit as Soon as Possible

Commit As Soon As Possible (Commit ASAP) allows users to schedule the commit process for workbook data so that it executes as soon as all the system resources are available. Commit ASAP is an option in the File menu of the RPAS Client. Procedures for using Commit ASAP are provided in the *Oracle Retail Predictive Application Server User Guide for the Classic Client*.

Commit ASAP takes a copy of the data to be committed. Unlike Commit Later, which adds a workbook commit process to a queue that is run in batch, the data that is eventually committed is the data that was present at the time the commit instruction was issued. With Commit Later, if the user makes further changes to the workbook and saves that workbook before the batch commit process is run, those changes will also get committed.

This chapter contains the following sections:

- [Using Commit ASAP](#)
- [Managing the Workbook Queue Using showWorkbookQueues](#)
- [Commit ASAP Settings Using configCommitAsap](#)

Using Commit ASAP

After attempting to commit a workbook using Commit ASAP (File\Commit ASAP), the user sees a message in the client that the workbook has been scheduled for a commit. The user can continue working. The system then tries to commit the workbook as soon as it can, taking into account any other scheduled commits. If the commit cannot be done prior to the domain's Commit ASAP deadline, it will be canceled and listed as failed.

There are four states for commit processes to be added to the Commit ASAP queue.

- Pending: The commit process is queued up to take place at some point in the future.
- Committing: The workbook is currently being committed.
- Success: The commit succeeded.
- Failed: The commit failed.

The status of each Commit ASAP process can be viewed by accessing a dialog window called Commit Status from the File menu. This dialog window displays all of the Commit ASAP processes with their respective status for all processes that have not been purged (see below). This dialog can be used to sort the tasks based on any of the columns.

Entries can be filtered in a variety of ways. If the check box **All Users** is not checked, users will only be able to view their own entries. If it is checked, any user will see the entries for all users. The check boxes in the **Status To Display** group allow a user to filter the output in order to see only the processes with the specified statuses. The window can be updated by using the **Refresh** button. The dialog remembers the settings based on the last use.

Note:

- If a user attempts to commit a workbook using the Commit ASAP option, and the workbook already has a process in the queue, the original processes is removed from the queue. That means that only one Commit ASAP can be pending in the queue for a given workbook/user/template name combination at any given time.
 - Data within the workbooks can be committed without saving the workbook. When users click Commit, and if the workbook has not been saved by then, RPAS Classic Client provides an option to Save, Commit, or perform both actions separately on the workbook.
-

Managing the Workbook Queue Using showWorkbookQueues

The RPAS utility `showWorkbookQueues` is used for viewing the status of Commit ASAP processes and for purging entries in the Commit ASAP status window.

The purge option requires a date before which entries will be removed, as well as specification for which entries to remove: succeeded, failed, or both.

showWorkbookQueue Usage

```
showWorkbookQueues -version
showWorkbookQueues -d domainPath -show
[all|pending|waiting|working|success|failed]*
showWorkbookQueues -d domainPath -purge date [success | failed]*
```

Table 12–1 provides descriptions of the arguments used by the `showWorkbookQueues` utility.

Table 12–1 *showWorkbookQueues Arguments*

Argument	Description
-version	Prints the RPAS version, revision, and build information of the utility.
-d domainPath	Specifies the path to the domain.
-show	Lists the contents of the queue in the order in which the parameter is specified. Possible values: all, pending, waiting, working, success, and failed.
all	Used with the -show parameter. This lists all of the workbooks in all statuses.
pending	Used with the -show parameter. This lists all workbooks that are waiting to be committed.
waiting	For Oracle Retail development use only.
success	Used with the -show parameter. This lists all workbooks that have been successfully committed.

Table 12–1 (Cont.) showWorkbookQueues Arguments

Argument	Description
failed	Used with the -show parameter. This lists all workbooks that did not successfully commit.
-purge <i>date</i>	Purges entries in the Commit ASAP status window. Entries before the date provided will be removed. The date should be a string of the following DateTime format: YYYYMMDDHHmm For example "200406071529" equals June 7, 2004 3:29 PM. An administrator must select to purge commit processes that either succeeded or failed.

Commit ASAP Settings Using configCommitAsap

There are two settings for Commit ASAP that are managed by an administrator. Both are set using the utility configCommitAsap.

- Maximum number of simultaneous commit processes (property MaxProcesses, default value is 4).
- Deadline by which all pending processes must be completed, after which they will be cancelled and marked as failed.

This deadline will likely be used by administrators before beginning the nightly batch processes (property deadline, default value is 00:01 [meaning 12:01 AM], in 24-hour time).

A commit process that starts before the deadline is reached will be processed. Commit requests that were in the queue before the deadline that did not get processed will be cancelled and marked as failed. Commit requests added to the queue after the deadline will use the deadline of the following day.

configCommitAsap Usage

```
configCommitAsap -d pathToDomain [-maxProcs numProcs]
[-deadline time] [-display]
```

Table 12–2 provides descriptions of the arguments used by the configCommitAsap utility.

Table 12–2 configCommitAsap Arguments

Argument	Description
version	Prints the RPAS version, revision, and build information of the utility.
-maxProcs <i>numProcs</i>	Sets the maximum number of concurrent commit processes where numProcs is an integer greater than 0. Workbooks can be committed in parallel if they do not require access to the same measure databases. If they do share databases, they will be committed sequentially.

Table 12–2 (Cont.) configCommitAsap Arguments

Argument	Description
-deadline <i>time</i>	Indicates the time of the day when all outstanding commit ASAP operations will time out. If a commit ASAP operation is submitted after this time, it will not timeout until the deadline time on the next day. This string must have the following format: HH:MM For example "13:30" refers to 1:30 PM.
-display	Displays the current commit ASAP settings.
-loglevel <i>level</i>	Sets the logger verbosity level. Possible values: all, profile, information, warning, error, or none.
-noheader	Disables timestamp header use.

Logging and Technical Information

A log file is available in the Commit ASAP directory that should be checked if a user reports an error with a Commit ASAP submission. The file is named `rpasServer.log` and is in the following directory: `<Path to domain>/commitAsapQueue`.

Another log file is generated for each Commit ASAP process and stored in a user's directory (`users/<userid>/asapLogs`). The format of the log file name is "`orig_<original workbook name>asap_<temporary workbook name>.log`." RPAS creates a temporary workbook in this process to capture a snapshot of the data that needs to be committed. Temporary workbooks are never viewed by a user. An administrator can use this log if something does not properly commit.

Note: These "snapshot" workbooks cannot be viewed or used in the RPAS Client.

An example of this log file is `orig_t1_asap_t5` where "t1" is the name of the original workbook and "t5" is the name of the snapshot workbook.

The following directories are used to store the copies of the workbook as they are processed through the system:

- Pending directory: Contains one file per submitted Commit ASAP that has not yet been processed. These files are, in general, binary and cannot be easily read.
- Working directory: Contains one file per submitted Commit ASAP that is currently in the commit process.
- Success directory: Contains one file per submitted Commit ASAP that has successfully completed its commit process.
- Failed directory: Contains one file per submitted Commit ASAP that either had a failure during its commit process or could not be committed prior to the deadline.
- Unknown directory: If the Commit ASAP process detects a corrupted queue file, a message gets logged and the file gets moved into the unknown directory.

Batch Processes and RPAS Utilities

Included with an RPAS installation is a collection of stand-alone executables and scripts that are used for a variety of operations. RPAS utilities are run directly against a domain. In a global domain environment, most utilities can only be run on the master domain. RPAS utilities can be categorized into the following groupings:

- **Hierarchy management:** The loading and refreshing of hierarchies, and the process of updating the data structures in the domain to reflect hierarchy changes
- **Measure data:** Utilities for loading, exporting, and moving data within and between domains
- **Miscellaneous:** A variety of utilities for performing certain procedures in batch and for setting a number of parameters on an environment or a domain
- **Information RPAS utilities:** A variety of utilities that retrieve information about a domain, data, the RPAS Server code, or an object used by the server

This chapter contains the following sections:

- [CSV File Format](#)
- [RPAS Utilities Logging Options](#)
- [Using Shell Scripts to Run Batch Processes](#)
- [Common Information and Parameters for RPAS Utilities](#)
- [RPAS Intraday Enabler](#)

CSV File Format

For those utilities that use a comma-separated value (CSV) file, the following formatting applies for any commas or double quotation marks in the data:

- If the data does not contain any commas or double quotation marks, it does not need any special formatting.
- If the data contains a comma, the string must be enclosed between opening and closing double quotation marks.
- If the data contains quotation marks, the string must be enclosed between opening and closing double quotation marks and any embedded quotation marks must be paired.

[Table 13-1](#) shows examples of the formatting.

Table 13–1 CVS File Format

Data	Formatted Data
Item 001	Item 001
Item 001, Soda	"Item 001, Soda"
"Large Screen" TV	""Large Screen"" TV"
Item 002, "Generic Brand" Cereal	"Item 002, ""Generic Brand"" Cereal"

RPAS Utilities Logging Options

RPAS has a number of applications used to control or process data. Currently there are no unified methods for logging output, controlling the level of logging, or directing logging to a particular file. Instead each utility has its own methods, although many are similar. The current behavior for each utility follows.

Log Levels

The standard log levels are controlled by the `-loglevel` option. Not all programs use these levels, but most do. The default logging level is Warning, which means that any log messages that are specified as a warning or higher will be output.

- All: Forces all log levels to be output
- Profile: Performance profiling information
- Audit: User-specific domain and workbook activities. These activities include the following:
 - Workbook build, calculation, save, commit, and custom menu operations
 - User login and logout to the domain
- Information: General status messages that are not problematic. Outputs status and progress of the operation, in addition to the error and warning messages.
- Warning: Messages indicating a potential problem, but not one that is fatal. Outputs warning messages, in addition to error messages.
- Error: Messages relating to a fatal problem. Outputs only error messages.
- None: No messages. No output is provided if the utility successfully executes.

Each of the lines that contain the above types of feedback is normally preceded with a code that indicates what type of information is being output. Each code should have an angle bracket (<) in front of it.

- E indicates that the message is an error.
- W indicates that the message is a warning.
- I indicates that the message is informational.
- U indicates that the message is audit-relevant information.
- P indicates that the message is a performance profile.

Note: Audit information related to workbook activities is recorded in `rpas.log` under each user's working directory. Information related to domain activities, such as user sign-on and sign-off, is recorded in `DomainDaemon.log`.

Utilities with Standard Logging

A number of utilities allow for the `-loglevel` option to control which messages are output to the screen. There is no way to log to a file directly. [Table 13–2](#) displays the utilities that can use the `-loglevel` option.

Table 13–2 Utilities That Can Use the `-loglevel` Option

alertmgr	regfunction
checkDomain	regmeasattr
configCommitAsap	regtemplate
createdb	regTokenMeasure
createGlobalDomain	reguserdim
dattrmgr	rpasverison
dbdiff	rtkappcnfgmeas
dimensionMgr	showWorkbookQueues
domaininfo	syncNAValue
ldrule	updateArray
listDb	updatestyles
mapData	upgradeDomain
moveDomain	usermgr
printArray	wbmgr
printMeasure	

Scripts

Shell scripts cannot use standard logging, but may execute the following programs that use it:

convertDomain

All output to the screen.

createRpasDomain

The `-v` option controls the type of messages sent to the screen.

Utilities with Multi-Process Logging

Some utilities are based on the multi-processes domain utility framework. These utilities send messages to the screen and a log file `master.log`. Any child processes output messages to a log file in the `domain/output` directory named `subdomain0000.log` where the number indicates the subdomain being processed. This directory will contain all log files created during the run of that utility. This change has been updated so that the controlling process logs to the screen as well as to a file in that directory. The newly created directory name is formatted as `APPNAMEYYYYMMDDHHMMIbXX`, where `APPNAME` is the utility name, `YYYY` is the year, `MM` is the month, `DD` is the day, `HH` is the hour, `MM` is the minute. The character `b` and `XX` (two digits) are used to make the directory name unique. The framework attempts to limit the number of directories created for any single utility to eight. The parameter `-loglevel` can be used to control the type of messages sent to the screen and log file.

These utilities are as follows:

- exportData
- loadmeasure
- loadHier
- exportMeasure
- reconfigGlobalDomainPartitions
- wbbatch
- reindexDomain

domainprop

The domainprop utility only provides logging to the screen.

hierarchyMgr

The hierarchyMgr utility only provides logging to the screen.

configCommitAsap

This utility should be started from the RpasDbServer application when the client requests a workbook to be committed.

Utilities with Special Logging

These utilities may use standard logging with additional features or may use entirely different logging methods.

DomainDaemon

The DomainDaemon uses standard logging. It logs output to a file (see below). The file is created either in the current working directory or in the directory specified by the RPAS_LOG_PATH environment variable.

The file name depends on the RPAS_LOG_BACKUPS environment variable. If it is set to 1 or greater, then:

The log file name is Daemon_Dyyyymmddhhmmmbxx.log where yyyy is the current year, mm is the current month, dd is the current day, hh is the current hour, mm is the current minute. The character b and xx (a two-digit number) are used to make the file name unique.

The number of these log files is limited to the number provided in the environmental variable RPAS_LOG_BACKUPS.

Otherwise, the log file name is Daemon.log. Any existing log file is renamed to Daemon.old.

At midnight, the current log file is closed and a new one opened, named as described above.

RpasDbServer

This should only be started from the DomainDaemon as a part of a client request to start an RPAS session. The logging level is controlled by the client's RPAS_LOG_LEVEL environment variable. If it is not set, then it defaults to logging messages at the warning level.

This utility creates log files in the domain/users/client directory, where the domain is the current domain path and the client is the current client. The actual file name used

is either `rpas Dyyyymmddhhmmmbxx.log` or `rpas.log`, based on the environmental variable `RPAS_LOG_BACKUPS` (c.f., `DomainDaemon`, above).

loadHier

The `loadHier` utility uses standard logging. This utility performs part of its processing in child processes. `loadHier` provides a list of all hierarchy positions that have been changed since the previous hierarchy load. The resulting directory name is:

```
<utility><YYMMDDHHMISS><pXXXXX><bYY>
```

where `utility` is the name of the program (for example, `-loadmeasure`), followed by a time and date stamp, then the process id (`pXXXX`), and then the character `b` and a 2-digit number to avoid conflicts (`bYY`).

If any problems occur when loading specific records that belong to the partition hierarchy, they are reported in the format as shown below. Note that the record is completely reproduced in this error report in the log.

```
<E 2008Jul02 12:04:52.196> Could not find position '90000044' in line number 3:
'2001052090000044 1000 7 '. Skipping!
```

Problems with records along non-partitioned hierarchies are reported as shown in the following example:

```
<I 2008Jul02 12:04:55.482> MeasureLoader::loadDataFromFile() Loading '.ovr' file
'/vol.nas/u09/rpasqc/qc_testing/aix/1208rc2_test/RDF_12/ldom1/input/psal.ovr'
<D 2008Jul02 12:04:55.514> Error on line 1: '2001031110000044 STR1000 8 '
.Position name: STR_STR1000 not found.
<D 2008Jul02 12:04:55.514> Error on line 2: '2011041510000044 1000 9 ' .Position
name: DAY20110415 not found.
<D 2008Jul02 12:04:55.964> 2 lines had problematic data.
```

locked

Messages are sent only to the screen.

mace

The `mace` utility uses standard logging. The `-debugRuleEngine` option logs some messages to the file `mace.log` in the current working directory.

reconfigGlobalDomainPartitions

Uses standard logging. The `loadHier` process may be spawned as a child process. See "[loadHier](#)" for additional details. When the `loadHier` utility is started as a child process it remaps the screen output of to the log file `loadHier.log` contained in the current working directory.

renamePositions

Uses standard logging. The `-log` option overwrites the default log file name of `hierName` and `Rename.log` in the current working directory. The `-loglevel` parameter does not control the types of messages written to this log file.

regmeasureServer

This application should only be started from the RPAS libraries to process measure registration and de-registration. Each process creates a log file in a newly created directory in the domain output directory. The newly created directory name is formatted as `regServerYYYYMMDDHHMMbXX`, where `YYYY` is the year, `MM` is the month, `DD` is the day, `HH` is the hour, `MI` is the minute. The character `b` and `XX` (two

digits) are used to make the directory name unique. The RPAS libraries attempt to create at most eight directories for any single application.

Utilities and Database Locks

On UNIX and Linux platforms, before running any RPAS program, it is important to run "umask 0". Otherwise the RPAS databases may be permanently locked out to access by other users or groups. This applies to any program that accesses any RPAS database, including, but not limited to:

- convertDomain
- createRpasDomain
- DomainDaemon
- domaininfo
- exportMeasure
- ldrule
- listDb
- loadhier
- loadmeasure
- mace
- printArray
- printMeasure
- printMeasureInfo
- regmeasure
- rpasInstall
- RPASODBCAgent
- RPASODBCDataService
- usermgr

If database lock files (*.db.lck) have been created with limited permissions that block other users' access, a workaround is to run the following command when logged in as the lock-file-owning user:

```
$ find {domainPath} -user {owner} -name \*.lck -exec rm -f {} ';' 
```

where {domainPath} is the path to the domain and {owner} is the UNIX user that owns the lock files.

Using Shell Scripts to Run Batch Processes

Batch processes should be written using scripts that call the RPAS 11 binaries found in the \$RPAS_HOME/bin/ directory. Any log files generated by scripts are in the [DOM]/scripts/err/ directory. Examples of tools include Korn Shell, Python, and Perl.

A Sample Shell Script

The following sample shell script loads the product and location hierarchies into a domain. It is assumed that this script is invoked from the [DOM]/scripts/ directory.

```

1  #!/bin/ksh
2  loadHier -d .. -load prod > ./err/loadhier.prod.log
3  loadHier -d .. -load loc >> ./err/loadhier.loc.log

```

Line 1 defines the shell that executes the script. In this example, it is defined to be the Korn shell. Therefore, this script is always executed from the Korn shell, even if the user's shell is different.

Lines 2 and 3 call the loadHier utility to load the latest product and location hierarchy information. Depending on the batch process to be performed by the shell script, lines 2 and 3 can be replaced by one or more lines to call one or more RPAS utilities.

Common Information and Parameters for RPAS Utilities

A number of standard arguments are available for most RPAS utilities. Check the usage of a specific utility to verify whether or not it is available.

Table 13-3 *Standard Arguments for RPAS Utilities*

Argument	Description
-version	Obtains the version information of the utility (for instance, RPAS 13.2.0). It does not require -d domainPath.
-d <i>pathtodomain</i>	Specifies the path to the domain against which the utility will run or from which data will be used.
-loglevel	See "RPAS Utilities Logging Options" for more information.
-n	Certain utilities contain this parameter to perform a dry run. Use this option to see what would change without making actual changes to the system or data. See the usage of a specific utility to see whether this option is applicable.
-noheader	Disables the use of a timestamp in the header of the log file.
-help -? -usage	Any of these arguments output the utility information and syntax to the terminal window. This can also be accomplished by running the utility with no arguments.

Logger verbosity levels determine how much information is generated on the terminal for a given utility. An administrator can set these levels for each RPAS utility. The available logger verbosity levels are as follows:

- none: No output if the utility successfully executes
- error: Outputs only error messages
- warning: Outputs warnings in addition to error messages
- information: Outputs status and progress of the operation in addition to the error and warning messages
- all: Outputs all available information generated by the utility, including error, warning, and informational messages

Each line, which contains the above type of feedback, is normally preceded with a code that indicates what type of information is being output. Each code should have an angle bracket (<) in front of it. E indicates the message is an error. W indicates the message is a warning. I indicates the message is informational.

Configuration Tools Log Files

For the RPAS Configuration Tools, information is logged in the files stderr.txt and stdout.txt, which are located in the bin subdirectory of the Tools directory. If a problem with the configuration tools is encountered, send these two files to Oracle Retail Customer Care, along with a description of the problem.

RPAS Intraday Enabler

The RPAS Intraday Enabler (ride) functionality enables batch operations to be run over an RPAS domain while users are accessing workbooks and completing workbook operations.

This functionality enables batch operations to be executed over a domain, but does not prevent users from accessing other components that do not affect or interfere with the batch operations. The running of an exclusive batch process does not cause any pre-existing workbook operations that require domain access to fail or terminate. Users in domains that are not part of the exclusive process are not affected in any way.

In domains that have been locked by an exclusive batch process, the users are still able to perform operations that only require access to the workbook. The operations include the following:

- Workbook Edits
- Workbook Calculations
- Workbook Saves
- Workbook Opens
- Workbook Navigation
- Commit ASAP Entry
- Commit Later Entry

Workbooks entered into the commit ASAP queue while an exclusive lock is in place are processed after the exclusive process is complete.

In these same domains, users cannot perform operations that require access to data within the domain. The access can be either read or write. The operations that are prevented include the following:

- Workbook Build
- Workbook Refresh
- Workbook Custom Menu (unless configured as intraday-concurrent)
- Workbook Commit Now/Later
- Insert Measure
- Dynamic Position Maintenance (DPM)

When users try to access one of these operations after the exclusive lock is obtained, they see a message stating that an exclusive process is running. A default message is provided, which can be replaced with a message as part of the call to the ride utility.

When users are working with a workbook in the master domain, a lock is required on the master domain and all local domains that are needed for the operation. The workbook operations in this domain are blocked when the master or local domain data is accessed by the ride process. See ["Scenarios"](#) for more details on domain access during different types of ride processes.

Configuration functionality is provided so that a custom menu can be flagged to run concurrently with a ride process (intraday-concurrent).

Note: See the *RPAS Configuration Tools User Guide* for details on how to configure a custom menu to be intra-day-concurrent. A custom menu that is configured to run concurrently with a ride process should only access workbook data, run a script that uses the ride utility, and run commits using the commit ASAP functionality. Custom menus that update or read directly from the domain should not be configured as intra-day-concurrent as this can conflict with the ride process.

In order for a batch job to run over a domain without interference from an online activity, exclusive domain access must be granted to the job that is running. This is achieved by creating a domain access control using a dual-lock control. The domain access control manages the lock request from workbooks and ride processes.

Figure 13–1 Domain Access Control




Figure 13–1 shows the process control that is in place with the locking schema. In this case, an administrator requests exclusive access to a domain in order to run a batch job. This requires an exclusive lock on the domain so that the job can run. After the lock is received, no other workbook operations can obtain write access to the domain until the process is complete. If the exclusive lock cannot be obtained, the process should time out and the administrator notified based on the output of the ride utility. When the ride utility times out, the domainStatus utility is automatically run to provide details of the user workbook operations that are blocking the ride process. See the "[Domain Lock Status Using domainStatus](#)" section on for details on that procedure.

ride Usage

```
ride -d domain -process pname|-script sname -args args {-message messageString}
{-timeout minutes} {-wait minutes} {-partitions pos1,pos2,...} {-masterInBatch}
```

Table 13–4 provides descriptions of the arguments used by the ride utility.

Table 13–4 *ride Utility Arguments*

Argument	Description
-d <i>domain</i>	Refers to a simple or master domain. When the master domain is specified, all local domains in the global domain environment, as well as the master domain, are locked.
-process <i>pname</i>	The name of the process to execute. This parameter cannot be used with the -script parameter.
-script <i>sname</i>	The name of the script to execute. This parameter cannot be used with the -process parameter.
-args <i>args</i>	Process arguments passed to the script or process to be executed. The -args parameter must be the last parameter or switch for this application. All parameters or switches after the -args parameter are passed on to the process or script to be started.
-message <i>messageString</i>	This optional override message is presented to the user who is trying to perform an operation blocked by an intra-day batch process. A default message is provided to the user if this argument is not provided.
-timeout <i>minutes</i>	The utility will time out if it cannot get access to the domains during this time. By default, there is no timeout. The timeout starts when the control utility is executed.
-wait <i>minutes</i>	The time to wait before starting the process or script. Even if domain access is granted, the process does not start until the end of wait time. The clock starts when the control utility is executed. The default is 0.
-partitions <i>pos1,pos2,...</i>	Partition positions (such as dept1, dept2, and so on) that determine the local domains that are accessed by the process or script.
-masterInBatch	When running over a global domain environment, the master domain is accessed by the process or script in addition to any local domains selected.

Note: If neither -partitions nor -masterInBatch are provided on the command line, the entire domain will be processed when running over a global domain environment (that is, all subdomains and the master).

Scenarios

This section provides several scenarios that are possible with the ride utility. The scenarios explain the domain access based on how the ride utility is executed. The tables indicate whether the specific operations are blocked or allowed.

Scenario 1

This scenario describes running ride and specifying only the master domain. This locks all domains.

Usage: ride -d master -script script.ksh

Table 13–5 Running the ride Utility Specifying Only the Master Domain

Workbook Operation		Master Domain	Local Domain 1	Local Domain 2	Local Domain 3
Build	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Blocked	Blocked
Refresh	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Blocked	Blocked
Commit Now	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Blocked	Blocked
Custom Menu (not ride concurrent)	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Blocked	Blocked
Insert Measure	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Blocked	Blocked
DPM		NA	Blocked	Blocked	Blocked

Scenario 2

This scenario describes running ride and specifying only one local domain (local domain with partition position 100).

Usage: ride -d master -script script.ksh -partitions 100

Table 13–6 Running the ride Utility Specifying One Local Domain

Workbook Operation		Master Domain	Local Domain 1	Local Domain 2	Local Domain 3
Build	HBI Measures	Blocked	Blocked	Allowed	Allowed
	No HBI Measures	Blocked	Blocked	Allowed	Allowed
Refresh	HBI Measures	Blocked	Blocked	Allowed	Allowed
	No HBI Measures	Blocked	Blocked	Allowed	Allowed
Commit Now	HBI Measures	Blocked	Blocked	Allowed	Allowed
	No HBI Measures	Blocked	Blocked	Allowed	Allowed
Custom Menu (not ride concurrent)	HBI Measures	Blocked	Blocked	Allowed	Allowed
	No HBI Measures	Blocked	Blocked	Allowed	Allowed
Insert Measure	HBI Measures	Blocked	Blocked	Allowed	Allowed
	No HBI Measures	Blocked	Blocked	Allowed	Allowed
DPM		NA	Blocked	Allowed	Allowed

Scenario 3

This scenario describes running ride and specifying only one local domain (local domain with partition position 100) and the masterInBatch option.

Usage: ride -d master -script script.ksh -partitions 100 -masterInBatch

Table 13–7 Running the ride Utility Specifying One Local Domain and masterInBatch

Workbook Operation		Master Domain	Local Domain 1	Local Domain 2	Local Domain 3
Build	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Allowed	Allowed
Refresh	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Allowed	Allowed
Commit Now	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Allowed	Allowed
Custom Menu (not ride concurrent)	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Allowed	Allowed
Insert Measure	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Allowed	Allowed
DPM		NA	Blocked	Blocked	Blocked

Scenario 4

This scenario describes running ride and specifying two local domains (local domain 1 with partition position 100 and local domain 2 with partition position 200).

Usage: ride -d master -script script.ksh -partitions 100,200

Table 13–8 Running the ride Utility Specifying Two Local Domains

Workbook Operation		Master Domain	Local Domain 1	Local Domain 2	Local Domain 3
Build	HBI Measures	Blocked	Blocked	Blocked	Allowed
	No HBI Measures	Blocked	Blocked	Blocked	Allowed
Refresh	HBI Measures	Blocked	Blocked	Blocked	Allowed
	No HBI Measures	Blocked	Blocked	Blocked	Allowed
Commit Now	HBI Measures	Blocked	Blocked	Blocked	Allowed
	No HBI Measures	Blocked	Blocked	Blocked	Allowed
Custom Menu (not ride concurrent)	HBI Measures	Blocked	Blocked	Blocked	Allowed
	No HBI Measures	Blocked	Blocked	Blocked	Allowed
Insert Measure	HBI Measures	Blocked	Blocked	Blocked	Allowed
	No HBI Measures	Blocked	Blocked	Blocked	Allowed
DPM		NA	Blocked	Blocked	Allowed

Scenario 5

This scenario describes running ride and specifying two local domains (local domain 1 with partition position 100 and local domain 2 with partition position 200) and the masterInBatch option.

Usage: ride -d master -script script.ksh -partitions 100,200 -masterInBatch

Table 13–9 Running the ride Utility Specifying Two Local Domains and the masterInBatch Option

Workbook Operation		Master Domain	Local Domain 1	Local Domain 2	Local Domain 3
Build	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Blocked	Allowed
Refresh	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Blocked	Allowed
Commit Now	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Blocked	Allowed
Custom Menu (not ride concurrent)	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Blocked	Allowed
Insert Measure	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Blocked	Allowed
DPM		NA	Blocked	Blocked	Blocked

Scenario 6

This scenario describes running ride to lock the master domain.

Usage: ride -d master -script script.ksh -masterInBatch

Table 13–10 Running the ride Utility to Lock the Master Domain

Workbook Operation		Master Domain	Local Domain 1	Local Domain 2	Local Domain 3
Build	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Allowed	Allowed	Allowed
Refresh	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Allowed	Allowed	Allowed
Commit Now	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Allowed	Allowed	Allowed
Custom Menu (not ride concurrent)	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Allowed	Allowed	Allowed
Insert Measure	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Allowed	Allowed	Allowed
DPM		NA	Blocked	Blocked	Blocked

Domain Lock Status Using domainStatus

The domainStatus utility provides a report on the processes that are locking the domains. The purpose is to identify the user activities that are preventing a ride process from running. This information is important since the ride process does not

terminate any existing workbook operation. The utility provides output that includes the process ID, user ID, operation type, and operation start time.

With the output of this utility, the system administrator can determine why the ride process is not running. This should provide enough information to either identify the user who is causing a delay or to manually terminate the process. That will be driven by the specific client and their processes.

Note that this utility reports on domain-level locks used during the ride process. Low-level data locks are not exposed by this utility.

domainStatus Usage

```
domainStatus -d domain -autoRefresh refreshPeriod
```

[Table 13–11](#) provides descriptions of the arguments used by the domainStatus utility.

Table 13–11 *domainStatus Usage*

Argument	Description
-d <i>domain</i>	Refers to a simple or master domain. When master domain is specified, all local domains in the global domain environment are locked as well as the master domain.
-autoRefresh <i>refreshPeriod</i>	Refreshes the lock status information every number of seconds specified by the refreshPeriod.

RPAS ODBC/JDBC Driver

The RPAS ODBC/JDBC Driver provides a SQL interface to the Oracle RPAS embedded database (OREDB), which includes both domain data and workbook data. This driver presents OREDB as a relational database to ODBC and JDBC client applications. The RPAS ODBC/JDBC Driver enables ODBC 3.51 and JDBC 3.0 compatible applications to connect to OREDB. Connectivity has been verified with the following applications:

- Oracle Business Intelligence Enterprise Edition
- Interactive SQL (ISQL) Utility
- JDeveloper

The RPAS ODBC/JDBC Driver enables system users to read measure data for stored measures in an RPAS domain.

- In a global domain environment, connection to local domains is not supported. Access to local domain data is possible through queries in global domains.
- The ODBC/JDBC Driver does not provide support for Forced Non-HBI (FNHBI) and non-materialized measures.
- The ODBC/JDBC Driver reports only external position names in both dimension tables and fact tables. Internal position names are not reported.
- Limited support is provided for conditional queries on measure data.
- The driver is not intended to replace the exportData utility, which is used for high-speed data export to ASCII files.

Note: For information on installing the RPAS ODBC/JDBC Driver, see the *Oracle Retail Predictive Application Server Installation Guide*.

- The ODBC driver inherits the RPAS user privilege. Note that if a user is not allowed to view any measures in RPAS, the user cannot view any measures through the ODBC driver either.

ODBC Configuration

On Windows, UNIX, and Linux platforms, configuring the system to connect the ODBC drivers and a domain environment consists of the following steps.

1. Install the ODBC server components. Refer to the *RPAS Installation Guide*.
2. Install the ODBC client components. Refer to the *RPAS Installation Guide*.

3. Start the RPAS ODBC Agent.
4. Configure the ODBC server components.
5. Configure the ODBC client components.
6. Create the ODBC data source name (DSN). This enables ODBC applications, such as OBIEE, to connect to the domain environments configured in the ODBC server and client configuration.
7. Start the RPAS ODBC Data Service.
8. Test the connection using Interactive SQL.

Defining the ODBC Server Configuration Settings

On UNIX and Linux platforms, upon completion of the ODBC server installation, a directory named `odbcserver` should be created under `$RPAS_HOME`. An RPAS ODBC Agent process should be started automatically. This Agent process works with the GUI ODBC Management Console installed on Windows PC to perform management and configuration tasks. The Windows version of RPAS ODBC Server must be installed on the Windows PC to make the ODBC Management Console available.

To define the ODBC Server configuration settings, do the following:


Adding RPAS ODBC Manager in the Management Console

To add RPAS ODBC Manager in the Management Console:

1. From the **Start** menu, select **Oracle RPAS ODBC Server** and then **Management Console**. The Oracle RPAS ODBC Management Console appears.


Note: When you start the Management Console for the first time, an error message appears indicating that the snap-in is not registered.

Figure 14–1 Console Manager Window Opened for First Time




2. To create a new work space, select **File** and then **New**.

Figure 14–2 Console Manager Window Showing Console Root Directory




To facilitate connection to a remote host, the Management Console is installed on a Windows PC and the Agent service is installed on a remote server.


3. To add a new snap-in, select **File** and then **Add/Remove Snap-in**. This starts the wizard to add a new snap-in.

Figure 14–3 Add/Remove Snap-In Dialog Box

4. To add a new snap-in, click **Add**.


Figure 14–4 Add Standalone Snap-in Dialog Box

5. Select **Oracle Retail RPAS ODBC Server Manager** and click **Add**.

Figure 14–5 Add Manager Dialog Box Showing Selection of the Mode for a New Manager

6. Select Local Host or Remote Host, depending on the location of the server you want to manage, and then click **Next**.

Figure 14–6 Add Manager Dialog Box Showing Connection Information to be Added for the Agent


The image shows a Windows-style dialog box titled "Add Manager" with a blue title bar and a red close button in the top right corner. The main area has a light beige background. At the top, under the heading "Agent", there is a message: "Please enter the connection information to your Agent. Click Finish to complete the wizard." Below this, a horizontal line separates the header from the input fields. The text "Specify the host on which your Agent is running. The host can be specified by either the host name or the host address." is followed by a label "Host:" and a text input field. Below that, the text "The port is the TCP port on which your Agent is listening. The TCP port must be in the range 1024 to 65534." is followed by a label "TCP Port:" and a text input field. At the bottom, there is a label "Encrypted (SSL):" followed by an unchecked checkbox. At the very bottom, there are three buttons: "< Back", "Finish", and "Cancel".

7. On the Add Manager window, enter the address of the server host machine and the TCP port number at which the RPAS ODBC Agent is listening.
8. If the Agent is not SSL enabled (which is the default), uncheck "Encrypted(SSL)". If the Agent is SSL enabled, check "Encrypted(SSL)".
9. Click **Finish**. You are returned to the previous window. You can click **Close** or **OK** to close the windows.

Configuring the Use of the ODBC Manager


From the Management Console, use the ODBC Manager to perform the required configuration.

1. To configure for a remote host, click + to expand the ODBC Manager you just added.

Figure 14–7 Console Manager Window Showing Expanded Console Root Directory

2. To connect to the Agent service, click + in front of "not connected to Unknown Service". If the Logon to Service dialog box is displayed, log on using the user name and password of a user who can administer the service on the local or remote server.
3. Expand Services by clicking +.

Figure 14–8 Console Manager Window Showing Expanded Connected to RPAS_ODBC_Agent Directory




4. Set up the environment variables.

Note: This step only applies to RPAS_Data_Service running on a Windows platform.

The PATH environment variable must be configured to include the following paths:

- Path to the bin subdirectory of RPAS_Home
- Path to the ODBC server binary
- Path to the ODBC server IP binary
- a. Navigate to the following screen. [Figure 14–9](#) shows the exact screen you should see.

Figure 14–9 Console Manager Window for Configuration



- b. Right-click on the blank space of the right panel.
- c. Select **New/Attribute** in the menu.
- d. Select **ServiceEnvironmentVariable** from the list for Attribute. In the Value field, enter the following:

```
PATH={pathToRpasHomeLib};{pathToODBCServerBin};{pathToODBCServerIPBin}
```

In the sample shown in [Figure 14–10](#), the following values are set:

- {pathToRpasHomeLib} is set to D:/src/rpas_head/rpasHome/lib
- {pathToODBCServerBin} is c:/odbcserver/bin
- {pathToODBCServerIPBin} is c:/odbcserver/ip/bin


Figure 14–10 New Attribute Dialog Box

- e. Click **OK**.
5. Expand Data Source Settings on the Console Manager window. You should see a pre-configured data source named gdom.

Note: When you create a new data source, make sure that most parameters of the new data source are identical to the parameters of the pre-configured data source gdom, except for DataSourceIPProperties and DataSourceIPCustomProperties.


The DataSourceIPCustomProperties attribute contains pre-registered property names, which must not be modified.

6. Under gdom, select IP Parameters.

Figure 14–11 Console Manager Window with IP Parameters Selected Under gdom

- On the right panel of the window, double click DataSourceIPProperties attribute. The following window appears.

Figure 14–12 DataSourceIPProperties Properties Dialog Box




Make sure the Value field has the keyword "DOMAIN_PATH=" and the string after the "=" sign is the absolute path to the RPAS domain you want connect to on the server side.

The following is an example of the path used to connect to a domain on the local host: "DOMAIN_PATH=C:\RPAS\11\Test\ODBC\nt_testGlobalDomain"

- On the right panel of the window, double click the DataSourceIPCustomProperties attribute. The following window appears.

Figure 14–13 DataSourceIPCustomProperties Properties Dialog Box



Copy the value from the DataSourceIPCustomProperties of the sample "gdom" data source to make sure all pre-registered properties are included.

Importing the Configuration Changes

1. To save the configuration, right click **Services** in the Console Manager window. In the menu, select **All Tasks** and then **Save Configuration**.
2. To save the snap-in configuration, select **File** and then **Save As**. Select `<installdir>\admin\oaadmin60.msc` for the file name. This overwrites the original file (which was basically empty).
3. Stop and start the RPAS Data Service.

Right-click RPAS_Data_Service. In the menu, click **Stop RPAS_Data_Service** or **Start_RPAS_Data_Service**. If you have made changes to any of the service attributes, you need to restart the Data Service.

The RPAS ODBC Data Service is now ready to accept connections.

Defining the ODBC Client Configuration for Windows


To define the ODBC Client configuration settings, do the following:

1. From the **Start** menu, select "ODBC Administrator" under the "Oracle RPAS odbc driver" menu item.

Note: Once the RPAS ODBC Client is successfully installed, a sample DSN named SampleRPAS is automatically created and configured to connect to the data source gdom on the server side. To create a new DSN, see the following steps.


2. In the ODBC Data Source Administrator window, select **Add** to add a data source.

Figure 14-14 ODBC Data Source Administrator Dialog Box



3. In the Create New Data source window, select **Oracle RPAS**.

Figure 14-15 Create New Data Source Dialog Box



4. In the Oracle Retail RPAS ODBC Driver Setup window, enter the following information:
 - Name and description of the ODBC data source.
 - In the Service Host field, enter the name of the server. If connecting to a service running on a local host, enter **localhost** or the name of the local host server. If the Agent service is running on a remote server, enter the name of the remote host server.
 - In the Service Port field, enter the port number that the data service is listening on.


Note: This is the port number of the data service and not the Agent service port number.

- In the Service Data Source field, enter the name of the service data source that has been configured for the data service. The default for this field is gdom.
- If the data service is not SSL enabled, uncheck **Encrypted(SSL)**. If the data service is SSL enabled, check **Encrypted(SSL)**.
- Enter custom properties, if needed. Custom properties are entered in the format of [name]=[value]. Multiple properties should be separated by a semicolon. For example: LANGUAGE=Japanese;WORKBOOK_SCHEMA=DOMAIN_T0


Table 14–1 lists the available custom properties (all optional):

Table 14–1 Available Custom Properties

Property Name	Description
LANGUAGE	Name of the language you use. The RPAS ODBC/JDBC driver is multi-language enabled. If data is in any language other than English, the LANGUAGE property should be set to the name of that language. If multiple languages are used in the domain, set this property to the name of the language other than English. For example, if some position names are in English and some are in Japanese, then LANGUAGE should be set to Japanese. The default is English.
WORKBOOK_SCHEMA	Name of the workbook you wish to connect to. If not set, the driver connects to the domain.
SHORT_DATE_FORMAT	Valid short date format used in RPAS.
DEFAULT_SCHEMA	Default schema name if the table name in the query is not qualified. This property is set to DOMAIN by the default configuration.
AGG_TABLE_NAMES	This property can be set to a list of valid aggregate table names separated by commas. When this is set, the driver presents the tables specified in the system tables. When this property is not set, the valid aggregate tables can still be queried even though they do not exist in the system tables.
NORMALIZE_DIM_TABLES	Valid values are Yes and No. The default value is No. If set to Yes, the dimension tables only contain columns for this dimension and its immediate parent dimension. If set to No, the dimension tables contain columns for this dimension and all parent dimensions within the hierarchy.

Figure 14–16 Oracle Retail RPAS ODBC Driver Setup Dialog Box

5. Click **Test Connect**. If the user security of the gdom Data source setting has been set to DBMSLogon, the Logon to Data Source dialog is displayed. Enter the data source user name and password configured for the data source.

Figure 14–17 Logon to Data Source Dialog Box

If the connection is successful, the following dialog box is displayed.

Figure 14–18 Oracle Retail RPAS ODBC Driver Setup Dialog Box for Successful Connection

6. Click **OK**.


Starting the RPAS ODBC Server Process

The RPAS ODBC Agent and Data Services should have automatically started after successful completion of the server installation.

The RPAS ODBC Data Service should be stopped and restarted using the Management Console.

1. To stop a data service, right-click the service name and then click **Stop** in the menu.
2. Right-click RPAS_Data_Service and then click **Start RPAS_Data_Service** in the menu.

Figure 14–19 Console Manager Window with Data Service Selected



Testing the Connection Using Interactive SQL

After the ODBC Server and ODBC Client have been configured, you can test the connection using Interactive SQL. The RPAS ODBC Server process must be running.

1. Select **Start, All Programs, Oracle RPAS ODBC driver, and then Interactive SQL (ODBC)**. The Interactive SQL command window appears.
2. Enter 'connect <user name>*<password>@<dsn_name>' where <dsn_name> is the name of the connection defined in the ODBC Server and ODBC Client configuration. The following is an example.

```
'connect adm*adm@SampleRPAS'
```

If the configuration is defined correctly, no errors are displayed.

ODBC Client Configuration for UNIX

Configuring the UNIX system to connect the ODBC drivers and a domain environment consists of the following steps:

1. Install the ODBC Server components. Refer to the *RPAS Installation Guide*.
2. Install the ODBC Client components. Refer to the *RPAS Installation Guide*.
3. Configure the ODBC Server components.
4. Configure the ODBC Client components.

5. Start the RPAS Data Service if it is not already started.
6. Test the connection using Interactive SQL.

Client Configuration

Both 32-bit and 64-bit ODBC Clients are available. They are delivered in directories named `odbcclient32` and `odbcclient64` respectively. The configuration steps are identical for the 32-bit and 64-bit ODBC Client.

Note: For the remainder of this chapter, the 32-bit and 64-bit ODBC Client are referred to as ODBC Client, and `odbcclient32` and `odbcclient64` are referred to as `odbcclient`.

If it comes with RPAS, then `odbcclient` directory is under your `$RPAS_HOME`. If it comes separately, the installer determines its location.

1. Set up the environment for the ODBC Client.

If the ODBC client does not come with RPAS (meaning the `odbcclient` directory is not under `$RPAS_HOME`), edit the `oaodbc.sh` file (`oaodbc64.sh` for 64-bit Client) in `odbcclient`:

- a. Make sure the following environment variables are set correctly:
 - `LIBPATH` and `OASDK_ODBC_HOME` are set to the full path of the lib directory inside the `odbcclient` directory
 - `ODBCINI` is set to the full path of the `odbc.ini` file (`odbc64.ini` for 64-bit Client), including the file name, inside the `odbcclient` directory
- b. Source `oaodbc.sh` by running the following command in the `odbcclient` directory:


```
../oaodbc.sh
```

2. Create and configure the data sources in `odbc.ini`.

The `odbc.ini` file in the `odbcclient` directory has three sections: `[ODBC]`, `[ODBC Data Sources]`, and `[SampleRpas]`, which is a section for the sample RPAS data source.

- a. Edit the `[ODBC]` section: Set `TraceDll` to the full path to `lib/odbctrac.so` and `InstallDir` to the full path of the `odbcclient` directory.
- b. Edit the `[ODBC Data Sources]` section: Add an entry for the new data source you are creating. The entry has the following format:

```
MyRPASDataSource= Oracle RPAS ODBC Driver
```
- c. Create a new `[MyRPASDataSource]` section: The `[SampleRPAS]` section can be copied and modified. In the new `[MyRPASDataSource]` section:
 - Set `Driver` to the full path of `odbcclient/lib/ivoa22.so`.
 - Set `Host` to the name or IP address of the server.
 - Set `Port` to the port number the `RPAS_Data_Service` listens at. This is not the port number used by the RPAS ODBC Agent.
 - Set `ServerDataSource` to the name of the data source you created in the server configuration. For `[SampleRPAS]`, this entry is set to `"gdom"`, since that is the data source created on the server as an example. This is shown in the following figure.

Figure 14–20 Console Manager Window with gdom Selected

Testing the Connection

After the ODBC Server and ODBC client have been configured, you can test the connection using Interactive SQL. The RPAS ODBC Data Service must be started.

1. In the odbcclient directory, source oaodbc.sh if you have not already done so.
2. Change to the tools directory and run the executable odbccisql. Then, at the ISQL prompt, enter 'connect <user Name>*<password>@<DataSourceName>' where <DataSourceName> is the name of the data source you defined in odbc.ini, as described in the previous section. The following is an example.

```
'connect adm*adm@MyRPASDataSource'
```

If the configuration is correctly defined, no errors are displayed.

Installing and Using the RPAS JDBC Driver

This section describes how to install, set up, and use the RPAS JDBC driver on all UNIX, Linux, and Windows platforms.


The RPAS JDBC driver is delivered in a single (zipped) jdbcclient directory. The user has full control on the location of the jdbcclient directory. For installation guidelines, refer to the *RPAS Installation Guide*.

After the JDBC driver is installed on your system, you need to update the CLASSPATH environment variable. This variable ensures that the JDBC client can access the appropriate Java classes needed to connect to the database.

Updating Environment Variables for the JDBC Driver on Windows


1. Open **System** in the Control Panel. The System Properties window appears.
2. On the **Advanced** tab, click **Environment Variables**. The Environment Variables dialog appears.

Figure 14–21 Environment Variables Dialog Box



3. Select the **CLASSPATH** from the **System variables** list and click **Edit**. The Edit System Variable dialog box appears.

Figure 14–22 Edit System Variable Dialog Box



4. Add the current working directory ".", driver_home/driver/lib/ORjc.jar, driver_home/driver/ORssl14.jar, and driver_home/driver/iaik_jce_full.jar to the CLASSPATH environment variable and click **OK**.

Note that driver_home is the location where the jdbcclient was installed. If your jdbcclient was installed in C:/jdbcclient, you see the following in your CLASSPATH:

```
.;C:/jdbcclient/driver/lib/ORjc.jar; C:/jdbcclient/driver/lib/ORssl14.jar;  
C:/jdbcclient/driver/lib/iaik_jce_full.jar
```

Note: Separate paths with semi-colons (;).

5. After updating the environment variable, restart your PC.

After you have updated the environment variables and restarted your PC, you are ready to use the RPAS JDBC driver with any JDBC client.

Updating Environment Variables for JDBC Driver on UNIX and Linux

On UNIX and Linux systems, use export (or set, depending what shell you use) to add the following to your CLASSPATH:

```
export CLASSPATH=.:jdbc_home/driver/lib/ORjc.jar: jdbc_
home/driver/lib/ORssl14.jar: jdbc_home/driver/lib/iaik_jce_full.jar:$CLASSPATH
where jdbc_home is the full path of the directory where jdbcclient is installed. If you
installed jdbcclient at /usr/products/oracle, then you should replace jdbc_home with
/usr/products/oracle/jdbcclient.
```

The above export command can be added to your .profile.

Using the RPAS JDBC Driver

Any JDBC client needs the following information to use a JDBC driver to connect to a database:

- A driver class
- A URL to the database specified in a form that the particular JDBC driver understands

For the RPAS JDBC driver, this information is specified as follows:

- Driver Class: com.oracle.ard.jdbc.openaccess.OpenAccessDriver
- URL: "jdbc:RPAS://<host>:<port>;ServerDataSource=<DataSourceName>"

<host> is the name or IP address of the server, <port> is the port number the RPAS Data Service listens at, and <DataSourceName> is the name of data source you created for the RPAS Data Server (it is "gdom" in the default configuration).

Enabling Spy for the RPAS JDBC Driver

Spy is a logging facility for the JDBC driver. To enable spy for the RPAS JDBC connection, do the following:

1. Add jdbc_home/spy/lib/ORy.jar to your CLASSPATH where jdbc_home is the installation directory of jdbcclient.
2. Set your driver class to com.oracle.ard.jdbcspy.SpyDriver.
3. Use the following URL:

```
"jdbc:spy:{jdbc:RPAS://<host>:<port>;ServerDataSource=<DataSourceName>};load=
com.oracle.ard.jdbc.openaccess.OpenAccessDriver;[key=value];..."
```

<host> is the name or IP address of the server, <port> is the port number the RPAS Data Service listens at, and <DataSourceName> is the name of data source you created for the RPAS Data Server. (It is "gdom" in the default configuration.) The key and value pairs are the attributes of the Spy class.

Table 14-2 lists the available attributes:

Table 14–2 Attributes Available

Key and Value	Description
log=System.out	Redirects logging to the Java output standard, System.out.
log=(file) <i>filename</i>	Redirects logging to the file specified by <i>filename</i> . For example, C:\temp\spy.log
linelimit= <i>numberofchars</i>	Specifies the maximum number of characters that Spy logs on one line. When set to no (default), there is no maximum limit on the number of characters.
logLobs={yes no}	Specifies whether Spy logs activity on Blob or Clob. The initial default is no.
logIS={yes no nosingleread}	Specifies whether Spy logs activity on InputStreams. When logIS=nosingleread, logging on InputStream and Reader objects is active; however, logging of the single-byte read InputStream.read or single-character Reader.read is suppressed to prevent generating large log files that contain single-byte or single character read messages. When set to no (default), Spy does not log activity on InputStreams.
logTName={yes no}	Specifies whether Spy logs the name of the current thread. When set to no (default), Spy does not log the name of the current thread.
timestamp={yes no}	Specifies whether a timestamp should be included on each line of the Spy log. When set to no (default), Spy does not include a timestamp on each line.

Using the jdbcisql Utility Provided with the RPAS JDBC Driver

Oracle Retail suggests that you use the jdbcisql.bat (for Windows) or jdbcisql.sh (for UNIX and Linux) located in jdbcclient/isql to start jdbcisql. Edit jdbcisql.bat or jdbcisql.sh to make sure it uses the appropriate URL (the argument of the -u option):

```
java jdbcisql -d com.oracle.ard.jdbc.openaccess.OpenAccessDriver -u
"jdbc:RPAS://<host>:<port>;ServerDataSource=gdom"
```

Note: The value of the ServerDataSource setting in the URL is case-sensitive on Windows and UNIX systems and has to match the "Data source setting" defined for the data service.

When the application starts, enter the following to log in and make the connection (user name = adm; password = adm) as shown below.

```
Connect adm*adm@
```

To enable Spy for jdbcisql, use a command line similar to the following:

```
java jdbcisql -d com.oracle.ard.jdbc.spy.SpyDriver -u "jdbc:spy:{jdbc:RPAS://
:<host>:<port>;ServerDataSource=gdom};load=com.oracle.ard.jdbc.openaccess.OpenA
ccessDriver;log=(file)C:\temp\spy.log;logIS=yes;logTName=yes;timestamp=yes"
```

Using Oracle SQL Developer

Create an XML file with the following content:


```

<?xml version = '1.0'?>
<!DOCTYPE connections>
<connections>
  <connection>
    <URL>jdbc:RPAS://<host>:<port>;ServerDataSource=<DataSourceName></URL>
    <ConnectionName>MyConnection</ConnectionName>
    <user>adm</user>
    <ConnectionType>OTHER_JDBC</ConnectionType>
    <JdbcDriver>com.oracle.ard.jdbc.openaccess.OpenAccessDriver</JdbcDriver>
  </connection>
</connections>

```

- The URL should correspond to the URL specification required by the RPAS JDBC Driver, as specified in the preceding sections.
- ConnectionName can be anything you like. This field can be changed later using the client application.
- Enter a user name for the connection. This field can also be changed later using the application.
- Leave the remaining information as shown in the code sample above.

To set up the connection:

1. Save this XML file with any name you like.
2. In SQL Developer, using the Tools/Preferences/Database/Third Party Drivers, add the ORjc.jar, Orssl14.jar, and iaik_jce_full.jar files to the list of third party drivers used by SQL Developer. (SQL Developer does not look in the classpath for drivers.)
3. Go to the Connection Navigator and right-click on Connections. Select **Import Connections**.
4. Browse to the XML file. The dialog displays the list of connections you specified in the file. Choose your connection in the sample code, MyConnection.

Using Oracle JDeveloper

Perform the following procedure to use Oracle JDeveloper with the JDBC driver:

1. Start JDeveloper.
2. From the JDeveloper left panel, select the **Connections** tab.
3. Right-click on **Databases** and select **New Database Connection**. The Create New Database Connection wizard appears.
4. On the first screen of the Create New Database Connection wizard, enter a connection name and choose **Third Party JDBC driver** for **Connection Type**.
5. On the second screen, enter the user name and password and then click **Next**.
6. On the third screen, perform the following:
 - a. Click **New** to add the driver.
 - b. Locate the library ORjc.jar, Orssl14.jar, and iaik_jce_full.jar files and their path. These jar files are available from the installation of the RPAS JDBC Client.
 - c. Enter the RPAS JDBC Driver connection URL, as specified at the beginning of this section.

- d. In the **Driver Class** field, enter
com.oracle.ard.jdbc.openaccess.OpenAccessDriver.
7. Follow the instructions to finish creating the connection.

Using a Java Program

You can instantiate oadriver in your application using one of following methods:

- new oadriver();
- Class.forName
("com.oracle.ard.jdbc.openaccess.OpenAccessDriver").newInstance();

Make sure *driver_home*/driver/lib/ORjc.jar, *driver_home*/driver/lib/Orssl14.jar, and *driver_home*/driver/lib/iaik_jce_full.jar are included in the CLASSPATH.

The Java code snippet below shows you how you can write a program that uses the driver.

Java Code Sample:

```
import java.sql.*;
public class RPASDriverTest {
    public RPASDriverTest() {}
    public static void main(String[] args)
    {
        try
        {
            if (args.length != 3)
            {
                System.out.println("Format:\n" +
                    "java RPASDriverTest <Database> <UID> <PWD>\n");
                return;
            }
            Connection conn = null;
            Driver d =
                (Driver)Class.forName("com.oracle.ard.jdbc.openaccess.OpenAccessDriver").newInstan
                ce();
            String url = "jdbc:RPAS://";
            String database = args[0];
            String uid = args[1];
            String pwd = args[2];
            url += database;
            System.out.println("Trying to connect to url: " + url );
            conn = DriverManager.getConnection(url, uid, pwd);
            DatabaseMetaData dma = conn.getMetaData();
            System.out.println("\nConnected to " + dma.getURL());
            System.out.println("Driver " +
                dma.getDriverName());
            System.out.println("Version " +
                dma.getDriverVersion());
            System.out.println("");
            // sample query
            String query = "SELECT * FROM DIM_YEAR";
            Statement stmt = conn.createStatement();
            ResultSet rs = stmt.executeQuery(query);
            rs.close();
            stmt.close();
        }
        catch (SQLException ex)
        {
```

```

System.out.println ("\n*** SQLException caught ***\n");
while (ex != null) {
    System.out.println ("SQLState: " + ex.getSQLState ());
    System.out.println ("Message: " + ex.getMessage ());
    System.out.println ("Vendor: " +
        ex.getErrorCode ());
    ex = ex.getNextException ();
    System.out.println ("");
}
}
catch (java.lang.Exception ex)
{
    //Got some other type of exception. Dump it.
    ex.printStackTrace ();
}
}
}

```

Running the Program

After compilation, run the program as:

```
java RPASDriverTest.class "<host>:<port>;ServerDataSource=<DSN>" <uid> <pwd>
```

Where <host> is the IP of the server box where RPAS ODBC Server is running and where <port> is port number of the ODBC Server. <DSN> is the data source name that is created on the server. Note the double quotes must be included because of the semicolon.

Data Query

This section provides the details of data query, including the limitations, metadata, and dimension tables.

Limitations

Note the following limitations when performing data queries.

Contention

Workbook commit requests issued by RPAS users compete with real-time intra-day reports (ODBC queries) when accessing the domain data. Commit requests imply 'write' locks on certain measures in the domain. While such 'write' locks are in place, reporting tools cannot access the same data. In such situations both reporting and workbook users can experience latencies or, in more severe scenarios, they can encounter system feedback informing them that their last operation did not succeed and they need to re-execute their last request. The likelihood of such concurrency issues depends on a number of factors, such as hardware capabilities (CPU and IO capacity), reporting volume, reporting granularity, number of reporting users, number of workbook users, and the commit data volume. To minimize the chance of such concurrency issues, all workbook commits should leverage the commit ASAP framework. Additionally, RPAS allows reporting from workbooks that is less prone to concurrency issues that may be experienced while reporting from domains.

Workbook Queries

All workbook queries must be performed against saved workbooks. The workbook can be open or closed. The user must save the workbook before reporting.


Metadata

The following figure shows the metadata tables available in a domain or workbook. These tables can be used to examine the structure of the domain, such as:

- Which measures and dimensions exist within the database
- Which hierarchies exist and their rollup structure
- Which fact tables are available
- Which measures exist at the intersections that they represent

When connected to a domain, an additional table (MD_WORKBOOK_SCHEMAS) is available to list all accessible workbooks within the domain with their schema names.

Figure 14–23 Database Diagram for All Metadata Tables in a Domain or in Each Workbook




Note: The MD_WORKBOOK_SCHEMAS table is not included in workbooks.

Fact and Dimension Tables

The following figure shows an example of the structure of fact and dimension tables and the relationships between them. A fact table represents an intersection where one or more measures' data is stored. Each measure is represented by a column in the table. Additionally, each dimension on the intersection is represented by a column. A record in the fact table is uniquely identified by a unique combination of position names for the intersecting dimensions.

A dimension table represents a dimension. It includes a column to list all position names, their labels, and their rollup mapping to each dimension at higher levels in the hierarchy.



characters long, and if the dimension name is less than four characters, it is padded with '_' characters to make it four characters long.

For information on limitations when performing queries, see the [Limitations](#) section.

Measure Security in the ODBC Driver

Before the existence of the ODBC/JDBC driver, an RPAS user could only use RPAS workbooks to access measures. Consequently, the ODBC/JDBC driver emulates the RPAS workbook security model to manage access rights to RPAS measures. It allows users to view all measures that they could view using the templates to which they have access.

This means that when users connect to a domain, they can see all the measures that they could insert into a workbook. These include all measures that their templates have access rights to (managed through the use of the Workbook Template Measure Rights worksheet in the Security Administration workbook) and all measures to which the users have explicitly been given access rights using the Measure Rights worksheet in the Security Administration workbook. All other measures are not accessible to the users.

When users connect to a workbook, they can access all measures in the workbook, irrespective of how those measures were brought into the workbook and irrespective of whether access rights to some of those measures were removed after the workbook was created. Since those measures exist in the workbook that the users can access, those measures (their workbook copies) are accessible to the users.

Use Cases

This section describes some use cases of RPAS.

Using Metadata Tables to Explore the Structure of a Domain or a Workbook

To explore the structure of a domain or a workbook, do the following:

1. Fetch the list of workbook schema names. In this example, the workbook is owned by 'USER01', built using the template 'TestTemplate', and labeled 'MyWorkbook'.

```
Select
    SCHEMA_NAME, CREATION_TIME, MODIFICATION_TIME
From
    MD_WORKBOOK_SCHEMAS
Where
    OWNER_NAME = 'USER01' and
    WORKBOOK_LABEL = 'MyWorkbook' and
    WORKBOOK_TEMPLATE = 'TestTemplate'
```

The SCHEMA_NAME obtained using this query can be directly used in the custom properties of the driver configuration to enable direct connection to a workbook instead of a domain.

2. List all measures in the domain or workbook (default schema).

```
Select
    *
From
    MD_MEASURES
```

3. List all measures in a specific schema (for example, 'DOMAIN_T0').

```

Select
    *
From
    DOMAIN_T0.MD_MEASURES

```

4. List all dimensions in the domain or workbook (default schema).

```

Select
    *
From
    MD_DIMENSIONS

```

5. List all fact tables in the domain or workbook (default schema).

```

Select
    *
From
    MD_FACT_TABLES

```

6. List all hierarchies in the domain or workbook (default schema).

```

Select
    *
From
    MD_HIERARCHIES

```

7. List all fact tables with the measures that are represented in those tables (default schema).

```

Select
    *
From
    MD_TABLE_MEASURES

```

List all fact tables with the dimension table names that intersect in the fact table (default schema).

```

Select
    *
From
    MD_TABLE_INTERSECTIONS

```

8. Use the following to see the structure of a particular hierarchy (for example: CLND). It lists the hierarchy with each of its dimensions and the roll up dimension name for each one of them (default schema).

```

Select
    *
From
    MD_HIERARCHY_SPECS
Where
    HIERARCHY_NAME = 'CLND'

```

Querying Fact Data

To query fact data, do the following:

1. Query fact data for all measures at the STR-SKU-DAY intersection with the unique position names for these dimensions.

```

Select

```

```

*
From
    FACT_STR_SKU_DAY_

```

2. Query fact data for specific measures at the STR-SKU-DAY intersection and list them with the position labels for each dimension.

```

Select
    DS.POSITION_LABEL, DU.POSITION_LABEL, DD.POSITION_LABEL, R_EX_LBLOWINTR, R_EX_
    UBLOWINTR
From
    FACT_STR_SKU_DAY_ F,
    DIM_STR DS,
    DIM_SKU DU,
    DIM_DAY DD
Where
    DS.POSITION_NAME = F.STR and
    DU.POSITION_NAME = F.SKU and
    DD.POSITION_NAME = F.DAY
Hint Join (FACT_STR_SKU_DAY_, DIM_STR, DIM_SKU, DIM_DAY);

```

Note: The optional Hint clause in the above SQL statement is not ANSI SQL standard, but the ODBC/JDBC Driver supports it. This Hint tells the driver to process the join tables in the specified order (fact table first, and then dimension tables).

Connecting to a Workbook

Complete the following steps to connect to a workbook:

1. Select **Start, Settings, Control Panel, Administrative Tools**, and then **Data Sources (ODBC)**.
2. Select the **System DSN** tab. Select the appropriate DSN and click **Configure**.
3. In the Options frame, enter **WORKBOOK_SCHEMA=<workbook schema name>**. Replace '<workbook schema name>' with the workbook schema name for the workbook to which you want to connect. The workbook schema names can be obtained by first connecting to the domain and then examining the MD_WORKBOOK_SCHEMAS table to obtain the schema name for the appropriate workbook (may be identified by owner name, template, creation and last modification time). For example: 'WORKBOOK_SCHEMA=DOMAIN_T0' or 'WORKBOOK_SCHEMA=SD0_T0'
4. Click **OK**.

Requesting Additional Aggregate Tables

To request additional aggregate tables, do the following:

1. Select **Start, Settings, Control Panel, Administrative Tools**, and then **Data Sources (ODBC)**.
2. Select the **System DSN** tab. Select the appropriate DSN and click **Configure**.
3. In the **Options** frame, enter **AGG_TABLE_NAMES=<comma-separated list of any additional aggregate fact table names>**.

By default, the database includes every fact table (a fact table represents an intersection) that one or more measures have as their base intersection. Any other

fact tables can be specifically requested by adding a comma-separated list as the value for this custom property. For example, to see a fact table for the intersections 'DEPT' and 'DEPT_YEAR', use the following value of this custom property: 'AGG_TABLE_NAMES=FACT_DEPT, FACT_DEPT_YEAR'.

4. Click **OK**.

If entering more than one connection property (that is, both the WORKBOOK_SCHEMA and AGG_TABLE_NAMES properties), separate the property key value pairs by a semicolon. Using the examples above, the content of the custom properties input box should be as follows: WORKBOOK_SCHEMA=DOMAIN_T0;AGG_TABLE_NAMES= FACT_DEPT, FACT_DEPT_YEAR

Clients

This section lists some sample ODBC/JDBC client applications that can connect to the RPAS datastore through the RPAS ODBC/JDBC Driver. The examples in this section do not include all client applications that can connect to the RPAS ODBC/JDBC Driver.

Note: In a client/server configuration, the server (executable) must be started before a client can connect to it.

Oracle Business Intelligence Enterprise Edition (OBIEE)

This section outlines how to connect to the defined DSN using the OBIEE Administration Tool and how to import data from the DSN. For more information about OBIEE, refer to OBIEE documentation. The user must install and configure the ODBC client first on the OBIEE server host (refer to section **ODBC Client Configuration for UNIX**) and test the connection. The ODBC client and the OBIEE server must both be 32-bit or 64-bit. The administrator must source the oaodbc.sh or oaodbc64.sh script under the ODBC client home directory before starting or re-starting the OBIEE server.

Configuring the ODBC Client for OBIEE

The following example provides a sample of a configuration of the ODBC client for OBIEE. This example was developed for OBIEE on AIX, but the process is the same for other environments.

1. Open the \$BIEE_HOME/setup/odbc.ini file, where \$BIEE_HOME is the directory where OBIEE is installed.
2. Set the TraceDll to the odbctrac.so that comes with RPAS odbcclient. Set InstallDir to the RPAS odbcclient installation directory.
3. In the [ODBC Data Sources] section, insert an entry for RPAS domain.

Example:

```
rpas_domain=This is the name of the data source for RPAS.
The name here (rpas_domain) should be the same as the data source name
configured in the RPAS ODBC Server.
```

4. Create a section in the file for the rpas_domain. The following example is subject to changes. Refer to the [SampleRPAS] section in odbc.ini or odbc64.ini under ODBC client home directory for all up-to-date settings.

Example:

```
[rpas_domain]
Driver= absolute_path_to_odbc_client/lib[64]/ivoa22.so
Description=Oracle Retail RPAS ODBC Driver
Host=<RPAS ODBC Server host>
Port=<odbc_data_service_port>
ServerDataSource=<data_source_name>
UseLDAP=0
DistinguishedName=
Encrypted=0
LoadBalancing=0
AlternateServers=
ConnectionRetryCount=0
ConnectionRetryDelay=3
CustomProperties=
```

Save your changes to the file.

Connecting OBIEE to an RPAS Domain

To connect OBIEE to a predefined DSN for an RPAS Domain:

1. Make sure the following Windows services are running:
 - Oracle BI Java Host
 - Oracle BI Server
2. Start the OBIEE Administration Tool. Select **Start, All Programs, Oracle Business Intelligence**, and then **Administration**.
3. From the File menu, select **open - online**. A window appears to use for entering login credentials.
4. Enter the administrator's user name and password and then click **open**.
Three panels now appear in the Admin Tool window: **Presentation, Business Model and Mapping**, and **Physical**.
5. From File menu, select **Import-From database**. A window appears to select the connection type and RPAS user information.
6. In order to optimize the OBIEE queries issued to the RPAS ODBC driver, you must select the appropriate connection type when configuring OBIEE. Selecting the RPAS-specific connection type will cause OBIEE to generate physical queries that delegate data aggregations to RPAS. For more, see *Oracle Fusion Middleware Integrator's Guide for Oracle Business Intelligence Enterprise Edition / 11g Release 1 (11.1.1)*.

Once you select the connection type, the RPAS schemas and tables appear in a new window.

7. Select the objects you want to import and then click **import**. After the import is complete, click **Close**.
A new physical model is created and listed in the **Physical** panel of the Admin Tool window.
8. Expand the physical model. Double-click on **connection** to open "connection" properties. Make sure the **Connection Type** is set to **ODBC 3.5**. Click **OK** to exit.
9. In the Admin Tool window, click **Save** to save your physical model.

Now that you have a basic physical model, you can build the business model and presentation layer on top of it. For more information on the business model, presentation layer, and OBIEE Web interface, refer to the OBIEE documentation.

Microsoft Access

To connect using Microsoft Access, do the following:

1. Start Microsoft Access.
2. Create a new (or open an existing) Access file (.mdb file).
3. From the File menu, select **Get External Data- Link Tables** (or **Import** if you want to import the data from RPAS datastore to Access). A dialog box appears.
4. In the Files of type box, select **ODBC Databases()**.
5. Click the **Machine Data Source** tab and then double-click the pre-configured ODBC data source from which you want to link.
6. At the logon prompt, enter your user ID and password and then click **OK**.

At this point, MS Access connects to the RPAS data source and displays the list of schemas and tables that you can import or link.

7. Click each table that you want to import or link and then click **OK**. If you are linking a table that does not have an index that uniquely identifies each record, then you will see a list of the fields in the linked table. Select a field, or a combination of fields, that uniquely identify each record and then click **OK**.

JDeveloper

JDeveloper works best with a native JDBC driver, which is included in the RPAS ODBC/JDBC Driver package.


To connect using JDeveloper:

1. Start JDeveloper.
2. On the JDeveloper left panel, select the **Connections** tab.
3. Right-click on **Databases** and select **New Database Connection**.
4. On the first screen of the Create New Database Connection wizard, enter a connection name and select **Third Party JDBC driver** for Connection Type.
5. On the second screen, enter the user name and password and then click **Next**.
6. On the third screen, click **New** to add the driver. It opens up the following dialog. You need to find the jar files ORjc.jar, ORssl14.jar, iaik_jce_full.jar, and ORy.jar and their paths (they are made available from the installation of the JDBC Driver). Then, create a library group "RPAS JDBC Driver" with those four jar files.

In the Driver Class field, enter

com.oracle.ard.jdbc.openaccess.OpenAccessDriver.

Then, enter in the URL field: **jdbc:RPAS://{host_name}:{port_number};ServerDataSource={data_source_name}**, where **host_name** is the host name or IP address of the ODBC server, **port_number** is the RPAS_Data_Service port number, and **data_source_name** is the Data Source Name that is configured on the ODBC server.

Figure 14–25 Register JDBC Driver Dialog Box

7. Follow the instructions to finish creating the connection.

After the connection is established, you can expand the Connection (and the nodes under the Connection) to browse the objects in the RPAS datastore. You can also open a SQL worksheet (by selecting **SQL Worksheet** from the Tools menu) to write or execute SQL statements.

XML Publisher

This section describes how to make the connection from XML Publisher to RPAS using the JDBC driver. (XML only supports JDBC.)

1. Install and configure the JDBC client driver and ODBC/JDBC server. Start the server.
2. Copy the JDBC client jar files ORjc.jar, ORssl14.jar, iaik_jce_full.jar, and ORy.jar (from JDBC client installation) to D:\OraHome_1\oc4j\j2ee\home\applib, where D:\OraHome_1 is the root directory where XML Publisher was installed.
3. Start the XML Publisher server. (Select **Start, All Programs, Oracle XML Publisher Server, OUIHome1**, and then **Oracle XML Publisher Enterprise Start**.)

4. Start a Web browser and go to the URL: <http://localhost:15101/xmlpserver/>

This URL is only an example. Contact the XML Publisher administrator/installer for the actual URL. The actual URL is recorded in D:\OraHome_1\xmlpserver\setupinfo.txt file of the XML Publisher server machine.

5. Log in as admin/admin.
6. Select the **Admin** tab and then select **JDBC Connection** under Data Sources to create a JDBC connection.
7. Click **Add Connection** to create a new connection. Provide the following information:
 - Enter a display name for Data Source Name.
 - Enter `jdbc:RPAS://{host_name}:{port_number};ServerDataSource={data_source_name}` for the URL, where **host_name** is the host name or IP address of the ODBC server, **port_number** is the RPAS_Data_Service port number, and **data_source_name** is the Data Source Name that is configured on the ODBC server.
 - Enter **adm** for user name and password.
 - Enter `"com.oracle.ard.jdbc.openaccess.OpenAccessDriver"` for Database Driver Class.

8. Click **Test Connection**. The confirmation message: "connect established successfully" should appear.
9. Click **Apply** to save the connection.

Interactive SQL (ISQL) Utility

ISQL is an interactive SQL tool that is provided by the ODBC/JDBC SDK.

To connect to the remote ODBC/JDBC server, use `odbcisql.exe` (with ODBC client installed) or `jdbcisql.class` (with JDBC client installed).

Note: Users are expected to know basic SQL in order to use ISQL.

To connect to the ODBC/JDBC server using `odbcisql`, start `odbcisql` and then, at the SQL prompt, issue the connect command as follows:

```
connect john/does@rpasDomain
```

where john/does is a predefined administrator account in RPAS and `rpasDomain` is a pre-configured Data Source Name.

Issue the connect command for `jdbcisql.class` as follows:

```
connect john*does@rpasDomain
```

After you are connected to the server, you can issue various SQL DML and DDL statements to inspect and modify the data in the RPAS datastore.

Supported and Unsupported SQL Functions

This section contains the following information:

- Detailed descriptions of various functions supported by the RPAS ODBC Driver.
- Descriptions of the SQL92 and SQL99 functionality that is not supported.

Supported SQL Functions

Use caution when applying functions to any dimension name or label columns, because the driver is not able to use the corresponding internal indexes to optimize row selection when functions are applied to those columns (which could be a significant performance hit).

It is suggested that users avoid applying functions to dimension name or label columns whenever possible.

Consider the following query:

```
Select * from fact_str_sku_day
where convert(day, SQL_DATE) = curdate();
```

Even though this query selects the data for only one day, the driver has to scan the entire fact table and then apply the convert function to every row of the table.

Working with OBIEE, the same can be achieved using a variable, which holds the converted string value of the current date (in the same format as the "day" column). The query then becomes:

```
Select * from fact_str_sku_day
```

Where day = @curDateString;

The driver only reads the rows that meet the condition.

Numeric Functions

Table 14–3 Numeric Functions

Function	Description
<code>ABS(numeric_exp)</code>	Returns the absolute value of <i>numeric_exp</i> . For example: SELECT ABS(-1.0), ABS(0.0), ABS(1.0) FROM emp WHERE empno = 1; This returns three result columns with values 1, 0, and 1.
<code>ACOS(float_exp)</code>	Returns the arccosine of <i>float_exp</i> as an angle, expressed in radians. For example: SELECT ACOS(-1) FROM emp WHERE empno = 1; This returns 3.14159.
<code>ASIN(float_exp)</code>	Returns the arcsine of <i>float_exp</i> as an angle, expressed in radians. For example: SELECT ASIN(-1.0) FROM emp WHERE empno = 1; This returns -1.57079.
<code>ATAN(float_exp)</code>	Returns the arctangent of <i>float_exp</i> as an angle, expressed in radians. For example: SELECT ATAN(45.0) FROM emp WHERE empno = 1; This returns 1.54857.
<code>ATAN2(float_exp1, float_exp2)</code>	Returns the arctangent of the x and y coordinates, specified by <i>float_exp1</i> and <i>float_exp2</i> , respectively, as an angle, expressed in radians. For example: SELECT ATAN2(35.175, 129.44) FROM emp WHERE empno = 1; This returns 0.2653399.
<code>CEILING(numeric_exp)</code>	Returns the smallest integer greater than or equal to <i>numeric_exp</i> . The return value is of the same data type as the input parameter. For example: SELECT CEILING(123.45), CEILING(-123.45), CEILING(0.0) FROM emp WHERE empno = 1; This returns 124, -123 and 0.
<code>COS(float_exp)</code>	Returns the cosine of <i>float_exp</i> , where <i>float_exp</i> is an angle expressed in radians. For example: SELECT COS(14.78) FROM emp WHERE empno = 1; This returns -0.59946542.
<code>COT(float_exp)</code>	Returns the cotangent of <i>float_exp</i> , where <i>float_exp</i> is an angle expressed in radians. For example: SELECT COT(124.78) FROM emp WHERE empno = 1; This returns -0.82045588.
<code>DEGREES(numeric_exp)</code>	Returns the number of degrees converted from <i>numeric_exp</i> radians. For example: SELECT DEGREES(3.143) FROM emp WHERE empno = 1; This returns 180.0806.

Table 14–3 (Cont.) Numeric Functions

Function	Description
EXP(float_exp)	Returns the exponential value of <i>float_exp</i> . For example: SELECT EXP(378.615) FROM emp WHERE empno = 1; This returns 2.69404760606322E+164
FLOOR(numeric_exp)	Returns the largest integer less than or equal to <i>numeric_exp</i> . The return value is of the same data type as the input parameter. For example: SELECT FLOOR(123.45), FLOOR(-123.45) FROM emp WHERE empno = 1; This returns 123 and -124.
LOG(float_exp)	Returns the natural logarithm of <i>float_exp</i> . For example: SELECT LOG(5.175643) FROM emp WHERE empno = 1; This returns 1.64396358.
LOG10(float_exp)	Returns the base 10 logarithm of <i>float_exp</i> . For example: SELECT LOG10(145.175643) FROM emp WHERE empno = 1; This returns 2.161893758. SELECT LOG10(0), LOG10(-1), LOG10(1) FROM emp WHERE empno = 1; This returns -1.#INF, -1.#IND and 0
MOD(integer_exp1, integer_exp2)	Returns the remainder (modulus) of <i>integer_exp1</i> divided by <i>integer_exp2</i> . For example: SELECT mod(empno, 2) FROM emp WHERE empno = 11; This returns 1.
PI()	Returns the constant value of pi as a floating-point value. For example: SELECT PI() FROM emp WHERE empno = 1; This returns 3.14159265358979.
POWER(numeric_exp, integer_exp)	Returns the value of <i>numeric_exp</i> to the power of <i>integer_exp</i> . For example: SELECT POWER(2, -5), POWER(2, 5) FROM emp WHERE empno = 1; This returns 0, 32.
RADIANS(numeric_exp)	Returns the number of radians converted from <i>numeric_exp</i> degrees. For example: SELECT RADIANS(45.0) FROM emp WHERE empno = 1; This returns 0.785398.
RAND([integer_exp])	Returns a random floating-point value using <i>integer_exp</i> as the optional seed value. For example: SELECT RAND(0) FROM emp WHERE empno = 1; This returns 38.

Table 14–3 (Cont.) Numeric Functions

Function	Description
ROUND(<i>numeric_exp</i> , <i>integer_exp</i>)	Returns <i>numeric_exp</i> rounded to <i>integer_exp</i> places right of the decimal point. If <i>integer_exp</i> is negative, <i>numeric_exp</i> is rounded to $integer\_exp$ places to the left of the decimal point. For example: <pre>SELECT ROUND(123.344, 2), ROUND(123.345, 2) FROM emp WHERE empno = 1;</pre> This returns 123.34 and 123.35. <pre>SELECT ROUND(748.58, -1), ROUND(748.58, -2), ROUND(748.58, 3), FROM emp WHERE empno = 1;</pre> This returns 750, 700 and 1000.
SIGN(<i>numeric_exp</i>)	Returns the positive (+1), zero (0), or negative (-1) sign of the given expression. For example: <pre>SELECT SIGN(empno) FROM emp WHERE empno = 11;</pre> This returns 1. <pre>SELECT SIGN(-1 * empno), SIGN(0) FROM emp WHERE empno = 1;</pre> This returns two result columns with values -1 and 0.
SIN(<i>float_exp</i>)	Returns the sine of <i>float_exp</i>, where <i>float_exp</i> is an angle expressed in radians. For example: <pre>SELECT SIN(1.570796) FROM emp WHERE empno = 11;</pre> This returns 0.999999.
SQRT(<i>float_exp</i>)	Returns the square root of <i>float_exp</i>. For example: <pre>SELECT SQRT(45.35) FROM emp WHERE empno = 11;</pre> This returns 6.7342.
TAN(<i>float_exp</i>)	Returns the tangent of <i>float_exp</i>, where <i>float_exp</i> is an angle expressed in radians. For example: <pre>SELECT TAN(0.785398) FROM emp WHERE empno = 11;</pre> This returns 0.999999.
TRUNCATE(<i>numeric_exp</i> , <i>integer_exp</i>)	Returns <i>numeric_exp</i> truncated to <i>integer_exp</i> places right of the decimal point. If <i>integer_exp</i> is negative, <i>numeric_exp</i> is truncated to $integer\_exp$ places to the left of the decimal point.
NCHAR(<i>code</i>)	Returns the Unicode character that has the specified code as a SQL_WCHAR value. The value of code should be between 0 and 65535. Example: "SELECT NCHAR(945)" returns the character d.

String Functions

Table 14–4 String Functions

Function	Description
ASCII(<i>string_exp</i>)	Returns the ASCII code value of the leftmost character of <i>string_exp</i> as an integer. For example: <pre>SELECT ASCII(ename) FROM emp WHERE ename = 'Adam';</pre> This returns 65, which is the ASCII code of A.

Table 14–4 (Cont.) String Functions

Function	Description
BIT_LENGTH(<i>string_exp</i>)	Returns the length in bits of the string expression. For example: SELECT BIT_LENGTH(ename) FROM emp WHERE ename = 'John'; This returns 32, which is the number of bits.
CHAR(<i>code</i>)	Returns the character that has the ASCII code value specified by code. The value of code should be between 0 and 255; otherwise, the return value is data source-dependent. For example: SELECT CHAR(65) FROM emp; This returns A, which is the character for ASCII code A.
CHAR_LENGTH(<i>string_exp</i>) CHARACTER_LENGTH(<i>string_exp</i>)	Returns the length in characters of the string expression, if the string expression is of a character data type; otherwise, returns the length in bytes of the string expression (the smallest integer not less than the number of bits divided by 8). (This function is the same as the CHARACTER_LENGTH function.) For example: SELECT CHAR_LENGTH(ename) FROM emp where ename = 'John'; This returns 4.
CONCAT(<i>string_exp1</i> , <i>string_exp2</i>)	Returns a character string that is the result of concatenating string_exp2 to string_exp1. If either of string_exp1 or string_exp2 is NULL value, it returns NULL string. If either of string_exp1 or string_exp2 is wide character string, the return value is a wide character string. For example: SELECT CONCAT('Name is: ', ename) FROM emp WHERE ename = 'John'; This returns 'Name is: John' SELECT CONCAT(N'Name is: ', ename) FROM emp WHERE ename = N'John'; This returns wide character string N'Name is: John'.
INSERT(<i>string_exp1</i> , <i>start</i> , <i>length</i> , <i>string_exp2</i>)	Returns a character string where length characters have been deleted from string_exp1, beginning at start, and where string_exp2 has been inserted into string_exp, beginning at start. If string_exp1 is wide character string, the return value is a wide character string. Offsets (start and length) must be specified in number of characters. For example: SELECT INSERT(ename, 1, 0, 'Name is: ') FROM emp WHERE ename = 'John'; This returns 'Name is: John' SELECT INSERT(ename, 1, 0, N'Name is: ') FROM emp WHERE ename = N'John'; If ename is a column of wide character data type, this returns wide character string: N'Name is: John'.
LCASE(<i>string_exp</i>) LOWER(<i>string_exp</i>)	Returns a string equal to that in string_exp, with all uppercase characters converted to lowercase. For example: SELECT LCASE(ename) FROM emp WHERE ename is 'John'; This returns 'john'.

Table 14–4 (Cont.) String Functions

Function	Description
LEFT(<i>string_exp</i> , <i>count</i>)	Returns the leftmost count characters of <i>string_exp</i>. If <i>string_exp</i> is wide character string, the return value is a wide character string. Offset (<i>count</i>) must be specified in number of characters. For example: <pre>SELECT LEFT(ename, 2) FROM emp WHERE ename = 'John';</pre> This returns 'jo'. <pre>SELECT LEFT(ename, 2) FROM emp WHERE ename = N'John';</pre> If <i>ename</i> is a column of wide character data type, this returns wide character string N'jo'.
LENGTH(<i>string_exp</i>)	Returns the number of characters in <i>string_exp</i>, excluding trailing blanks. If <i>string_exp</i> is wide character string, the return value is a number of wide characters in <i>string_exp</i>. Trailing blanks are not checked in wide character implementation. For example: <pre>SELECT LENGTH('John '), LENGTH('John') FROM emp;</pre> This returns 4 for both result columns as trailing blanks are excluded. <pre>SELECT LENGTH(N'John '), LENGTH(N'John') FROM emp;</pre> This returns 7 for the first result column and 4 for the second result column. Trailing blanks are not checked in wide character implementation.
LOCATE(<i>string_exp1</i> , <i>string_exp2</i> [, <i>start</i>])	Returns the starting position of the first occurrence of <i>string_exp1</i> within <i>string_exp2</i>. The search for the first occurrence of <i>string_exp1</i> begins with the first character position in <i>string_exp2</i> unless the optional argument, <i>start</i>, is specified. If <i>start</i> is specified, the search begins with the character position indicated by the value of <i>start</i>. The first character position in <i>string_exp2</i> is indicated by the value 1. If <i>string_exp1</i> is not found within <i>string_exp2</i>, the value 0 is returned. If <i>string_exp2</i> is a wide character string, returns the starting position of the first occurrence of <i>string_exp1</i> within the wide character string <i>string_exp2</i>. Offset (<i>start</i>) must be specified in number of characters. If <i>string_exp2</i> is a wide character exp, the result is computed by treating both arguments as wide character string. For example: <pre>SELECT LOCATE('h', 'John', 1) FROM emp;</pre> This returns 3 as 'h' is the found at the third position. <pre>SELECT LOCATE(N'h', N'John', 1) FROM emp;</pre> This returns 3 as N'h' is the found at the third position.
LTRIM(<i>string_exp</i>)	Returns the characters of <i>string_exp</i>, with leading blanks removed. For example: <pre>SELECT LTRIM(' ABC') FROM emp;</pre> This returns 'ABC'.
OCTET_LENGTH(<i>string_exp</i>)	Returns the length in bytes of the string expression. The result is the smallest integer not less than the number of bits divided by 8. For example: <pre>SELECT OCTET_LENGTH(ename) FROM emp WHERE ename = 'John';</pre> This returns 4.

Table 14–4 (Cont.) String Functions

Function	Description
<code>POSITION(character_exp1 character_exp2)</code>	Returns the position of the first character expression in the second character expression. The result is an exact numeric with an implementation-defined precision and a scale of 0. If <code>character_exp1</code> and <code>character_exp2</code> are wide character strings, returns the position of the first wide character expression in the second wide character expression. If <code>character_exp2</code> is a wide character string, the result is computed by treating both arguments as wide character strings. For example: <pre>SELECT POSITION('abc', '1234abc def') FROM emp;</pre> This returns 5. <pre>SELECT POSITION(N'abc', N'1234abc def') FROM emp;</pre> This returns 5.
<code>REPEAT(string_exp, count)</code>	Returns a character string composed of <code>string_exp</code> repeated <code>count</code> times. If <code>string_exp</code> is wide character string, the return value is a wide character string. For example: <pre>SELECT REPEAT(ename, 2) FROM emp WHERE ename = 'John';</pre> This returns 'JohnJohn' <pre>SELECT REPEAT(ename, 2) FROM emp WHERE ename = N'John';</pre> If <code>ename</code> is a column of wide character data type, this returns N'JohnJohn'.
<code>REPLACE(string_exp1, string_exp2, string_exp3)</code>	Search <code>string_exp1</code> for occurrences of <code>string_exp2</code>, and replace with <code>string_exp3</code>. If <code>string_exp1</code> is wide character string, the return value is a wide character string. For example: <pre>SELECT REPLACE(address, 'San Francisco', 'SFO') FROM emp where address = '100 Vanness, San Francisco';</pre> This returns '100 Vanness, SFO'. <pre>SELECT REPLACE(address, N'San Francisco', N'SFO') FROM emp WHERE address = N'100 Vanness, San Francisco';</pre> If <code>address</code> is a column of wide character data type, this returns N'100 Vanness, SFO'.
<code>RIGHT(string_exp, count)</code>	Returns the right-most <code>count</code> characters of <code>string_exp</code>. If <code>string_exp</code> is wide character string, the return value is a wide character string. Offset (<code>count</code>) must be specified in number of characters. For example: <pre>SELECT RIGHT(ename, 2) FROM emp WHERE ename = 'John';</pre> This returns 'hn'. <pre>SELECT RIGHT(ename, 2) FROM emp WHERE ename = N'John';</pre> If <code>ename</code> is a column of wide character data type, this returns N'hn'.
<code>RTRIM(string_exp)</code>	Returns the characters of <code>string_exp</code> with trailing blanks removed. For example: <pre>SELECT RTRIM('abc ') FROM emp;</pre> This returns 'abc'.

Table 14–4 (Cont.) String Functions

Function	Description
SPACE(<i>count</i>)	Returns a character string consisting of <i>count</i> spaces. For example: SELECT ename+space(5)+ename FROM emp WHERE ename = 'John'; This returns 'John John'.
SUBSTRING(<i>string_exp</i> , <i>start</i> , <i>length</i>) SUBSTR(<i>string_exp</i> , <i>length</i>)	Returns a character string that is derived from <i>string_exp</i> , beginning at the character position specified by <i>start</i> for <i>length</i> characters. If <i>string_exp</i> is wide character string, the return value is a wide character string. Offset (<i>start</i> and <i>length</i>) must be specified in number of characters. For example: SELECT SUBSTR(ename, 1, 3) FROM emp WHERE ename = 'John'; This returns 'Joh' SELECT SUBSTR(ename, 1, 3) FROM emp WHERE ename = N'John'; If <i>ename</i> is a column of wide character data type, this returns N'Joh'
UCASE(<i>string_exp</i>) UPPER(<i>string_exp</i>)	Returns a string equal to that in <i>string_exp</i> , with all lowercase characters converted to uppercase. For example: SELECT UCASE(ename) FROM emp WHERE ename = 'John'; This returns 'JOHN'.
UNICODE(<i>string_exp</i>)	Returns the Unicode code of the first character of the <i>string_exp</i> as a SQL_INTEGER value. Example: "SELECT UNICODE('dej')" returns an integer value of 945.

Time and Date Functions

Table 14–5 Time and Date Functions

Function	Description
CURDATE()	Returns the current date. For example: SELECT CURDATE() FROM emp; Returns the current date as: 2008-10-25
CURTIME()	Returns the current local time. For example: SELECT CURTIME() FROM emp; Returns the current time as: 10:20:05
CURTIMESTAMP()	Returns the current local date and local time as a timestamp value. For example: SELECT CURTIMESTAMP() FROM emp; Returns current date and time as: 2003-03-31 14:08:57

Table 14–5 (Cont.) Time and Date Functions

Function	Description
DATEADD(<i>datepart</i> , <i>number</i> , <i>date</i>)	Returns a new date time value based on adding an interval to the specified date. The return date-time data type is same as the input <i>date</i> value.
TIMESTAMPADD(<i>datepart</i> , <i>number</i> , <i>date</i>)	<i>datepart</i>: the parameter that specifies on which part of the date to return a new value. Both ODBC notation and SQL Server notation for <i>datepart</i> are supported. Datepart Abbreviations Year SQL_TSI_YEAR , year, yy, yyyy, quarter SQL_TSI_QUARTER, quarter, qq, q Month SQL_TSI_MONTH, month, mm, m dayofyear DAYOFYEAR, dy, y Day SQL_TSI_DAY, day, dd, d Week SQL_TSI_WEEK, week, wk, ww Hour SQL_TSI_HOUR, hour, hh minute SQL_TSI_MINUTE , minute, mi, n second SQL_TSI_SECOND, second, ss, s The current implementation does not support millisecond and fractional second specifications. <i>number</i>: the value used to increment the <i>datepart</i>. If value is not an integer, the fractional part of the value is discarded. For example, if you specify day for <i>datepart</i> and 1.75 for <i>number</i>, <i>date</i> is incremented by 1. <i>date</i>: an expression that returns a date or timestamp value or a character string in a date-time format.

Table 14–5 (Cont.) Time and Date Functions

Function	Description
DATEDIFF(<i>datepart</i> , <i>startdate</i> , <i>enddate</i>) TIMESTAMPDIFF(<i>datepart</i> , <i>startdate</i> , <i>enddate</i>)	Returns the number of date and time boundaries crossed between two specified dates. <i>startdate</i> is subtracted from <i>enddate</i>. If <i>startdate</i> is later than <i>enddate</i>, a negative value is returned. <i>datepart</i>: the parameter that specifies on which part of the date to calculate the difference. Both ODBC notation and SQLServer notation for <i>datepart</i> are supported. Datepart Abbreviations Year SQL_TSI_YEAR, year, <i>yy</i>, <i>yyyy</i>, quarter SQL_TSI_QUARTER, quarter, <i>qq</i>, <i>q</i> Month SQL_TSI_MONTH, month, <i>mm</i>, <i>m</i> dayofyear DAYOFYEAR, <i>dy</i>, <i>y</i> Day SQL_TSI_DAY, day, <i>dd</i>, <i>d</i> Week SQL_TSI_WEEK, week, <i>wk</i>, <i>ww</i> Hour SQL_TSI_HOUR, hour, <i>hh</i> minute SQL_TSI_MINUTE, minute, <i>mi</i>, <i>n</i> second SQL_TSI_SECOND, second, <i>ss</i>, <i>s</i> The current implementation does not support dayofyear, millisecond and fractional second specifications. For example: <pre>SELECT DATEDIFF(year, hiredate, curdate()) FROM emp WHERE hiredate = '2000-10-01'; SELECT DATEDIFF(SQL_TSI_YEAR, hiredate, curdate()) FROM emp WHERE hiredate = '2000-10-01';</pre> This returns 2 (assuming that the <i>curdate()</i> returns year 2002).
DAYNAME(<i>date_exp</i>)	Returns a character string containing the data source-specific name of the day (for example, Sunday through Saturday or Sun. through Sat. for a data source that uses English, or Sonntag through Samstag for a data source that uses German) for the day portion of <i>date_exp</i>. For example: <pre>SELECT DAYNAME('2002-01-01'), DAYNAME('2002-01-02') FROM emp;</pre> This returns 'Tuesday' and 'Wednesday'.
DAYOFMONTH(<i>date_exp</i>)	Returns the day of the month based on the month field in <i>date_exp</i> as an integer value in the range of 1-31. For example: <pre>SELECT DAYOFMONTH('2002-01-05') FROM emp;</pre> This returns 5.
DAYOFWEEK(<i>date_exp</i>)	Returns the day of the week based on the week field in <i>date_exp</i> as an integer value in the range of 1-7, where 1 represents Sunday. For example: <pre>SELECT DAYOFWEEK('2002-01-05') FROM emp;</pre> This returns 7.
DAYOFYEAR(<i>date_exp</i>)	Returns the day of the year based on the year field in <i>date_exp</i> as an integer value in the range of 1-366. For example: <pre>SELECT DAYOFYEAR('2002-01-05') FROM emp;</pre> This returns 5.

Table 14–5 (Cont.) Time and Date Functions

Function	Description
HOUR(<i>time_exp</i>)	Returns the hour based on the hour field in <i>time_exp</i> as an integer value in the range of 0-23. For example: SELECT HOUR('22:20:20') FROM emp; This returns 22.
MINUTE(<i>time_exp</i>)	Returns the minute based on the minute field in <i>time_exp</i> as an integer value in the range of 0-59. For example: SELECT MINUTE('22:21:20') FROM emp; This returns 21.
MONTH(<i>date_exp</i>)	Returns the month based on the month field in <i>date_exp</i> as an integer value in the range of 1-12. For example: SELECT MONTH('2002-01-05') FROM emp; This returns 1.
MONTHNAME(<i>date_exp</i>)	Returns a character string containing the data source-specific name of the month (for example, January through December or Jan. through Dec. for a data source that uses English, or January through Dezember for a data source that uses German) for the month portion of <i>date_exp</i> . For example: SELECT MONTHNAME('2002-01-05') FROM emp; This returns January.
NOW()	Returns current date and time as a timestamp value. For example: SELECT NOW() FROM emp; This returns the current date and time: 2002-10-25 10:20:05.
QUARTER(<i>date_exp</i>)	Returns the quarter in <i>date_exp</i> as an integer value in the range of 1-4, where 1 represents January 1 through March 31. For example: SELECT QUARTER('2002-01-05') FROM emp; This returns 1.
SECOND(<i>time_exp</i>)	Returns the second based on the second field in <i>time_exp</i> as an integer value in the range of 0-59. For example: SELECT SECOND('22:21:20') FROM emp; This returns 20.
Returns the week of the year based on the week field in <i>date_exp</i> as an integer value in the range of 1-53. For example: SELECT WEEK('2002-01-05') FROM emp; This returns 1.	WEEK(<i>date_exp</i>)
YEAR(<i>date_exp</i>)	Returns the year based on the year field in <i>date_exp</i> as an integer value. For example: SELECT YEAR('2002-01-01') FROM emp; This returns 2002.

System Functions

Table 14–6 System Functions

Function	Description
DATABASE()	Returns the name of the database corresponding to the connection handle. (The name of the database is also available by calling SQLGetConnectOption with the SQL_CURRENT_QUALIFIER connection option.)
USER()	Returns the user name in the DBMS. (The user name is also available using SQLGetInfo by specifying the information type: SQL_USER_NAME.) This can be different than the login name.

Aggregate Functions

Aggregate functions return a single row based on groups of rows, rather than on single rows.

Table 14–7 Aggregate Functions

Function	Description
AVG([ALL DISTINCT] <i>expression</i>)	Returns the average of the values in a group. Null values are ignored.
SUM([ALL DISTINCT] <i>expression</i>)	Returns the sum of all the values, or only the DISTINCT values, in the expression. SUM can be used with numeric columns only. Null values are ignored.
COUNT([([ALL DISTINCT] <i>expression</i>) *])	Returns the number of items in a group. COUNT(*) returns the number of items in a group, including NULL values and duplicates. COUNT(ALL <i>expression</i>) evaluates <i>expression</i> for each row in a group and returns the number of non-null values. COUNT(DISTINCT <i>expression</i>) evaluates <i>expression</i> for each row in a group and returns the number of unique, non-null values.
MAX([ALL DISTINCT] <i>expression</i>)	Returns the maximum value in the expression.
MIN([ALL DISTINCT] <i>expression</i>)	Returns the minimum value in the expression.

Aggregate functions can appear in SELECT lists and HAVING clauses. If you use the GROUP BY clause in a SELECT statement, OpenAccess SDK divides the rows of a queried table or view into groups. In a query containing a GROUP BY clause, all elements of the SELECT list must be expressions from the GROUP BY clause, expressions containing aggregate functions, or constants. OpenAccess SDK applies the aggregate functions in the SELECT list to each group of rows and returns a single result row for each group.

If you omit the GROUP BY clause, OpenAccess SDK applies aggregate functions in the SELECT list to all the rows in the queried table or view.

Many aggregate functions accept these options:

- DISTINCT causes an aggregate function to consider only distinct values of the argument expression.
- ALL causes an aggregate function to consider all values, including all duplicates.

For example:


```
SELECT max(sal), MIN(sal), AVG(sal) FROM emp;
SELECT deptno, MAX(sal), SUM(sal) FROM emp GROUP BY deptno;
SELECT deptno, COUNT(empno) FROM emp GROUP BY deptno;
```

Other Functions

DECODE

Syntax

```
DECODE (expr, [search, result]..., default)
```

Example

```
SELECT DECODE (deptno,10, 'ACCOUNTING',
               20, 'RESEARCH',
               30, 'SALES',
               40, 'OPERATION',
               'NONE')
FROM dept
```

To evaluate this expression, the OpenAccess SDK SQL engine compares expr to each search value one by one. If expr is equal to a search, the OpenAccess SDK SQL engine returns the corresponding result. If no match is found, the OpenAccess SDK SQL engine returns default, or if default is omitted, returns null. The return value is the same data type as the first result expression. The search, result, and default values can be derived from expressions.

The OpenAccess SDK SQL engine evaluates each search value only before comparing it to expr, rather than evaluating all search values before comparing any of them with expr. Consequently, OpenAccess SDK SQL engine never evaluates a search if a previous search is equal to expr.

The OpenAccess SDK SQL engine automatically converts expr and each search value to the datatype of the first search value before comparing. The OpenAccess SDK SQL engine automatically converts the return value to the same datatype as the first result. If the first result has the datatype CHAR or if the first result is null, then the OpenAccess SDK SQL engine converts the return value to the datatype of CHAR.

In a DECODE expression, the OpenAccess SDK SQL engine considers two nulls to be equivalent.

If expr is null, the OpenAccess SDK SQL engine returns the result of the first search that is also null. The maximum number of components in the DECODE expression, including expr, searches, results, and default is 255.

Examples:

```
SELECT DECODE(empno, 1, 'E1', 2, 'E2', 'DEFAULT') FROM emp;

# First Result expression is NULL. Result should be type
XO_TYPE_CHAR
SELECT DECODE(empno, 1, NULL, 2, 'E2', 'DEFAULT') FROM emp;

# Input expression is NULL, Result should match the result
of NULL search expr
SELECT DECODE(ename, 'Bob', 'My Bob', 'Mary', 'My Mary',
NULL, 'New Name', 'Default Name') FROM emp;

# no default value, so return NULL for non-match values
SELECT DECODE(empno, 1, 'E1', 2, 'E2') FROM emp;
```

```
# ERROR CHECKING
# Invalid number of arguments
# Invalid syntax used with scalar function:DECODE. Function
expects 3 arguments.
SELECT DECODE() FROM emp;
SELECT DECODE(empno, 1) FROM emp;

# Conversion errors
# decode() Error converting value of result expression to
XoType:<4>

SELECT DECODE(empno, 1, 10, 2, 20, 'abc') FROM emp;
```

IFNULL, ISNULL, NVL

These functions allow the NULL value to be replaced by a default value. OpenAccess SDK supports IFNULL as defined by ODBC, ISNULL as defined by SQL Server, and NVL as defined by Oracle.

Syntax

```
IFNULL (expr, default_val)
ISNULL (expr, default_val)
NVL (expr, default_val)
```

The OpenAccess SDK SQL engine evaluates the input expression and returns the expression value if it is non-NULL. If the expression value is NULL, default_val is returned. The return value is of the same data type as the input expression.

Example

```
SELECT ename, IFNULL (sal, 1000) FROM emp;
SELECT ename, ISNULL (sal, 1000) FROM emp;
SELECT ename, NVL (sal, 1000) FROM emp;

SELECT ename, IFNULL (hiredate, '2001-01-01') FROM emp;
```

CAST

Syntax

```
CAST (value_exp AS data_type)
```

Example

```
SELECT empno, CAST(empno AS VARCHAR) FROM emp
SELECT empno, CAST(empno AS SMALLINT) FROM emp
```

The function returns the value specified by value_exp converted to the specified *data_type*, where *data_type* is one of the following:

- CHAR
- NUMERIC
- DECIMAL
- INTEGER
- SMALLINT
- FLOAT

- REAL
- DOUBLE
- DATE
- TIME
- TIMESTAMP
- VARCHAR
- LONGVARCHAR
- BINARY
- VARBINARY
- LONGVARBINARY
- TINYINT
- BIT
- WCHAR,
- WVARCHAR
- WLONGVARCHAR

[Table 14–8](#) defines the precision, length, and scale keywords of the CAST function.

Table 14–8 *CAST Function*

Keyword	Value
CHAR	255
BINARY	255
BIT	1
DATE	6
DOUBLE	8
FLOAT	8
INTEGER	4
LONGVARBINARY	1000000
LONGVARCHAR	1000000
NUMERIC	34
SMALLINT	2
REAL	4
TIME	6
TIMESTAMP	16
TINYINT	1
VARBINARY	1024
VARCHAR	1024
WLONGVARCHAR	2000000
WVARCHAR	512
WVARCHAR	2048

Table 14–8 (Cont.) CAST Function

Keyword	Value
CHAR, BINARY	255
DATE	10
DOUBLE	15
FLOAT	15
INTEGER	10
LONGVARBINARY	1000000
LONGVARCHAR	1000000
NUMERIC	32
REAL	7
SMALLINT	5
TIME	6
TINYINT	3
VARBINARY	1024
VARCHAR	1024
WCHAR	255
WVARCHAR	1024
WLONGVARCHAR	1000000
NUMERIC	5
All other types	0

CONVERT**Syntax**

CONVERT (value_exp, data_type)

Example

```
SELECT empno, CONVERT(empno, SQL_VARCHAR) FROM emp
SELECT empno, CONVERT(empno, SQL_SMALLINT) FROM emp
```

The function returns the value specified by value_exp converted to the specified data_type, where data_type is one of the following:

- SQL_CHAR
- SQL_NUMERIC
- SQL_DECIMAL
- SQL_INTEGER
- SQL_SMALLINT
- SQL_FLOAT
- SQL_REAL
- SQL_DOUBLE
- SQL_DATE

- SQL_TIME
- SQL_TIMESTAMP
- SQL_VARCHAR
- SQL_LONGVARCHAR
- SQL_BINARY
- SQL_VARBINARY
- SQL_LONGVARBINARY
- SQL_TINYINT
- SQL_BIT
- SQL_WCHAR
- SQL_WVARCHAR
- SQL_WLONGVARCHAR

Table 14–9, Table 14–10, and Table 14–11 define the length, precision, and scale keywords of the CONVERT function.

Table 14–9 CONVERT Function-Length

Keyword	Length
SQL_CHAR	256
SQL_BINARY	256
SQL_BIT	1
SQL_DATE	6
SQL_DOUBLE	8
SQL_FLOAT	8
SQL_INTEGER	4
SQL_LONGVARBINARY	1000000
SQL_LONGVARCHAR	1000000
SQL_NUMERIC	34
SQL_SMALLINT	2
SQL_REAL	4
SQL_TIME	6
TIMESTAMP	16
TINYINT	1
VARBINARY	1024
VARCHAR	1024
WLONGVARCHAR	2000000
WVARCHAR	512
WVARCHAR	2048

Table 14–10 CONVERT Function-Precision

Keyword	Precision
SQL_BINARY	255
SQL_BIT	1
SQL_CHAR	255
SQL_DATE	10
SQL_DOUBLE	15
SQL_FLOAT	15
SQL_INTEGER	10
SQL_LONGVARBINARY	1000000
SQL_LONGVARCHAR	1000000
SQL_NUMERIC	32
SQL_REAL	7
SQL_SMALLINT	5
SQL_TIME	8
SQL_TINYINT	3
SQL_VARBINARY	1024
SQL_VARCHAR	1024
SQL_WCHAR	255
SQL_WVARCHAR	2048
SQL_WLONGVARCHAR	1000000

Table 14–11 CONVERT Function-Scale

Keyword	Scale
SQL_NUMERIC	5
All other types	0

Unsupported SQL Functions

The SQL engine of the RPAS ODBC Server implements a large portion of the entry level SQL as defined in the X3.135-1992, "Database Language SQL" specification and commercial databases like SQL Server and Oracle. It is compliant with the ODBC minimal grammar specification.

This section describes the unsupported features.

Handling of NULLS

NOT IN should return FALSE if any member of the set is NULL. When evaluating the IN condition, the OpenAccess SDK SQL engine treats the comparison of any value with NULL as FALSE, so NOT IN will become TRUE.

```
SELECT * FROM emp WHERE job NOT IN
(SELECT job FROM emp WHERE job IS NULL)
```

This example should return no results if there is an emp record with a NULL value for Job.

Schema Information

The SQL engine of the RPAS ODBC Driver does not support the following:

- Collate sequence and character set
- DEFAULT clause for column values

Data Definition Language (DDL)

The only DDL that is supported is "create view".

Insert

- Insert statements are not supported.
- Update measure data on fact table is supported.

SELECT Syntax

Subqueries are not supported in a SELECT list.

Example:

SELECT

(SELECT a.empno FROM emp a WHERE a.deptno = b.deptno)

FROM

dept b

Value Expressions

Special Values: The SQL engine does not support the use of special values (CURRENT_USER, SESSION_USER, CURRENT_TIMESTAMP) in value specification.

Value Functions

The SDK SQL engine does not support the following functions:

- TRANSLATE
- TRIM

Note: The OpenAccess SDK SQL engine supports LTRIM and RTRIM.)

- DIFFERENCE
- SOUNDEX

Date and Time Functions

The OpenAccess SDK SQL engine does not support the following Date and Time functions:

- CURRENT_TIME[(time-precision)] - The SQL engine does not support time-precision argument.
- CURRENT_TIMESTAMP[(time-precision)]

- `EXTRACT(extract-field FROM extract-source)`

Advanced Value Expressions

NULLIF

NULLIF is shorthand for a frequently used variation of CASE.

Syntax

```
NULLIF(value1, target_value)
is equivalent to
CASE
    WHEN value1 = target_value THEN NULL
ELSE value1
END
Example:
.. WHERE sales_revenue / NULLIF(our_cost, -1) > 50
```

COALESCE

Coalesce is shorthand for a frequently used variation of CASE.

Syntax:

```
COALESCE (value1, value2, value3)
is equivalent to:
CASE
    WHEN value1 IS NOT NULL THEN value1
    WHEN value2 IS NOT NULL THEN value2
    ELSE value3
END
```

Example:

```
SELECT name, job_title, COALESCE (salary, commission,
subsistence)
FROM job_assignments
```

Row Value Constructor

A row value constructor is a parenthesized list of values.

Example:

The following expression:

```
WHERE c1=CA AND c2=CB AND c3=CC
```

can be written using row value constructor as:

```
WHERE (c1, c2, c3) = (CA, CB, CC)
```

Predicates

The SQL engine does not support the following predicates.

- `OVERLAPS` predicate: Determines whether two intervals of time overlap with one another.event-information `OVERLAPS` event-information
- `MATCH` predicate

Join Operators

This section explains which Join operations are supported by the RPAS SQL engine, and which are not.

Supported Join Operators

The OpenAccess SDK SQL engine supports the following join operations:

- Implicit JOIN. The WHERE clause explicitly specifies the join condition.
- INNER JOIN. All joins that are not OUTER JOINS are considered in SQL terminology as INNER joins. The use of keyword INNER has no additional effects, but helps the statement to be completely self-documenting.
 - `SELECT * FROM t1 INNER JOIN t2 ON t1.c1 = t2.c3 WHERE search-LEFT OUTER JOIN` - This join preserves unmatched rows from the left table.
 - `SELECT * FROM t1 LEFT OUTER JOIN t2 ON t1.c1 = t2.c3 WHERE search-condition`
- UNION and UNION ALL operators. UNION is used for combining two result tables that are union compatible.

```
SELECT c1, c2 FROM t1 UNION SELECT c3, c4 FROM t2
```

Unsupported Join Operators

The following join operations (syntax) are not supported in this release:

- CROSS JOIN: Functionally similar to the implicit joins.


```
SELECT * FROM t1 CROSS JOIN t2
```
- NATURAL JOIN: Also referred to as natural, equi-join selects rows from the tables that have same value for columns with the same name.


```
SELECT * FROM t1 NATURAL JOIN t2
```
- Condition JOIN: Uses the keyword ON to specify the JOIN condition between tables. The scope of fields referred in the ON condition is restricted.


```
SELECT * FROM t1 JOIN t2 ON t1.c1 = t2.c3 WHERE search-condition
```
- Column Name JOIN: Specifies a more restricted form of NATURAL join. NATURAL joins use all columns with the same names to manage the matching process. The column name JOIN specifies which column values should be matched.


```
SELECT * FROM t1 JOIN t2 USING (c1, c2)
```
- RIGHT OUTER JOIN: Preserves unmatched rows from the right table.


```
SELECT * FROM t1 RIGHT OUTER JOIN t2 ON t1.c1 = t2.c3 WHERE search-condition
```
- FULL OUTER JOIN: Preserves unmatched rows from both the left and right tables.


```
SELECT * FROM t1 FULL OUTER JOIN t2 ON t1.c1 = t2.c3 WHERE search-condition
```
- UNION JOIN: Creates a new virtual table with the union of all columns from the source tables. The UNION join has no provision for column matching.

Publishing Measure Change Events


Event-driven planning requires the ability to identify events when they arise. This includes events that result from changes in plans and those that arise because of advancement in the planning activity, for example, approval of a plan or creation of new items. The ability to get notification of the event when it occurs is therefore essential. In the context of RPAS applications, many measure changes result from business activities and therefore fall into the category of notification-worthy events.

RPAS Publishing Measure Changes (PMC) provides a mechanism to monitor measure changes and receive notification messages through a standard JMS messaging service. A measure change event is defined as the measure data being written by any means. The following is the list of the sources for PMC events:

- MACE: all left-hand side measures in the expression or expressions of the rule group specified in the command line.
- Workbook Commit: all left-hand side measures in the commit rule group.
- loadMeasure: all successfully loaded measures.
- loadHier: measures whose base intersection contains at least one affected dimension. "Affected" means adding and deleting one or more logical positions.
- Dynamic Position Management (DPM): measures whose base intersection contains a dimension affected by DPM. Adding or removing positions to a dimension or its lower level children constitutes a change in the measure.

Figure 15–1 demonstrates a JMS system in the context of RPAS. The JMS system consists of four basic components:

- JMS Provider: The central daemon that accepts connections from clients and routes and queues messages. It is sometimes called JMS Broker.
- Java Naming and Directory Interface (JNDI) Service: A service that provides mapping between names and objects. It adds abstraction to the more complex structure of underlying objects.
- JMS Client (Publisher): A standalone daemon process running with the other RPAS server side processes such as DomainDaemon and RpasDbServer.
- JMS Client (Subscriber): Any applications interested in getting notifications of measure changes.

Figure 15–1 JMS System in RPAS

Configuring Subjects and Measures for Monitoring

A subject is defined as a logical grouping of measures. A subject is mapped to one or more measures in one domain, while a measure can belong to one or more subjects. For example, a subject named PlanningMeasures can include all measures associated with planning.

The configuration file of measure-change monitoring serves two purposes:

- Defines the mappings between subjects and measures.
- Defines the inclusion filters for monitored measures.

By default, no measures are monitored unless they are specifically included in the configuration.

For each domain, there is one measure-change monitor configuration file. It is named `MeasureChangeMonitor.properties` and must be located under `{domainPath}/config/`.

Note: There is only one property file for a global domain. It is placed under the config directory of the master domain.

This configuration file is in standard Java properties file format. Each line defines the relationship between one subject and one measure:

```
Subject.MeasureName=true
```

- `true` means inclusion in monitoring.
- `false` is used to exclude a subject or measure from monitoring. Use of `false` is not recommended since, by default, a measure is not monitored and such a line can be deleted or commented out from the file. Comments are any contents on a line from `"#"` to the end of the line.

Example:

```
TestMeasures.R_EX_DEMOA=true
TestMeasures.R_EX_DEMOB=true
LanguageMeasures.R_MsgLabel=true
```

The following are defined in the above example:

- Two subjects: "TestMeasures" and "LanguageMeasures".
- The subject "TestMeasures" represents two measures: "R_EX_DEMOA" and "R_EX_DEMOB".
- The subject "LanguageMeasures" represents one measure: "R_MsgLabel".

A subject name must be a valid hierarchical variable name, that is, consisting of only alphanumeric characters, underscores, and periods. RPAS does not enforce any naming convention for a subject. It is up to the retailer to define an individual naming convention.

After the configuration file is modified, RPAS processes can detect that the file has changed and automatically reload it. There is no need to restart any RPAS processes.

Configuring the RPAS JMS Publisher

JMS Publisher for measure change events is implemented in Java and runs as a standalone process. It is decoupled from any other RPAS server side processes in terms of inter-process communications and domain data file locking.

Each JMS Publisher process is tied to a domain and a JMS topic. When the publisher detects an event for its domain, it generates a JMS message and sends the message to a JMS provider. The format of the JMS message is defined by a template file, which can contain any of the macros listed in [Table 15–1](#). The macros are replaced by actual values at run time.

Table 15–1 Configuration Macros

Macro Name	Format	Notes
__EventDateTime__	YYYY-MM-DDThh:mm:ss	Local time of the server.
__SourceURI__	RPAS/JMS/{hostname}	{hostname} is the name of the server where the publisher is running.
__SUBJECT__	String	Subject of the event as defined in MeasureChangeMonitor.properties.
__TYPE__	"MeasureChange"	Type is a constant string.
__DOMAIN__	String	Path to the domain.
__MEASURE__	String	RPAS internal measure name.
__ORIGINUSER__	String	RPAS User ID, only available for workbook commit or DPM. Use "-" if not available.

Here are two examples of a JMS message template.

Example 1: Simple name/value pairs:

```
type=__TYPE__
domain=__DOMAIN__
measure=__MEASURE__
time=__EventDateTime__
```

```
user=__ORIGINUSER__
```

Example 2: XML-based Notification Event Architecture for Retail (NEAR) format:

```
<?xml version="1.0" encoding="UTF-8" ?>
<AlertEvent MajorVersion="1" MinorVersion="0" TypeCode="RPASEvent"
  Priority="0" Severity="Information" Mode="Test" FixVersion="0"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"

xsi:schemaLocation="http://www.retail.oracle.com/workspace/alerts/AlertEventV1.0.0
.xsd"
  xmlns="http://www.retail.oracle.com/workspace/alerts/">
  <SequenceNumber>0</SequenceNumber>
  <EventDateTime>__EventDateTime__</EventDateTime>
  <EventDescription>RPAS Measure Change Event</EventDescription>
  <SourceName>RPAS</SourceName>
  <SourceURI>__SourceURI__</SourceURI>
  <Instance>1</Instance>
  <RoutingInfo TypeCode="SubjectInfo">__SUBJECT__</RoutingInfo>
  <AlertData><![CDATA[<type>__TYPE__</type>
<domain>__DOMAIN__</domain>
<measure>__MEASURE__</measure>
<originUser>__ORIGINUSER__</originUser>]]></AlertData>
</AlertEvent>
```

Command Line

The following command line is used for JMS Publisher:

```
java -cp {classpath} oracle.rpas.pmc.MCPublisher -d {domainPath} [-c
{configFileName}] [name1=value1 [name2=value2 [...]]]
```

Table 15–2 provides descriptions of the arguments used in the command line.

Table 15–2 Arguments Used in the Command Line

Argument	Description
classpath	Should include rpaspmc.jar and jms.jar under \$RPAS_HOME/lib and any vendor-specific JMS implementation jar files.
-d <i>domainPath</i>	Path to the monitored domain. It must be a simple or master domain.
-c <i>configFileName</i>	Configuration file in Java properties file format. This argument is optional.
name1=value1 name2=value2 ...	The name/value pairs. If the configuration file is not specified, the name/value pairs are used. If the configuration file is specified, the name/value pairs are added to the configuration and overwrite any value with the same name already present in the file. This argument is optional.

Configuration Settings

The configuration settings for the JMS Publisher can be categorized into general and vendor-specific settings.

General Settings

Table 15–3 lists all vendor-neutral configuration settings for the JMS Publisher:

Table 15–3 Vendor-Neutral Configuration Settings for the JMS Publisher

Setting Name	Description	Required	Value
topic	JMS topic lookup name.	Yes	A string with only alphanumeric characters, underscores, and periods. Maximum length is 100 bytes.
topicConnectionFactory	JMS topic connection factory lookup name.	Yes	A string with only alphanumeric characters, underscores, and periods. Maximum length is 100 bytes.
messageTemplate	Template file name for the JMS message. This file must be placed under the config directory of the domain.	Yes	Standard template file PMCMessageTemplate.xml is available under \$RPAS_HOME/domain/config. This file must be copied to the config directory of the domain.
logLevel	Logging level. A log level is a cut-off level which means logs with a lower level are filtered out.	No	Specify one of the following values, in low to high order. "VERBOSE": All logs. "DEBUG": Debug logs. "INFO": Informational logs. "WARN": Warning logs. "ERROR": Error logs. "SUPPRESS": No logs. If not specified, the default is "INFO".
logFile	The path to the log file. Can be a relative or absolute path.	No	If not specified, output to the console.
restartAfterException	Flag to restart after encountering any JMS related exceptions.	No	If true, the publisher will try to restart after catching any JMS related exceptions. Interval between retries is 180 seconds. To stop the publisher from trying to restart, the process must be ended manually. Default is true.
message.deliveryMode	Delivery mode.	No	"NON_PERSISTENT" or "PERSISTENT". Default is "PERSISTENT"
message.priority	A priority number for the JMS message.	No	0 to 9. Default is 4.
message.timeToLive	Time to live in milliseconds.	No	0 or any positive integer. Default is 0 (unlimited).

Vendor-Specific Settings

Sun Open Message Queue is used as an example for these settings. For further details, consult the vendor documentation for your JMS implementation.

File-Based JNDI Object Store

File-based JNDI object store is used primarily for development and testing. It is very easy to set up, but has weak built-in security.

Table 15–4 File-Based JNDI Object Store

Property Name	Description	Required	Value
java.naming.factory.initial	Initial context for JNDI lookup.	Yes	For Open Message Queue, it must be: com.sun.jndi.fscontext.ReffFSContextFactory
java.naming.provider.url	Directory path to the object store.	Yes	Example: file:///C:/myapp/mqobjs

LDAP-Based JNDI Object Store

An LDAP server is the recommended object store for production JMS messaging systems. LDAP servers are designed for use in distributed systems and provide security features that are required in production environments.

Table 15–5 LDAP-Based JNDI Object Store

Property Name	Description	Required	Value
java.naming.factory.initial	Initial context for JNDI lookup.	Yes	For Open Message Queue, it must be: com.sun.jndi.ldap.LdapCtxFactory
java.naming.provider.url	Server URL and directory path to the object store.	Yes	Example: ldap://myD.com:389/ou=mq1,o=App where administered objects are stored in the directory /App/mq1
java.naming.security.principal	Identity of the principal for authenticating callers.	No	The format of this attribute depends on the authentication scheme. For example: uid=homerSimpson,ou=People,o=mq If this attribute is not specified, the behavior is determined by the LDAP service provider.
java.naming.security.credentials	Credentials of the authentication principal.	No	The value of this attribute depends on the authentication scheme. For example, it might be a hashed password, clear-text password, key, or certificate. If this property is not specified, the behavior is determined by the LDAP service provider.

Table 15–5 (Cont.) LDAP-Based JNDI Object Store

Property Name	Description	Required	Value
java.naming.security.authentication	Security level for authentication.	No	Specify one of these options: "none": No security "simple": Simple security "strong": Strong security For example, if you specify "simple", you are prompted for any missing principal or credential values. This enables a more secure way of providing identifying information. If this property is not specified, the behavior is determined by the LDAP service provider.

Here is a sample configuration file for RPAS JMS Publisher:

```
java.naming.factory.initial=com.sun.jndi.fscontext.RefFSContextFactory
java.naming.provider.url=file:///C:/Temp
topic=RPASEventTopic
topicConnectionFactory=RPASEventTopicConnectionFactory
messageTemplate=PMCMessageTemplate.xml
```

Configuring the RPAS JMS Subscriber

A sample JMS Subscriber is implemented to use for testing or troubleshooting. This subscriber is packaged and deployed along with the publisher. It simply writes the messages it receives to logging output, which can be console or a log file.

Command Line

Here is the command line for the sample JMS Subscriber:

```
java -cp {classpath} oracle.rpas.pmc.MCSubscriber [-c {configFileName}]
[name1=value1 [name2=value2 [...]]]
```

- The {classpath} should include rpaspmc.jar and jms.jar under \$RPAS_HOME/lib and any vendor-specific JMS implementation jar files.
- The {configFileName} is a configuration file in Java properties file format. It is optional.
- The name/value pairs are used if the configuration file is not specified. Otherwise, the name/value pairs are added to the configuration file and any value with the same name already present in the file is overwritten.

Configuration Settings

Vendor-specific settings are the same as the JMS publisher. Refer to ["Configuring the RPAS JMS Publisher"](#) for details.

[Table 15–6](#) lists all the general settings for the JMS Subscriber:

Table 15–6 General Settings for the JMS Subscriber

Property Name	Description	Required	Value
topic	JMS topic lookup name.	Yes	A string with only alphanumeric characters, underscores, and periods. Maximum length is 100 bytes.
topicConnectionFactory	JMS topic connection factory lookup name.	Yes	A string with only alphanumeric characters, underscores, and periods. Maximum length is 100 bytes.
logLevel	Logging level. A log level is a cut-off level which means logs with a lower level are filtered out.	No	Specify one of the following values, in low to high order. ■ "VERBOSE": All logs. ■ "DEBUG": Debug logs. ■ "INFO": Informational logs. ■ "WARN": Warning logs. ■ "ERROR": Error logs. ■ "SUPPRESS": No logs. If not specified, default is "INFO".
logFile	The path to the log file. Can be a relative or absolute path.	No	If not specified, output to the console.
clientID	JMS Client ID. Only required for a durable subscription. Not recommended for other use.	No	This must be a unique string for all JMS clients using the same connection factory.
durableSubscriptionName	A unique name for a durable subscription.	No	If a name is not provided, the subscription will be a transient subscription. This means messages are not queued up if the connection for the subscriber is lost or the subscriber is not running.

Here is a sample configuration file for RPAS JMS Subscriber:

```
java.naming.factory.initial=com.sun.jndi.fscontext.RefFSContextFactory
java.naming.provider.url=file:///C:/Temp
topic=RPASEventTopic
topicConnectionFactory=RPASEventTopicConnectionFactory
```

In-Context Launch

In-Context Launch (ICL) is an RPAS feature that provides retailers with an integrated UI planning solution that can be used to respond to events by adjusting the affected plans within the context of the changed circumstances and interacting plans. A planner's workspace can be developed in which relevant events are announced in role-specific dashboards. The planner can drill down into the event report to identify the context of the event. Appropriate planning applications can be launched to react to the event by either replanning or deriving further insight from the plans, in order to justify the event.

RPAS launches a planning application given a context, which is the template or workbook to launch and the wizard parameters that define the scope if a new workbook is to be built. Launching the planning application involves the following steps:

1. Start the RPAS client if it is not already running.
2. Log in the user, if the user is not already logged in.
3. Create the workbook, if a new workbook is being built.
4. Open the new or existing workbook, if that is what the context specified.

The following topics are covered in this chapter:

- Opening a pre-built workbook.
- Building a new workbook given a specification of the wizard parameters.
- Issuing the open or build requests from a web application running in a Web portal.

Launching RPAS

A system integrator is responsible for creating a web-based application that can launch RPAS. Creating an application involves the following steps:

1. Configure the navigational pane to launch RPAS.
2. Create static or dynamic web reports and dashboards, including OBIEE and other Oracle Fusion technologies, to present integrated data to the end-user with links or buttons that launch RPAS in-context.
3. Create alerts or alert reports, using BPM, OBIEE, or other Oracle Fusion technologies, that have links or other UI mechanisms that can launch RPAS in-context.

RPAS provides the framework that system integrators can use to launch any RPAS-based workbook, independent of whether the workbook uses a standard wizard or a custom wizard.

RPAS supports launching workbooks that are built using the dynamic-template API. This includes all templates built using RPAS Configuration Tools, but does not include templates that have been coded in C++ using an API lower than the Dynamic Template.

Note: Successfully launching RPAS administrative templates that are not built using RPAS Configuration Tools is not guaranteed.

Issuing a Launch Request from a Web Page

RPAS ICL is implemented as an extension of the RPAS Web Launch feature. As a prerequisite, a RPAS Web Launch Servlet and RPAS Client web installation package must be deployed to a web server and made available to system integrators. For more information, see the "RPAS Classic Client Web Deployment" chapter in the *RPAS Installation Guide*.

System integrators must design the report page to have a link, or other web GUI widget, that invokes client-side JavaScript code to initiate an RPAS ICL. The page embeds a Java applet from RPAS Web Launch (RWL Applet).

Note: Java 1.6_10 or above is required.

The applet performs the following operations:

1. Finds the required RPAS client installation on the user's workstation, if it was installed by the applet in a previous web session.

Note that a standalone RPAS client installation is not recognized because the installation is tracked by a browser cookie. However, the applet can be forced to use a pre-installed client by setting a parameter for the applet.
2. Downloads and installs the appropriate version of the RPAS client if the required RPAS client installation cannot be found or a new build of the client is available on the web server.
3. Starts the RPAS client.
4. Routes the ICL request to the RPAS server through the RPAS client and waits for a response. The response can be a success or failure code and a corresponding message.

The applet makes a call back to a JavaScript function to return the response. It is the responsibility of the system integrators to provide code to handle the response in the JavaScript function.

Embedding an RWL Applet on a Web Page

System integrators can embed the applet in a web page using the regular HTML <Applet> tag. The following attributes are required in the <Applet> tag:

Table 16–1 Attributes Required in the <Applet> Tag

Attribute	Description	Value
code	Applet class name	com.retek.mdap.client.launch.LaunchApplet.class

Table 16–1 (Cont.) Attributes Required in the <Applet> Tag

Attribute	Description	Value
codebase	URL location to download the applet archive	The URL must point to the RPAS Web Launch Web server instance, for example, "http://mspdev43.us.oracle.com:7777/RPAS4/"
archive	Applet archive file name	mdap.jar
mayscript	Enable JavaScript interaction	NA

The following parameters are used to control the behavior of the RWL Applet:

Table 16–2 Parameters Used to Control the Behavior of the RWL Applet

Parameter	Description	Required	Example
AutoStart	Launch the client automatically when the applet starts. Must be set to false for ICL.	Yes	Must be set to false if the launch is initiated by JavaScript code.
defaultInstallDir	Default installation location for the RPAS client. The client version number is automatically appended to the end of this string.	No	C:\Oracle RPAS Client
Debug	Debug flag. If true, more logging information is printed in the Java Console of the web browser.	No	true
launchPreinstalled Only	Launch the preinstalled RPAS client only. Do not install an RPAS client. The preinstalled path is determined by the server side setting in the configuration file clientPath.txt. For more information, refer to the <i>RPAS Installation Guide</i> .	No	false
supportMultipleVersions	Support multiple RPAS client versions on the same workstation. Must be set to true for ICL.	Yes	Must be set to true.
WebLaunchBase	The URL for the Web Launch Servlet. If not specified, the applet will generate a URL by appending "/web" to its codebase property.	No	http://mspdev43.us.oracle.com:7777/RPAS4/web


It is recommended that the applet be embedded with the height and width set to zero. When the height and width are set to non-zero values, a text area is shown that gets updated with status information from the RPAS client.

Figure 16–1 shows the applet UI if the RWL Applet is embedded in the web page with non-zero width and height. The message window in the applet is updated with status information from the RPAS client. Status information includes the success or failure of building or opening a workbook and any error messages returned by RPAS.

It is assumed that when using non-zero height and width, system integrators appropriately size the applet in order for the UI to be user-friendly. For example, a

width and height of 5 is unrealistic for any user interface returning messages to the end-user.

Figure 16–1 UI of the RWL Applet when Height and Width Values are Non-Zero



Using JavaScript with the RWL Applet

System integrators embed the RWL Applet in an HTML web page. The launch context can be passed to the applet by the JavaScript-enabled active content of the web page. Real-time status is passed back. The communication between the applet and the active content is achieved by a JavaScript invocation of an applet method.

To launch an instance of an RPAS client, the active content makes a JavaScript call:

```
document.appletName.launchRPASClient(id, xmlstring [, user, password]);
```

where:

- **id**: RPAS client process identifier, which can be any string.
- **xmlstring**: launch context in XML format encoded by Base64 scheme. Base64 encoding is a popular scheme used in e-mail contents, readily available in JavaScript public domain source code.

For purposes of quick manual testing or prototyping, the MIME Tools plug-in in Notepad++ 4.8.5 and above may be used to encode any text to Base64 and back.

- **user and password**: optional arguments. For details on ICL supported mechanisms for user authentication, see ["User Authentication"](#).

The applet notifies active content regarding the status of the RPAS client through a JavaScript callback function. The following JavaScript function may be defined if a custom callback is desired:

```
function RPASStatusListener(id, code, message)
{
    ...
}
```


where

- id: same RPAS client process identifier previously passed by the launchRPASClient() call.
- code: status code (0 - success, 1 - failure, -1 - applet exception)
- message: a descriptive status message.

If the callback function is not defined, the RWL Applet displays status messages on the status bar of the browser in the following format: "id - [OK | Error | Exception] - message".

The following figure shows the interactions between the Active Content, RWL Applet and the RPAS client processes. Note that one embedded RWL applet can be invoked multiple times using JavaScript to launch multiple RPAS client processes using different context XML data. [Figure 16-2](#) shows the case where two RPAS client processes are launched.

Figure 16-2 Two RPAS Clients Being Launched



Since the status listener function `RPASStatusListener(id, code, message)` can be invoked multiple times by multiple status listener threads in the applet, the JavaScript developer should use the function ID argument to identify where the status message came from and act accordingly.

The web page must stay active to enable the communication between the RPAS client processes and RWL Applet. The web page can become inactive, for example, if the user clicks the Back button of the web browser or closes the active tab or window.

If the web page becomes inactive, the communication channel between the applet and client process is terminated and cannot be recovered. If this happens, the RPAS client

process becomes orphaned and then acts like a regular RPAS client. It continues to function, but its status is not displayed on the web page.

The following sample code demonstrates the integration of RWL Applet and JavaScript-enabled active content. Note that the width and height of the applet are set to zero so that the applet is hidden on the web page. The JavaScript function `launchRPASClient()` simply invokes the corresponding applet method. The status callback function `RPASStatusListener()` is implemented to display the status information in the web browser status bar that is generally at the bottom of the browser window.

```
<HTML>  
<HEAD>  
<META http-equiv=Content-Type content="text/html; charset=UTF-8">  
<TITLE>Launch RPAS Client</TITLE>  
</HEAD>  
<SCRIPT>  
var xmlDoccomplete = "BuildComplete";  
var xmlcomplete =  
"PD94bwWgdmVyc2lrbj0ims4wIiBlbmNvZGluZz0iVVRLGTgiID8+DQo8UlbBU0xh  
dW5jaENvbHnRleHQ+ICAgDQo8U2VydmVypGOkCTXkZXNjcmlwdGlrbj05UZXRORGR9tY  
YWluPC9kZXNjcmlwdGlrbj04NCiAgICA8CjJvZmZsZT5mYXczZTwvcHJvZmZsZT4N\  
Cgk8ZGVidWc+ZmFsc2U8L2RlYnVnPgOKICAgIDxkYkRhZW1vblBvcnQ+MTIzNDg8\  
L2RiRGFlbnw9UG9ydD4NCgk8ZGJTZXJ2ZXJIb3NOtmfTZT5tc3BkZXYOMy5lcY5v\  
cmFjbGUuY29tPC9kYlNlcnZlckhvzc3ROYW1lPgOKCTXkYlNlcnZlclBvcnRTdGFy\  
dD4xMDI1PC9kYlNlcnZlclBvcnRTdGFyYD4NCgk8ZGJTZXJ2ZXJQb3J0RW5kPjY1\  
NTM1PC9kYlNlcnZlclBvcnRFbmQ+DQoJPGRlZmFlbHREYlVzXXIGLz4NCgk8ZGVm\  
YXVsdeRIUGFzc3dvcmQGzL4NCgk8ZG9tYWluUm9vdERpcmvjdG9yeSAvPgOKPC9T\  
ZXJ2ZXI+DTxDGbllbnQ+DSAGIDx2ZXJzaW9uPjEzLjEuMDdwvdmVyc2lrbj04NPC9D\  
bGllbnQ+DQo8RG9tYWluPgOKCTXkZXNjcmlwdGlrbj05UZXRORGR9tYWluPC9kZXNj\  
cmlwdGlrbj04NCgk8ZGlyZWNOb3J5TmfTZT4vdm9sLm5hcy9ycGFzX3NlL2hlYW5n\  
bWEVeZG9tYWlucy9UZXRORGR9tYWluPC9kaXJlY3Rvcnl0YW1lPgOKICAgPHVzXXI+\  
YWRtPC9lc2VyPgOKCTXwYXNZd29yZD5hZG08L3Bhc3N3b3JkPgOKPC9Eb21haW4+\  
DQo8V29ya2Jvb2s+DQoJPGRlZmFlbHREYlVzXXIGLz4NCgk8ZGVmYXVsdeRIUGFzc3dvcmQGzL4NCgk8ZG9tYWluUm9vdERpcmvjdG9yeSAvPgOKPC9TZXJ2ZXI+DTxDGbllbnQ+DSAGIDx2ZXJzaW9uPjEzLjEuMDdwvdmVyc2lrbj04NPC9DbGllbnQ+DQo8RG9tYWluPgOKCTXkZXNjcmlwdGlrbj05UZXRORGR9tYWluPC9kZXNjcmlwdGlrbj04NCgk8ZGlyZWNOb3J5TmfTZT4vdm9sLm5hcy9ycGFzX3NlL2hlYW5nbWEVeZG9tYWlucy9UZXRORGR9tYWluPC9kaXJlY3Rvcnl0YW1lPgOKICAgPHVzXXI+YWRtPC9lc2VyPgOKCTXwYXNZd29yZD5hZG08L3Bhc3N3b3JkPgOKPC9Eb21haW4+DQo8V29ya2Jvb2s+DQoJPGRlZmFlbHREYlVzXXIGLz4NCgk8ZGVmYXVsdeRIUGFzc3dvcmQGzL4NCgk8ZG9tYWluUm9vdERpcmvjdG9yeSAvPgOKPC9TZXJ2ZXI+DTxDGbllbnQ+DSAGIDx2ZXJzaW9uPjEzLjEuMDdwvdmVyc2lrbj04NPC9DbGllbnQ+DQo8RG9tYWluPgOKCTXkZXNjcmlwdGlrbj05UZXRORGR9tYWluPC9kZXNjcmlwdGlrbj04NCgk8ZGlyZWNOb3J5TmfTZT4vdm9sLm5hcy9ycGFzX3NlL2hlYW5nbWEVeZG9tYWlucy9UZXRORGR9tYWluPC9kaXJlY3Rvcnl0YW1lPgOKICAgPHVzXXI+YWRtPC9lc2VyPgOKCTXwYXNZd29yZD5hZG08L3Bhc3N3b3JkPgOKPC9Eb21haW4+DQo8V29ya2Jvb2s+DQoJPGRlZmFlbHREYlVzXXIGLz4NCgk8ZGVmYXVsdeRIUGFzc3dvcmQGzL4NCgk8ZG9tYWluUm9vdERpcmvjdG9yeSAvPgOKPC9TZXJ2ZXI+DTxDGbllbnQ+DSAGIDx2ZXJzaW9uPjEzLjEuMDdwvdmVyc2lrbj04NPC9DbGllbnQ+DQo8RG9tYWluPgOKCTXkZXNjcmlwdGlrbj05UZXRORGR9tYWluPC9kZXNjcmlwdGlrbj04NCgk8ZGlyZWNOb3J5TmfTZT4vdm9sLm5hcy9ycGFzX3NlL2hlYW5nbWEVeZG9tYWlucy9UZXRORGR9tYWluPC9kaXJlY3Rvcnl0YW1lPgOKICAgPHVzXXI+YWRtPC9lc2VyPgOKCTXwYXNZd29yZD5hZG08L3Bhc3N3b3JkPgOKPC9Eb21haW4+DQo8V29ya2Jvb2s+DQoJPGRlZmFlbHREYlVzXXIGLz4NCgk8ZGVmYXVsdeRIUGFzc3dvcmQGzL4NCgk8ZG9tYWluUm9vdERpcmvjdG9yeSAvPgOKPC9TZXJ2ZXI+DTxDGbllbnQ+DSAGIDx2ZXJzaW9uPjEzLjEuMDdwvdmVyc2lrbj04NPC9DbGllbnQ+DQo8RG9tYWluPgOKCTXkZXNjcmlwdGlrbj05UZXRORGR9tYWluPC9kZXNjcmlwdGlrbj04NCgk8ZGlyZWNOb3J5TmfTZT4vdm9sLm5hcy9ycGFzX3NlL2hlYW5nbWEVeZG9tYWlucy9UZXRORGR9tYWluPC9kaXJlY3Rvcnl0YW1lPgOKICAgPHVzXXI+YWRtPC9lc2VyPgOKCTXwYXNZd29yZD5hZG08L3Bhc3N3b3JkPgOKPC9Eb21haW4+DQo8V29ya2Jvb2s+DQoJPGRlZmFlbHREYlVzXXIGLz4NCgk8ZGVmYXVsdeRIUGFzc3dvcmQGzL4NCgk8ZG9tYWluUm9vdERpcmvjdG9yeSAvPgOKPC9TZXJ2ZXI+DTxDGbllbnQ+DSAGIDx2ZXJzaW9uPjEzLjEuMDdwvdmVyc2lrbj04NPC9DbGllbnQ+DQo8RG9tYWluPgOKCTXkZXNjcmlwdGlrbj05UZXRORGR9tYWluPC9kZXNjcmlwdGlrbj04NCgk8ZGlyZWNOb3J5TmfTZT4vdm9sLm5hcy9ycGFzX3NlL2hlYW5nbWEVeZG9tYWlucy9UZXRORGR9tYWluPC9kaXJlY3Rvcnl0YW1lPgOKICAgPHVzXXI+YWRtPC9lc2VyPgOKCTXwYXNZd29yZD5hZG08L3Bhc3N3b3JkPgOKPC9Eb21haW4+DQo8V29ya2Jvb2s+DQoJPGRlZmFlbHREYlVzXXIGLz4NCgk8ZGVmYXVsdeRIUGFzc3dvcmQGzL4NCgk8ZG9tYWluUm9vdERpcmvjdG9yeSAvPgOKPC9TZXJ2ZXI+DTxDGbllbnQ+DSAGIDx2ZXJzaW9uPjEzLjEuMDdwvdmVyc2lrbj04NPC9DbGllbnQ+DQo8RG9tYWluPgOKCTXkZXNjcmlwdGlrbj05UZXRORGR9tYWluPC9kZXNjcmlwdGlrbj04NCgk8ZGlyZWNOb3J5TmfTZT4vdm9sLm5hcy9ycGFzX3NlL2hlYW5nbWEVeZG9tYWlucy9UZXRORGR9tYWluPC9kaXJlY3Rvcnl0YW1lPgOKICAgPHVzXXI+YWRtPC9lc2VyPgOKCTXwYXNZd29yZD5hZG08L3Bhc3N3b3JkPgOKPC9Eb21haW4+DQo8V29ya2Jvb2s+DQoJPGRlZmFlbHREYlVzXXIGLz4NCgk8ZGVmYXVsdeRIUGFzc3dvcmQGzL4NCgk8ZG9tYWluUm9vdERpcmvjdG9yeSAvPgOKPC9TZXJ2ZXI+DTxDGbllbnQ+DSAGIDx2ZXJzaW9uPjEzLjEuMDdwvdmVyc2lrbj04NPC9DbGllbnQ+DQo8RG9tYWluPgOKCTXkZXNjcmlwdGlrbj05UZXRORGR9tYWluPC9kZXNjcmlwdGlrbj04NCgk8ZGlyZWNOb3J5TmfTZT4vdm9sLm5hcy9ycGFzX3NlL2hlYW5nbWEVeZG9tYWlucy9UZXRORGR9tYWluPC9kaXJlY3Rvcnl0YW1lPgOKICAgPHVzXXI+YWRtPC9lc2VyPgOKCTXwYXNZd29yZD5hZG08L3Bhc3N3b3JkPgOKPC9Eb21haW4+DQo8V29ya2Jvb2s+DQoJPGRlZmFlbHREYlVzXXIGLz4NCgk8ZGVmYXVsdeRIUGFzc3dvcmQGzL4NCgk8ZG9tYWluUm9vdERpcmvjdG9yeSAvPgOKPC9TZXJ2ZXI+DTxDGbllbnQ+DSAGIDx2ZXJzaW9uPjEzLjEuMDdwvdmVyc2lrbj04NPC9DbGllbnQ+DQo8RG9tYWluPgOKCTXkZXNjcmlwdGlrbj05UZXRORGR9tYWluPC9kZXNjcmlwdGlrbj04NCgk8ZGlyZWNOb3J5TmfTZT4vdm9sLm5hcy9ycGFzX3NlL2hlYW5nbWEVeZG9tYWlucy9UZXRORGR9tYWluPC9kaXJlY3Rvcnl0YW1lPgOKICAgPHVzXXI+YWRtPC9lc2VyPgOKCTXwYXNZd29yZD5hZG08L3Bhc3N3b3JkPgOKPC9Eb21haW4+DQo8V29ya2Jvb2s+DQoJPGRlZmFlbHREYlVzXXIGLz4NCgk8ZGVmYXVsdeRIUGFzc3dvcmQGzL4NCgk8ZG9tYWluUm9vdERpcmvjdG9yeSAvPgOKPC9TZXJ2ZXI+DTxDGbllbnQ+DSAGIDx2ZXJzaW9uPjEzLjEuMDdwvdmVyc2lrbj04NPC9DbGllbnQ+DQo8RG9tYWluPgOKCTXkZXNjcmlwdGlrbj05UZXRORGR9tYWluPC9kZXNjcmlwdGlrbj04NCgk8ZGlyZWNOb3J5TmfTZT4vdm9sLm5hcy9ycGFzX3NlL2hlYW5nbWEVeZG9tYWlucy9UZXRORGR9tYWluPC9kaXJlY3Rvcnl0YW1lPgOKICAgPHVzXXI+YWRtPC9lc2VyPgOKCTXwYXNZd29yZD5hZG08L3Bhc3N3b3JkPgOKPC9Eb21haW4+DQo8V29ya2Jvb2s+DQoJPGRlZmFlbHREYlVzXXIGLz4NCgk8ZGVmYXVsdeRIUGFzc3dvcmQGzL4NCgk8ZG9tYWluUm9vdERpcmvjdG9yeSAvPgOKPC9TZXJ2ZXI+DTxDGbllbnQ+DSAGIDx2ZXJzaW9uPjEzLjEuMDdwvdmVyc2lrbj04NPC9DbGllbnQ+DQo8RG9tYWluPgOKCTXkZXNjcmlwdGlrbj05UZXRORGR9tYWluPC9kZXNjcmlwdGlrbj04NCgk8ZGlyZWNOb3J5TmfTZT4vdm9sLm5hcy9ycGFzX3NlL2hlYW5nbWEVeZG9tYWlucy9UZXRORGR9tYWluPC9kaXJlY3Rvcnl0YW1lPgOKICAgPHVzXXI+YWRtPC9lc2VyPgOKCTXwYXNZd29yZD5hZG08L3Bhc3N3b3JkPgOKPC9Eb21haW4+DQo8V29ya2Jvb2s+DQoJPGRlZmFlbHREYlVzXXIGLz4NCgk8ZGVmYXVsdeRIUGFzc3dvcmQGzL4NCgk8ZG9tYWluUm9vdERpcmvjdG9yeSAvPgOKPC9TZXJ2ZXI+DTxDGbllbn
```



```

PC92YwX1zT4NCgkJCQk8L1N1cGVyVHJlZT4NCgkJCTwvV2l6YXJkUGFnZT4NCgkJ\
CTxXaXphcmRQYwdlIGlkPSJTS1UiIGZhbGxiYWNRPSJub25lIj4NCgkJCQk8U3Vw\
ZXJUCmVlIG5hbWU9InRyZWUuIj4NCgkJCQkJPkZhbHVlPg0KCQkJCQkJPGRpbSBu\
YW1lPSJDTFNTIj4NCgkJCQkJCQk8c2VsZWNOaW9uPjAyMTwvc2VsZWNOaW9uPg0K\
CQkJCQkJPc9kaW0+DQoJCQkJCTwvdmFsdWU+DQoJCQkJPc9TdXB1c1RyZWU+DQoJ\
CQk8L1dpemFyZFBhZ2U+DQoJCTwvV2l6YXJkUGFnZXM+DQoJCTxzYXZlPg0KCQkJ\
PGxhYmVsPlRlc3R3YjwvbGFiZWw+DQoJCQk8YWNjZXNzPkdYb3VwPC9hY2Nlc3M+\
DQoJCQk8Z3JvdXA+QURNSU48L2dyb3VwPg0KCQk8L3NhdmU+DQoJPC9idWlsZD4N\
CjwvV29ya2Jvb2s+DQo8L1JQVNMVXV2hDb250ZXh0Pg0K";
function launchRPASClient(id, xmlstring)
{
    document.LaunchApplet.launchRPASClient(id, xmlstring);
}
function RPASStatusListener(id, code, message)
{
    status = id + " - Status: code(" + code + ") message(" + message + ")";
}
</SCRIPT>
<BODY>
<CENTER>
<H3>Launch RPAS Client</H3>


Click on the button to launch RPAS client and build a new workbook!
<P>
<INPUT type="button" value="Build A New Measure Analysis Workbook"
    onClick = "launchRPASClient(idcomplete, xmlcomplete)">
<APPLET name="LaunchApplet"

code="com.retek.mdap.client.launch.LaunchApplet.class"
codebase="http://mspdev43.us.oracle.com:7777/RPAS4/"
archive="mdap.jar"
align="center" width="0" height="0" mayscript VIEWASTEXT>

<PARAM name="AutoStart" value="false">
<PARAM name="supportMultipleVersions" value="true">
<PARAM name="launchPreinstalledOnly" value="false">
<PARAM name="Debug" value="true">
<PARAM name="defaultInstallDir" value="C:\Oracle RPAS Client">
Your web browser is not currently Java-enabled.
</APPLET>
</CENTER>
</BODY>
</HTML>

```

Figure 16-3 shows this sample code rendered by a web browser. Note that the status bar shows the current status of the RPAS client.

Figure 16–3 Example HTML Code Rendered by a Web Browser

User Authentication

This section describes Oracle Single Sign-On (OSSO) and non-OSSO authentication of users.


OSSO-Based Authentication

If the HTML active content is protected by OSSO, it can be assumed that the user has already been authenticated when the applet is started. However, it cannot be assumed that the user has the privileges to launch the RPAS client. The privileges can only be confirmed by the RWL Servlet. The applet, therefore, makes a query to the servlet to confirm the privileges of the user before launching the RPAS client.

When OSSO is being used, the user name and password are not needed and should not be used for the call to `document.appletName.launchRPASClient()`. System integrators should rely on OSSO credentials and mechanisms to authenticate the user.

Figure 16–4 shows the flow of OSSO support for ICL.

Figure 16–4 OSSO Support Flowchart for ICL



Non-OSSO Authentication

When OSSO is not used, the authentication for the RPAS client login is done by RPAS itself. There are several ways to pass user credentials, user name and password, to the RPAS client through the RWL Applet:

- Directly embed the user name and password in the XML data. The domain section of launch context can contain user and password information.

Note: This may pose a security issue, depending on the life cycle of the launch contextual data.


- Do not pass the user name and password to the applet. The RPAS client prompts the user to log in every time the client starts. This may be inconvenient, but may be acceptable if only one RPAS session is needed.
- Use JavaScript code to pass in the user name and password in the applet method call `launchRPASClient(id, xmlstring, user, password)`. The user name and password can be obtained from the user input fields of the web page. The user only needs to enter the user name and password once for one browser session. Multiple RPAS client sessions can then be launched without entering the user name and password again.

Deploying In-Context Launch

Figure 16–5 shows how an RPAS In-Context Launch enabled application is deployed. Note that the RPAS Web Launch Servlet has to be deployed on a web server. This web server does not have to be the same as the web server that is used for any reporting, dashboard or BAM applications. In this release, RPAS ICL is able to interact with the

dynamic content of the web page using mechanisms described in the previous sections. As long as system integrators can ensure that the RPAS Web Launch Servlet is available, ICL can be achieved from any web-based application.

Figure 16–5 ICL Deployment Diagram



Use Cases

This section describes some use cases.

Launching without Context

A system integrator may want to launch RPAS without any context. This means starting up an RPAS client and logging on automatically, but without opening or building a workbook. This type of launch is required to support a generic link in the navigation pane where an end-user can simply launch RPAS to do planning using RPAS Today. The navigation pane only serves as a launchpad for RPAS, where RPAS can be started and logged into by clicking a static link or button.

To code this type of launch, the user must either remove the `<Workbook>` tag from the context or ensure that the `<Workbook>` tag does not include an `<open>` or `<build>` request.

Launching With Context

An RPAS launch that includes more than a simple login into RPAS is considered an in-context launch. This includes launching an RPAS client to open an existing workbook or build a workbook given the template and wizard parameters.

Opening an Existing Workbook

A system integrator may want to design an application that opens a pre-built workbook. This workbook may have been built by a system integrator, the user, workbook build batch, or previous in-context launch request.

It is assumed that system integrators are aware of either the internal, unique workbook name that RPAS uses or the workbook label, and optionally, the template name and creation date and time of the workbook. These parameters can be specified in the XML context specification to have RPAS open the workbook. For information, see ["Open Workbook"](#).

To enable opening the latest workbook from a nightly auto-workbook-build batch, system integrators can provide the label and template name. ICL automatically defaults to the latest workbook of the same label.

Note: The last workbook that was created is opened (not the last workbook that was modified).

Building a New Workbook

A system integrator may want to ensure that a new workbook is built for the end-user depending on what the end-user is or has been doing in the web application. The following examples illustrate this concept.

- If the user clicks a link for "Merchandise Financial Planning" in the navigation pane, RPAS is launched to build a workbook for the "Merchandise Financial Planning" template. The end-user is required to make all the wizard selections.
- If the user is navigating a report and drills down to identify some outliers at the SKU/Store level, RPAS is launched to automatically build a workbook for those SKU/Stores without requiring the end-user to reselect them in the wizard.

This can be achieved using the build context detailed in the "Build Workbook" section.

RPAS Launch Context Format

The launch context is described in XML format. It consists of four components: client, server, domain, and workbook.

- The client component identifies the RPAS client to be launched. It is required to support multiple client versions on the same workstation.
- The server component identifies the server to connect to.
- The domain component identifies the domain to log in to.
- The workbook component contains information for opening or building a workbook.

Client Settings

Table 16–3 *Client Settings*

Setting	Description	Type	Required
version	RPAS client version number. RPAS Web Launch supports multiple client versions. This version number is used to determine which client binary will be launched. If not specified, the default client for RPAS Web Launch is used.	String	No
allowConcurrentSessions	Allow the same user to open multiple concurrent RPAS client sessions connecting to the same domain without prompting for confirmation.	Boolean	No
sslWalletLocation	The absolute path to the Oracle wallet for the Classic Client.	String	No

Server Settings

Table 16–4 Server Settings

Setting	Description	Type	Required
description	Descriptive name for this server configuration.	String	No
profile	Enable profiling for the RPAS client.	Boolean	No
debug	Enable debug logging for the RPAS client.	Boolean	No
dbDaemonPort	The port number of DomainDaemon.	Integer	Yes
dbServerHostName	The host name where DomainDaemon is running.	Integer	Yes
dbServerPortStart	The starting port number of the RpasDbServer listening port range.	Integer	No
dbServerPortEnd	The ending port number of the RpasDbServer listening port range.	Integer	No
defaultDbUser	The default user ID for RpasDbServer.	String	No
defaultDbPassword	The encrypted password for the default user ID.	String	No
domainRootDirectory	The root directory of RPAS domains. If specified, it is prepended to the directoryName in Domain Settings.	String	No

Domain Settings

Table 16–5 Domain Settings

Setting	Description	Type	Required
description	Descriptive name for this domain configuration. It is displayed as the title of the RPAS client main window.	String	Yes
directoryName	The domain path.	String	Yes
user	The user ID used to log in. If not specified, a login dialog prompts for user input.	String	No
password	The encrypted password for the user ID. If not specified, a login dialog prompts for user input.	String	No

Workbook Settings

The Workbook section has two options to open or build a workbook. If the options or the whole Workbook section is not specified, the RPAS client logs in to the domain without any workbook operation. For information on opening a workbook, see ["Open Workbook"](#). For information on building workbook, see ["Build Workbook"](#).

XML Schema

Here is the full schema for RPAS Launch Context.

```
<?xml version="1.0" encoding="UTF-8" ?>

<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema">

  <xs:complexType name="BooleanDataType">
```

```

<xs:sequence>
  <xs:element name="value" type="xs:boolean" />
  <xs:element name="default" type="xs:boolean" minOccurs="0" />
</xs:sequence>
<xs:attribute name="name" type="xs:string" use="required" />
<xs:attribute name="fallback" type="xs:string" use="optional" />
</xs:complexType>

<xs:complexType name="StringDataType">
  <xs:sequence>
    <xs:element name="value" type="xs:string" />
    <xs:element name="default" type="xs:string" minOccurs="0" />
  </xs:sequence>
  <xs:attribute name="name" type="xs:string" use="required" />
  <xs:attribute name="fallback" type="xs:string" use="optional" />
</xs:complexType>

<xs:complexType name="SelectionType">
  <xs:sequence>
    <xs:element name="selection" type="xs:string" maxOccurs="unbounded" />
  </xs:sequence>
</xs:complexType>

<xs:complexType name="StringSetDataType">
  <xs:sequence>
    <xs:element name="value" type="SelectionType" />
    <xs:element name="default" type="SelectionType" minOccurs="0" />
  </xs:sequence>
  <xs:attribute name="name" type="xs:string" use="required" />
  <xs:attribute name="fallback" type="xs:string" use="optional" />
</xs:complexType>

<xs:complexType name="StringVectorDataType">
  <xs:sequence>
    <xs:element name="value" type="SelectionType" />
    <xs:element name="default" type="SelectionType" minOccurs="0" />
  </xs:sequence>
  <xs:attribute name="name" type="xs:string" use="required" />
  <xs:attribute name="fallback" type="xs:string" use="optional" />
</xs:complexType>

<xs:complexType name="DimDataType">
  <xs:sequence>
    <xs:element name="dim" maxOccurs="unbounded">
      <xs:complexType>
        <xs:complexContent>
          <xs:extension base="SelectionType">
            <xs:attribute name="name" type="xs:string" use="required" />
          </xs:extension>
        </xs:complexContent>
      </xs:complexType>
    </xs:element>
  </xs:sequence>
</xs:complexType>

<xs:complexType name="PQDDDataType">
  <xs:sequence>
    <xs:element name="pqd">
      <xs:complexType>
        <xs:simpleContent>

```

```

        <xs:extension base="xs:string">
            <xs:attribute name="type" type="xs:string" use="required" />
        </xs:extension>
    </xs:simpleContent>
</xs:complexType>
</xs:element>
</xs:sequence>
</xs:complexType>

<xs:element name="Checkbox" type="BooleanDataType" />
<xs:element name="RadioButton" type="BooleanDataType" />
<xs:element name="DatePicker" type="StringDataType" />
<xs:element name="DropDownList" type="StringDataType" />
<xs:element name="Edit" type="StringDataType" />
<xs:element name="Listbox" type="StringSetDataType" />
<xs:element name="SingleHierSelect" type="StringVectorDataType" />
<xs:element name="Tree" type="StringSetDataType" />

<xs:element name="SuperTree">
    <xs:complexType>
        <xs:choice>
            <xs:sequence>
                <xs:element name="value" type="DimDataType"/>
                <xs:element name="default" type="DimDataType" minOccurs="0" />
            </xs:sequence>
            <xs:sequence>
                <xs:element name="value" type="PQDDDataType"/>
                <xs:element name="default" type="PQDDDataType" minOccurs="0" />
            </xs:sequence>
        </xs:choice>
        <xs:attribute name="name" type="xs:string" use="required" />
        <xs:attribute name="fallback" type="xs:string" use="optional" />
    </xs:complexType>
</xs:element>

<xs:element name="open">
    <xs:complexType>
        <xs:choice>
            <xs:element name="name" type="xs:string" />
            <xs:sequence>
                <xs:element name="label" type="xs:string"/>
                <xs:element name="templateName" type="xs:string" minOccurs="0"/>
                <xs:element name="createdTime" type="xs:string" minOccurs="0"/>
            </xs:sequence>
        </xs:choice>
    </xs:complexType>
</xs:element>

<xs:element name="build">
    <xs:complexType>
        <xs:sequence>
            <xs:element name="templateName" type="xs:string" />
            <xs:element name="allowManualOverride" type="xs:boolean" minOccurs="0"/>
            <xs:element name="customErrorMessage" type="xs:string" minOccurs="0"/>
            <xs:element name="customWaitMessage" type="xs:string" minOccurs="0"/>
            <xs:element ref="WizardPages" />
            <xs:element ref="save" minOccurs="0"/>
        </xs:sequence>
    </xs:complexType>
</xs:element>

```



```

<xs:element name="save">
  <xs:complexType>
    <xs:sequence>
      <xs:element name="label" type="xs:string" />
      <xs:element name="access" type="xs:string" minOccurs="0"/>
      <xs:element name="group" type="xs:string" minOccurs="0"/>
    </xs:sequence>
  </xs:complexType>
</xs:element>

<xs:element name="WizardPage">
  <xs:complexType>
    <xs:sequence>
      <xs:choice minOccurs="0" maxOccurs="unbounded">
        <xs:element ref="Checkbox" />
        <xs:element ref="DatePicker" />
        <xs:element ref="DropDownList" />
        <xs:element ref="Edit" />
        <xs:element ref="Listbox" />
        <xs:element ref="RadioButton" />
        <xs:element ref="SingleHierSelect" />
        <xs:element ref="SuperTree" />
        <xs:element ref="Tree" />
      </xs:choice>
    </xs:sequence>
    <xs:attribute name="id" type="xs:string" use="required" />
    <xs:attribute name="option" type="xs:string" use="optional" />
    <xs:attribute name="fallback" type="xs:string" use="optional" />
  </xs:complexType>
</xs:element>

<xs:element name="WizardPages">
  <xs:complexType>
    <xs:sequence>
      <xs:element ref="WizardPage" maxOccurs="unbounded" />
    </xs:sequence>
  </xs:complexType>
</xs:element>

<xs:element name="Client">
  <xs:complexType>
    <xs:sequence>
      <xs:element name="version" type="xs:string" />
      <xs:element name="allowConcurrentSessions" type="xs:boolean"
minOccurs="0"/>
    </xs:sequence>
  </xs:complexType>
</xs:element>

<xs:element name="Server">
  <xs:complexType>
    <xs:sequence>
      <xs:element name="description" type="xs:string" />
      <xs:element name="profile" type="xs:boolean" />
      <xs:element name="debug" type="xs:boolean" />
      <xs:element name="dbDaemonPort" type="xs:integer" />
      <xs:element name="dbServerHostName" type="xs:string" />
      <xs:element name="dbServerPortStart" type="xs:integer" />
      <xs:element name="dbServerPortEnd" type="xs:integer" />
    </xs:sequence>
  </xs:complexType>
</xs:element>

```

```

        <xs:element name="defaultDbUser" type="xs:string" />
        <xs:element name="defaultDbPassword" type="xs:string" />
        <xs:element name="domainRootDirectory" type="xs:string" />
    </xs:sequence>
</xs:complexType>
</xs:element>

<xs:element name="Domain">
    <xs:complexType>
        <xs:sequence>
            <xs:element name="description" type="xs:string" />
            <xs:element name="directoryName" type="xs:string" />
            <xs:element name="user" type="xs:string" minOccurs="0"/>
            <xs:element name="password" type="xs:string" minOccurs="0"/>
        </xs:sequence>
    </xs:complexType>
</xs:element>

<xs:element name="Workbook">
    <xs:complexType>
        <xs:choice>
            <xs:element ref="open" />
            <xs:element ref="build" />
        </xs:choice>
    </xs:complexType>
</xs:element>

<xs:element name="RPASLaunchContext">
    <xs:complexType>
        <xs:sequence>
            <xs:element ref="Client" />
            <xs:element ref="Server" />
            <xs:element ref="Domain" />
            <xs:element ref="Workbook" minOccurs="0"/>
        </xs:sequence>
    </xs:complexType>
</xs:element>
</xs:schema>

```

Open Workbook

In any retail planning scenario, a particular workbook or, more likely, an auto-build workbook, may be used to plan for a specific grouping of products and locations. When an event or alert related to these products and locations occurs, the user may need to always open a particular workbook or auto-build workbook. System integrators configuring an application can configure and code an in-context opening of the specific workbook.

System integrators can programmatically specify the parameters that RPAS needs to open a pre-built workbook. Programmatically does not imply the use of an API, but instead implies program application logic that constructs the XML specification that ICL can use to identify and open the desired workbook.

System integrators can identify the target workbook by either of the following ways:

- By the RPAS name of the workbook.
- By the auto-build parameters, that is, user, template name, label, and creation date. Template name and creation date are optional parameters. If template name or creation date are not specified and multiple workbooks with the same label are

found, RPAS chooses the workbook that was created last and then further breaks the tie based on alphabetical ordering of the template name.

Specifications

Two options can be specified when a workbook is opened.

Option 1

Table 16–6 Option 1: Workbook Specification

Setting	Description	Type	Required
name	The RPAS name of the workbook, such as "t1" and "t18" which is the root directory name of the workbook)	String	Yes

Option 2

Table 16–7 Option 2: Workbook Specification

Setting	Description	Type	Required
label	The workbook label	String	Yes
templateName	The template name of the workbook	String	No
createdTime	The created time of the workbook in the format YYYY-MM-DD hh:mm:ss	String	No

Examples

Here is sample launch context for opening a workbook by RPAS name:

```
<?xml version="1.0" encoding="UTF-8" ?>
<RPASLaunchContext>
  <Server>
    <description>TestDomain</description>
    <profile>>false</profile>
    <debug>>false</debug>
    <dbDaemonPort>12348</dbDaemonPort>
    <dbServerHostName>mspdev43.us.oracle.com</dbServerHostName>
    <dbServerPortStart>1025</dbServerPortStart>
    <dbServerPortEnd>65535</dbServerPortEnd>
    <defaultDbUser />
    <defaultDbPassword />
    <domainRootDirectory /></Server>

  <Client>      <version>13.2</version></Client>

  <Domain>
    <description>TestDomain</description>
    <directoryName>/vol.nas/rpas_se/huangma/domains/TestDomain</directoryName>
    <user>adm</user>
    <password>...</password>
  </Domain>

  <Workbook>
    <open>
      <name>t1</name>
    </open>
  </Workbook>
</RPASLaunchContext>
```

Here is sample launch context for opening a workbook by label and template name:

```
<RPASLaunchContext>
... .. (Same as above, omitted here.)
<Workbook>
  <open>
    <label>wb2</label>
    <templateName>MEASUREANALYSIS</templateName>
  </open>
</Workbook>
</RPASLaunchContext>
```

Build Workbook

RPAS users generally build workbooks through a manual process using RPAS Wizard Pages from workbook templates. This process is automated in ICL by specifying the contents and selections of the wizard pages in the launch context.

The automated process enables system integrators to programmatically specify the parameters that RPAS needs to build a new workbook. The integrator can program the RPAS web launch to build virtually any workbook from any workbook templates, simple or complex, by constructing an XML string according to the launch context schema and the configuration of the wizard pages.

Users can think of the XML specifications as macro definitions of what they do when they build the same workbook manually. The specifications for building workbook start from a workbook template, which contains a list of wizard pages. Each wizard page has a set of pre-defined widgets:

- RadioButton
- Checkbox
- Edit
- DropDownList
- SingleHierSelect
- DatePicker
- Listbox
- Tree
- SuperTree

Each widget has its own specifications, depending on its type.

The fallback option can be specified on the widget and wizard page level. The widget level fallback option is considered first when the specifications for that widget are incomplete or incorrect. If a specification cannot be resolved on the widget level, the wizard page fallback option is used. If the fallback option on the wizard page is not specified or the option has errors, an error message is displayed to the user.

If any unrecoverable errors occur on a wizard page, the workbook build process falls back to manual mode if the allowManualOverride option is set to true. The wizard starts from the first page. All selections prior to the problematic page are preserved and the user can click **Next** to skip them and go to the page with errors, which is blank along with all subsequent pages.

Specifications

The specifications to build a workbook consist of basic settings and a list of wizard pages. Each wizard page contains a list of widget specifications.

Table 16–8 Basic Settings

Setting	Description	Type	Required
templateName	The template name for the workbook to be built.	String	Yes
allowManualOverride	Allow manual override in case of fatal errors in the XML data. The default is true. The wizard starts from the first page.	Boolean	No
customErrorMessage	Custom error message to display in case of fatal errors in the launch context data.	String	No
customWaitMessage	Custom wait message to display while the workbook is being built.	String	No

Table 16–9 Wizard Page Specifications

Setting	Description	Type	Required
option (attribute)	Wizard page handling option. Can be one of the following values: manual: forces manual selection for all widgets. This option forces manual input for this and all subsequent pages. However, the wizard still starts from the first page. skippable: this page can be skipped. Normally used for a conditional wizard page. This page only exists if some particular value is selected for some widget on the previous page.	String	No
fallback (attribute)	Wizard page level fall back option. Can be one of the following values: default: use values from the default settings for all widgets. cache: use selections from the cache selection file. Cache files are from last successful workbook build of the same template.	String	No

Table 16–10 Wizard Widget Specifications

RadioButton and Checkbox			
Setting	Description	Type	Required
name (attribute)	Widget name.	String	Yes
fallback (attribute)	Widget level fallback option. Can be: default: use value from the default setting.	String	No
value	Boolean type: selected value for the widget. Value must be either false or true.	XML section	Yes
default	Default value for the widget. Must be a Boolean type.	XML section	No

Table 16–11 Wizard Widget Specifications (2)

Edit, DropDownList, SingleHierSelect, and DatePicker

Setting	Description	Type	Required
name (attribute)	Widget name.	String	Yes
fallback (attribute)	Widget level fallback option. Can be: default: use value from the default setting.	String	No
value	String type: selected value for the widget.	XML section	Yes
default	Default value for the widget. Must be a String type.	XML section	No

Table 16–12 Wizard Widget Specifications (3)

Listbox and Tree

Setting	Description	Type	Required
name (attribute)	Widget name.	String	Yes
fallback (attribute)	Widget level fallback option. Can be one of the following values: ■ default: use value from the default setting. ■ all: select all available items or positions. ■ intersection: use intersection of the specified items/positions and all available items/positions if any of the specified items/positions do not exist.	String	No
value	String Set type: selected value for the widget. Value is a set of selections.	XML section	Yes
default	Default value for the widget. Must be a String type.	XML section	No

Table 16–13 Wizard Widget Specifications (4)

SuperTree

Setting	Description	Type	Required
name (attribute)	Widget name.	String	Yes
fallback (attribute)	Widget level fallback option. Can be one of the following values: ■ default: use value from the default setting. ■ all: select all available items or positions. ■ use intersection of specified positions and all available positions if any of the specified positions do not exist.	String	No

Table 16–13 (Cont.) Wizard Widget Specifications (4)

SuperTree			
Setting	Description	Type	Required
value	String to String Set Map Type/PQD File Type: selected value for the widget. Can be one of the following values: - Value is one or more dimensions, each having a set of selections. - Value is a Position Query Definition (PQD) file name in the format: <pqd type="{pqd_type}"> pqd_file_name </pqd> where {pqd_type} must be either PointAndClickPQD or RollingCalendarPQD.	XML section	Yes
default	Default value for the widget. Must be a String to String Set Map type or PQD file type.	XML section	No

Workbook Save Specifications

These save specifications are optional. If specified, the workbook is immediately saved after it is built.

Table 16–14 Workbook Save Specifications

Setting	Description	Type	Required
label	Workbook label.	String	Yes
access	Workbook access. Must be user, group, or world. If not specified, the default is user.	String	No
group	Workbook group. Required when workbook access is group. There is no default.	String	Conditional
customWaitMessage	Custom wait message to display while the workbook is being built.	String	No

Example

The following example shows a sample launch context for building a workbook. It specifies a new Measure Analysis workbook for a simple domain (TestDomain). The domain is created by using the RPAS utility createRpasDomain with the command line option -x.

```
<?xml version="1.0" encoding="UTF-8" ?>
<RPASLaunchContext>
  <Client>
    <version>13.2</version>
  </Client>
  <Server>
    <description>mspdev43-65535</description>
    <profile>>false</profile>
    <debug>>false</debug>
    <dbDaemonPort>12348</dbDaemonPort>
    <dbServerHostName>mspdev43.us.oracle.com</dbServerHostName>
    <dbServerPortStart>1025</dbServerPortStart>
    <dbServerPortEnd>65535</dbServerPortEnd>
    <defaultDbUser />
    <defaultDbPassword />
  </Server>
</RPASLaunchContext>
```

```

    <domainRootDirectory />
</Server>
<Domain>
  <description>TestDomain</description>
  <directoryName>/vol.nas/rpas_se/domains/TestDomain</directoryName>
</Domain>
<Workbook>
  <build>
    <templateName>MEASUREANALYSIS</templateName>
    <WizardPages>
      <WizardPage id="ExtraMeasuresPage" fallback="none">
        <Listbox name="lb1">
          <value>
            <selection>R_EX_DEMOA</selection>
            <selection>R_EX_DEMOB</selection>
            <selection>R_EX_DEMOC</selection>
          </value>
        </Listbox>
      </WizardPage>
      <WizardPage id="DAY" fallback="none">
        <SuperTree name="tree1" fallback="all">
          <value>
            <dim name="WEEK">
              <selection>W01_1997</selection>
              <selection>W01_1998</selection>
              <selection>W01_1999</selection>
              <selection>W01_2000</selection>
            </dim>
            <dim name="DAY">
              <selection>19970108</selection>
            </dim>
          </value>
          <default>
            <dim name="DAY">
              <selection>19970108</selection>
            </dim>
          </default>
        </SuperTree>
      </WizardPage>
      <WizardPage id="STR" fallback="none">
        <SuperTree name="tree1">
          <value>
            <dim name="STR">
              <selection>0102</selection>
              <selection>0144</selection>
              <selection>0152</selection>
              <selection>0557</selection>
              <selection>0594</selection>
              <selection>0959</selection>
            </dim>
          </value>
        </SuperTree>
      </WizardPage>
      <WizardPage id="SKU" fallback="none">
        <SuperTree name="tree1">
          <value>
            <dim name="CLSS">
              <selection>021</selection>
            </dim>
          </value>
        </SuperTree>
      </WizardPage>
    </WizardPages>
  </build>
</Workbook>

```



```

        </SuperTree>
    </WizardPage>
</WizardPages>
<save>
    <label>Testwb</label>
    <access>Group</access>
    <group>ADMIN</group>
</save>
</build>
</Workbook>
</RPASLaunchContext>

```

The buildWorkbookContext Utility

The buildWorkbookContext utility is a server side command line utility. It runs under both global domains and local domains.

Before running this utility, you should first manually create a workbook using the workbook template you intend to use in the context. This step generates selection files. In order to capture as much information as possible in the selection files, you can choose to select all for the widgets on wizard pages. This is especially useful for list box widgets.

After the XML file is generated, you can use the data in the file directly if you want to build an identical workbook. Or, you can use it as a starting point and tailor it to meet some specific requirements.

buildWorkbookContext Usage

```
buildWorkbookContext -d [domain_path] -user [user_name] -wbt [wbt_name] -out
[output_file]
```


where

- [domain_path]: The path to the global or local domain.
- [user_name]: The username whose selection files are used. This is the user who has manually run the workbook template wizard.
- [wbt_name]: The name of the workbook template for which the XML context is built. This is the internal name for the template. It can be identified by looking at the selection files under DOMAIN_PATH/users/[user]/selections. The names of the selection files are in this format: [templateName]_[wizpageID].sel.
- [output_file]: Output file name for the XML context.

Hybrid Storage Architecture

Hybrid Storage Architecture (HSA) combines the strengths of both relational and multi-dimensional data storage technologies. In an HSA environment, stored data is divided into two groups, based on their usage and integration characteristics. The majority of data is still stored in the domain, and Oracle database is leveraged for the storage of shared data across two or more applications. Multiple domains can share one Oracle database schema. The measure data in the Oracle database can be configured to be accessible by one or multiple domains. [Figure 17-1](#) illustrates a general scenario where multiple RPAS domains, reporting tools, and non-RPAS systems access shared data stored in Oracle database.

Figure 17-1 Hybrid Storage Architecture




The chief objective in an HSA deployment is to facilitate read-write access of shared data by multiple applications in near real time. Typically, the shared data is committed by an RPAS workbook and is read by another RPAS application, an external application, or a reporting tool. The use of the Oracle database mitigates RPAS data access contention issues that can occur due to the concurrent read-write operations.

The HSA architecture does not target bulk integration. In fact, storage of large volumes of writable data in Oracle database can have severe performance implications. Read-only bulk data that is not recalculated or modified by RPAS processes can be stored in the database; however, writable data such as data that is committed by workbooks or calculated during the batch cycle may be stored in the Oracle database only if the volume is relatively small.

An HSA deployment can benefit near real-time integration scenarios because the database provides transaction protection for the data stored in the database relational tables.

Figure 17–2 Hybrid Storage Architecture Integration




Current integration between RPAS applications is typically done through text data file export and import either through RPAS command line utilities or Oracle Data Integrator (ODI). It must be done through a batch cycle. In HSA, the measures in Oracle database can be shared by multiple RPAS applications. Any updates to the shared measures by one application are visible to other applications immediately after the transaction is committed in the Oracle database.

By moving data into the common Oracle database platform, the door is opened for standard database integration through Extraction, Transformation, and Loading (ETL) methodology with non-RPAS applications. RPAS HSA provides APIs that can be used in this kind of external integration.

The data stored in the Oracle database can be accessed by other Oracle database features or applications such as auditing, backup, and reporting. This can fit well with the customer's corporate enterprise data management strategy.

Figure 17–3 shows a physical deployment of the Hybrid Storage Architecture in a Fusion Client environment. The Oracle database server is introduced to host the RPAS Data Mart, adding a new data storage tier. The RPAS DB Server process acts as a client to the Oracle database and can access data in both the domain and the Data Mart. In a Fusion Client environment, the mid-tier WebLogic server remains unchanged. In a Classic Client environment, the WebLogic tier is not present as the Classic Client connects directly to the RPAS DB Server tier.

Figure 17-3 Hybrid Storage Architecture Deployment



All RPAS HSA enhancements are implemented on the RPAS server platform and are transparent to the RPAS client and end users. While the RPAS administrator has extra responsibilities in managing RDM configurations and operations, no extra training is required for the end users.

Data Model

As part of the HSA installation, the RDM is created, based on the integration and partition configuration. The RDM consists of eight schemas, one central data mart schema, and seven access schemas.

The data mart schema has the following schema objects:

- Tables
 - RPAS Metadata tables
 - Dimension-based tables
 - * Dimension tables
 - * Translation tables
 - * Index translation tables
 - Fact tables
 - Temporary and logging tables
- Packages

The dimension and fact-based tables' definitions depend on the configuration. The dimension and fact tables are partitioned based on the partitioning configuration. The remaining objects are fixed and are explained in this chapter.

Each of the access schemas has private synonyms for most of the data mart schema objects. In addition to this, based on their functionality, some access schemas may create staging tables dynamically.

This chapter describes the following components:

- Entity relationship diagram with sample dimension and fact configuration

- Data Mart's table description
- Data Mart's partitioning scheme
- Data Mart's sequence descriptions
- Data Mart's package descriptions
- Access schema descriptions
- Security considerations

Naming Convention

A standard naming convention is followed for the RPAS schema objects. The RPAS table naming conventions conforms to Fusion standards by adhering to the following guidelines.

- Schema object names start with the RP_ application short name.
- Names are grouped by functional or business area with a one- to two-letter abbreviation. If the object is used by multiple business areas, then the 'G_' abbreviation is used to identify generic object.
- Object names use [A-Z 0-9_] and no multiple consecutive underscores.

Here is the summary of all RPAS suffixes. Note that an object may have one or more suffixes applied at a given point in time.

Table 17-1 RPAS Suffixes

Suffix	Stands For...
MD	RPAS Metadata table
TL	Translated table
D	Dimension table
FT	Fact table
ITT	Index Translation table
MV	Materialized View
IDX	Index
PKG	Package
S	Sequence
ERR	Error table

Dimension-Fact Entity Relationship Diagram

This section provides an entity relationship diagram for the dimension and fact tables in the central data mart schema for a sample configuration.

Sample Integration Configuration

RPAS Configuration Tools allow customers to configure dimensions, measures (facts), business rules, and workbooks. Dimensions and facts impact the schema. A sample configuration described in [Figure 17-4](#) is used for illustration purposes in this chapter.


Three dimensions, product, location and time, are configured to have three levels each, as shown in [Figure 17-4](#).

The lowest level SKU rolls up to STYL (abbreviated words are used to respect the RPAS limit of four characters), which rolls up to CLSS for the Product dimension.

The lowest level STR rolls up to CITY, which rolls up to RGN for the Region dimension.


The lowest level DAY rolls up to WEEK, which rolls up to MNTH for the Time dimension.

Figure 17–4 Sample Configuration



Two fact groups are configured to have two facts each at the same intersection of sku, store, and week. The facts INV_BOP, INV_EOP, RECEIPT, and SALES are configured in two groups so that INV_EOP and INV_EOP belong to GroupA and RECEIPT and SALES belong to GroupB.

Figure 17–5 Fact Groups



Sample ERD for Dimension and Fact Tables

One table is created for each dimension level, and one table is created for each fact group in the central data mart schema, which is a relational representation of multi-dimensional data. The entity relationship model is termed a De-normalized Snowflake schema. The dimension tables are de-normalized so that each dimension table contains direct references not only to the immediate parent, but also to all higher level tables in the hierarchy. At the center of the snowflake are the fact tables. There can be multiple fact tables for each intersection, depending on the fact grouping

configuration. Figure 17-6 shows a sample ERD with the snowflake schema for two fact tables and related dimension tables, as illustrated above.

Figure 17-6 Sample ERD for Dimension and Fact Table




Table Descriptions

The data mart schema tables are broadly classified into the following logical categories:

- RPAS metadata tables
- Dimension-based tables
- Fact tables
- Temporary and logging tables

RPAS Metadata Tables

The RPAS metadata, that is, the metadata for the configured dimensions and facts and other RDM details, is stored in fixed RPAS metadata tables.

The metadata tables are named as 'RP_G_<Short_Desc>_MD'.

There are eleven RPAS metadata tables, as explained below, along with six metadata tables for the workbook commit process.

RP_G_DIMENSION_MD This table stores the shared dimensions with their configured order and purge age.

Table 17–2 RP_G_Dimension_MD

Name	Null	Type
DIMENSION ID	NOT NULL	NUMBER(12)
DIMENSION NAME	NOT NULL	VARCHAR2(256)
DIMENSION ORDER	NULL Allowed	NUMBER(12)
PURGE AGE	NULL Allowed	NUMBER(12)

RP_G_DOMAININFO_MD This table stores the ID, name, and path of the registered domains. The domain ID is generated from a sequence. The domain name is added from the integration configuration. The domain path is added only when the domain is registered.

Table 17–3 RP_G_Domaininfo_MD

Name	Null	Type
DOMAIN ID	NOT NULL	NUMBER(12)
DOMAIN NAME	NOT NULL	VARCHAR2(128)
DOMAIN PATH	NULL Allowed	VARCHAR2(512)

RP_G_FACT_INFO_MD This table stores the measure properties of shared facts as configured in the integration configuration.

Table 17–4 RP_G_Fact_Info_MD

Name	Null	Type
FACT ID	NOT NULL	NUMBER(12)
FACT NAME	NOT NULL	VARCHAR2(30)
FACT TYPE	NULL Allowed	NUMBER(12)
BASE INTX	NULL Allowed	VARCHAR2(24)
NA VALUE	NULL Allowed	VARCHAR2(512)
DEF AGG	NULL Allowed	VARCHAR2(24)
ALLOWED AGGS	NULL Allowed	VARCHAR2(128)
FACT GROUP	NULL Allowed	VARCHAR2(24)
PURGE AGE	NULL Allowed	NUMBER(12)
HYBRID AGG SPEC	NULL Allowed	VARCHAR2(512)

RP_G_INTEGRATION_MAP_MD This table contains entries for each registered domain having the shared measure. Since measure names need not be unique in all domains, this table provides a mapping to the measures in different domains that point to the same fact. This information must be provided in the integration configuration.

Table 17–5 RP_G_Integration_Map_MD

Name	Null	Type
FACT	NOT NULL	VARCHAR2(32)
DOMAIN	NULL Allowed	VARCHAR2(32)
MEASURE	NULL Allowed	VARCHAR2(32)

RP_G_LEVEL_MD This table stores the details of the shared levels of each shared dimension as configured.

Table 17-6 RP_G_Level_MD

Name	Null	Type
LEVEL ID	NOT NULL	NUMBER (12)
LEVEL NAME	NOT NULL	VARCHAR2(256)
DIMENSION ID	NOT NULL	NUMBER(12)
DIMENSION NAME	NOT NULL	VARCHAR2(256)
AGG FROM	NULL Allowed	VARCHAR2(256)
POSITION FORMAT	NULL Allowed	VARCHAR2(256)

RP_G_PARTITION_MD In a partitioned RDM, the partition level dimension table is partitioned by list. This table stores the mapping of those partition level positions with their respective partition IDs. This table is internally used by the RDM partition manager utility.

Table 17-7 RP_G_Partition_MD

Name	Null	Type
POSITION NAME	NOT NULL	VARCHAR2(256)
PARTITION_ID	NOT NULL	NUMBER(12)

RP_G_RDMINFO_MD This table is used internally to maintain some RDM properties for internal usage.

Table 17-8 RP_G_RDMINFO_MD

Name	Null	Type
PROPERTY NAME	NOT NULL	VARCHAR2(256)
PROPERTY VALUE	NOT NULL	VARCHAR2(256)

RP_G_RDM_HISTORY_MD This table stores RDM's history and versioning. It is used to record the history (creation and update) of RDM.

Table 17-9 RP_G_RDM_History_MD

Name	Null	Type
EVENT ID	NOT NULL	NUMBER(12)
SCHEMA VERSION	NOT NULL	NUMBER(12)
RPAS VERSION	NOT NULL	VARCHAR2(16)
LAST DEV TAG	NOT NULL	VARCHAR2(10)
UTILITY	NOT NULL	VARCHAR2(32)
ACTION	NOT NULL	VARCHAR2(32)
START TIME	NOT NULL	TIMESTAMP(6)
COMPLETE TIME	NOT NULL	TIMESTAMP(6)

RP_G_RDM_STATUS_MD In an integrated environment, interdependent utilities on different registered domains update the RDM status messages in this table. This is needed for better coordination among the interdependent utilities.

Table 17–10 RP_G_RDM_Status_MD

Name	Null	Type
SEQ ID	NOT NULL	NUMBER(12)
STATUS CODE	NOT NULL	NUMBER(12)
DOMAIN ID	NULL Allowed	NUMBER(12)
INFO	NULL Allowed	VARCHAR2(32)
CREATED TIME	NOT NULL	TIMESTAMP(6)

RP_G_RDM_ST_MSG_MD This is the parent table to the previous RD_G_RDM_STATUS_MD table and stores all the RDM status codes, along with their corresponding messages.

Table 17–11 RP_G_ST_MSG_MD

Name	Null	Type
STATUS CODE	NOT NULL	NUMBER(12)
STATUS MSG	NOT NULL	VARCHAR2(256)

RP_G_UPGRADE_MODULE_MD This table stores the information for each upgrade module applied to RDM.

Table 17–12 RP_G_Upgrade_Module_MD

Name	Null	Type
DEV TAG	NOT NULL	VARCHAR2(10)
EVENT ID	NOT NULL	NUMBER(12)
MODULE NAME	NOT NULL	VARCHAR2(32)
DESCRIPTION	NOT NULL	VARCHAR2(64)
START TIME	NOT NULL	TIMESTAMP(6)
COMPLETE TIME	NOT NULL	TIMESTAMP(6)

Commit Metadata Tables

RPAS maintains six metadata tables for the workbook commit process.

Table 17–13 Commit Metadata Tables

Table Name	Description
RP_G_COMMIT_DIMS_MD	This table stores the dimensions and levels and partitions that the workbook corresponds to for current commit (identified by commit ID). Used internally by the commit process.
RP_G_COMMIT_HISTORY_MD	This table maintains the history of state and start time for each commit.

Table 17–13 (Cont.) Commit Metadata Tables

Table Name	Description
RP_G_COMMIT_MASKS_MD	This table stores the conditional commit details: commit masks and the measures each mask applies to. Used internally by the commit process.
RP_G_COMMIT_MEASURES_MD	This table stores the measure names being committed along with their corresponding staging table. Used internally by the commit process.
RP_G_COMMIT_STATUS_MD	This table stores the states of the currently running commit process with various details like the workbook owner and template.
RP_G_COMMIT_STG_TBLS_MD	This table pairs the fact table/partition and the corresponding staging tables. The commit_measures table stores measure names in each staging table.

Dimension-Based Tables

For each shared level of the shared dimensions configured, RPAS creates one table for each of the following categories.

Dimension Table A dimension table is created for each dimension level named as 'RP_G_<level_Name>_D'.

The level tables are de-normalized so that the each level table has the reference to each higher level, so that the lowest level has reference to all higher level tables. Refer to [Figure 17–4](#).

The RPAS dimension tables have a position ID and business identifier for every level in the dimension. The position ID is an integer generated from a sequence, and the business identifier is an alpha-numeric public identifier, known as the external name in RPAS terminology.

Rollup dimension or shrunken dimension refers to taking a conformed subset of the dimension table to support aggregate data at higher fact grains. This design has the benefit of providing workbook build lookups with fewer rows.

Hierarchies and rollup paths may not be obvious to the design user simply by viewing the dimension tables if there are many alternate hierarchies. Metadata tables must be referenced for this that clearly define the hierarchies and levels for both the application and report users.

Each position of the dimension levels is assigned a surrogate key, which is a 32-bit integer and internally used as a unique identifier of the position. The records of fact tables are identified by combinations of the position IDs instead of positions. The position IDs are necessary for efficient data storage and query processing.

Each position has a status column, as LEVEL_STATUS, which identifies the status of the position (that is, formal, informal, place holder, or deleted).

The dimension tables have the following structure:

Table 17–14 Dimension Table Structure

Name	Null	Type
<LEVEL1_ID PRIMARY KEY>	NOT NULL	NUMBER(15)
<LEVEL1_STATUS>	NOT NULL	NUMBER(12)

Table 17–14 (Cont.) Dimension Table Structure

Name	Null	Type
<LEVEL1>	NOT NULL	VARCHAR2(30)
<LEVEL2_ID>	NOT NULL	NUMBER(15)
<LEVEL2>	NOT NULL	VARCHAR2(30)
<LEVEL3_ID>	NOT NULL	NUMBER(15)
<LEVEL3>	NOT NULL	VARCHAR2(30)

Here is an example for a SKU level dimension table from the sample configuration.

Table 17–15 RP_G_SKU_D (Example)

Name	Null	Type
SKU ID	NOT NULL	NUMBER(12)
SKU STATUS	NOT NULL	NUMBER(12)
SKU	NOT NULL	VARCHAR2(256)
STYL ID	NOT NULL	NUMBER(12)
STYL	NOT NULL	VARCHAR2(256)
CLSS ID	NOT NULL	NUMBER(12)
CLSS	NOT NULL	VARCHAR2(256)

Apart from the level dimension tables, the lowest level of each dimension has an age table that is used for purging functionality.

The dimension level age table is named as 'RP_G_<level_name>_AGE_D' and has following structure.

Table 17–16 RP_G_<level_name>_AGE_D

Name	Null	Type
<LEVEL_ID PRIMARY KEY>	NOT NULL	NUMBER(12)
STALE DATE	NULL Allowed	TIMESTAMP(6)

Here is an example for SKU level age table from the sample configuration.

Table 17–17 RP_G_SKU_AGE_D

Name	Null	Type
SKU ID	NOT NULL	NUMBER(12)
STALE DATE	NULL Allowed	TIMESTAMP(6)

Translation Table


The translation tables provide multi-language support. The translation table for each dimension level has the primary key and translated columns from the original dimension table and two special columns, LANGUAGE and SOURCE_LANG. Both LANGUAGE and SOURCE_LANG columns contain a value from the LANG column of the language reference table, RP_G_LANG_D.

LANGUAGE indicates the defined language of the row's translated columns. For example, when querying values for French, a program selects rows from the TL table where LANGUAGE = French in the RP_G_LANG_D.

The translation table stores translations. The column SOURCE_LANG stores the language code of the original language, and the LANG_CODE column stores the language code of the language of the translation. For example, if the label "Men's shirt" in English has a Chinese translation, then there is a record in the translation table for this Chinese translation, with SOURCE_LANG='English' and LANG_CODE='Chinese'.

Apart from one table for each dimension level, there is a table to store translations for each level label itself (RP_G_LEVEL_LABELS_TL); a table to store translations for each dimension label (RP_G_DIMENSION_LABELS_TL); and a table to store facts for each fact label. See Figure 17-7 for the translation tables.

Figure 17-7 Hybrid Storage Architecture ER Diagram for Translation



The translation table for each dimension level table is named as 'RP_G_<level_name>_TL_D' and has following structure:

Table 17-18 RP_G_<level_name>_TL_D

Name	Null	Type
<LEVEL_ID PRIMARY KEY>	NOT NULL	NUMBER(15)
LANG_CODE PRIMARY KEY	NOT NULL	VARCHAR2(20)
SOURCE_LANG	NULL Allowed	VARCHAR2(20)
<LEVEL_LABEL>	NULL Allowed	VARCHAR2(1000)

Here is an example of a SKU level translation table from the sample configuration.

Table 17-19 RP_G_SKU_TL_D

Name	Null	Type
SKU_ID	NOT NULL	NUMBER(12)
LANG_CODE	NOT NULL	VARCHAR2(20)
SOURCE_LANG	NULL Allowed	VARCHAR2(20)
SKU_LABEL	NULL Allowed	VARCHAR2(1000)

Index Translation Table

The Index Translation Tables (ITT) maintains the mapping between the position IDs and the RPAS integer indexes (also referred to as Physical ID) for all dimension level positions between the data mart and the RPAS domains. This mapping is important for efficient fact storage and optimal data transfer between the Data Mart and the RPAS domains.

Each table is partitioned by domain_id for faster access. Each level position has each domain's physical ID stored in the table. The column RENAMED is used internally by the renamePositions utility to communicate the information to other integrated domains.

ITT tables are named as 'RP_G_<level_name>_ITT' and have following structure:

Table 17-20 RP_G_<level_name>_ITT

Name	Null	Type
DOMAIN_ID	NOT NULL	NUMBER(12)
POSITION_NAME	NOT NULL	VARCHAR2(256)
PHYSICAL_ID	NOT NULL	NUMBER(12)
RENAMED	NOT NULL	NUMBER(12)

In addition to the ITT table, a materialized view is also provided for faster mapping, which provides direct mapping between the dimension level position ID and its physical ID.

The materialized view for a dimension level is named as 'RP_G_<level_name>_ITT_MV' and has following structure:

Table 17-21 RP_G_<level_name>_ITT_MV

Name	Null	Type
DOMAIN_ID	NOT NULL	NUMBER(12)
PHYSICAL_ID	NOT NULL	NUMBER(12)
SKU_ID	NOT NULL	NUMBER(12)

Fact Tables

Multiple facts at the same base intersection can be configured to belong to the same fact group. One fact table is created for one fact group. A fact table consists of the position IDs of intersection dimension levels and the facts belonging to the corresponding fact group. See [Figure 17-4](#).

N/A values are converted to NULL in the fact table to reduce data size.

Fact tables are named as 'RP_G_<fact_group>_FT' and have following structure:

Table 17-22 *RP_G_<fact_group>_FT*

Name	Null	Type
<LEVEL1_ID>	NULL Allowed	NUMBER(15)
<LEVEL2_ID>	NULL Allowed	NUMBER(15)
<LEVEL3_ID>	NULL Allowed	NUMBER(15)
<FACT1>	NULL Allowed	NUMBER(15)
<FACT2>	NULL Allowed	VARCHAR2(30)

Here is an example for the fact group GroupB from the sample configuration.

Table 17-23 *RP_G_GroupB_FT*

Name	Null	Type
SKU_ID	NOT NULL	NUMBER(12)
STR_ID	NOT NULL	NUMBER(12)
WEEK_ID	NOT NULL	NUMBER(12)
RECEIPT	NULL Allowed	NUMBER(12)
SALES	NULL Allowed	NUMBER(12)

Temporary and Logging Tables

A logging table, MERGE_LOG, logs SQL statements for the workbook commit process.

Schema Partitioning

The dimension and fact tables of RDM are partitioned based on the partition configuration provided at installation time. The ITT tables are partitioned based on the number of domains registered to the RDM.

Dimension and Fact Table Partitioning

The user data (dimensions and facts) is partitioned using the same logic as the RPAS domains: first, a partitioning level is chosen, then the positions on this level are grouped into partitions, and all levels lower and all facts at this level or lower inherit the same partitioning through the hierarchical roll-up relationships. The RDM schema partitioning cannot always be the same as that of the RPAS domains, especially when the domains are partitioned at different dimension levels or the partitioning positions are different. However, the proper mapping and alignment between the domains' partitions and the data mart's happens in such a way as to ensure that any local domain (of the multiple integrated domains) always maps to one or more data mart partitions, and any data mart partition is completely contained in one or zero local domain of any integrated domain. This partition mapping is essential for optimal performance in workbook operations, and RPAS batch processes that need to access the data mart because each data access (read or write) will have a well defined data source or destination.

In the data mart schema, the selected partition level dimension table is partitioned by list, and all its child dimension and fact tables are partitioned by reference.

A master metadata table RP_G_PARTITION_MD, as explained in the metadata tables in [Table 17-7](#), stores the mapping of the partition level positions and their corresponding partition IDs.

The partition level dimension table has an extra column for PARTITION_ID, corresponding to the same column in the RP_G_PARTITION_MD table. The RP_G_PARTITION_MD is provided to help with partition management. With this table, the partition level dimension table need not to be partitioned by the list of positions but is partitioned by the list of partition IDs.

ITT Table Partitioning

The Index Translation Tables are also partitioned by each integrated domain. In other words, each domain has its own partition of the ITT, so that each process initiated from any domain only needs to access one partition, which is a small portion of the entire table. Also, fast data updating and refreshing are performed at the partition level.

Sequence Descriptions

Sequences follow the naming convention of 'RP_G_<short_description or table_name>_S'.

[Table 17–24](#) lists all the sequences used by the data mart schema, along with their descriptions.

Table 17–24 Data Mart Schema Sequences

Sequence Name	Description
RP_G_<level_name>_D_S	Each dimension level table has a corresponding sequence to automatically generate the position ID.
RP_G_DOMAININFO_MD_S	This sequence is used to generate the domain ID for the RP_G_DOMAININFO_MD table.
RP_G_DPMP_S	This sequence is used for the DPM process to generate a unique sequence ID.
RP_G_DPMPH_S	This sequence is used to generate a placeholder position ID for the DPM process.
RP_G_ETL_S	This sequence is used by the ETL process to create the staging tables.
RP_G_RDM_STATUS_S	This is used to generate SEQ_ID in the RP_G_RDM_STATUS_MD table.
RP_G_STAGING_S	This sequence is used to generate unique staging table names internally by loadFactData and transferFactData utilities.

Package Descriptions

The PL/SQL packages are stored in the data mart schema and are named as 'RP_<functional_description>_PKG'.

The packages are described in [Table 17–25](#).

Table 17–25 Data Mart Schema Package Descriptions

Package Name	Description
RP_ETL_PKG	This package is used for the external data integration of non-RPAS applications with the RPAS application.
RP_G_DATALOAD_API_PKG	This package has the common APIs used for data load functionalities, such as loadFactData, transferFactData and external integration.

Table 17–25 (Cont.) Data Mart Schema Package Descriptions

Package Name	Description
RP_GLOBAL_PKG	This is a common package that has some static definitions and procedures that are used globally.
RP_GUTS_PKG	This package is used by SE diagnostic tool's checkRDM validations.
RP_HIERMGR_PKG	This package handles the hierarchy operations such as DPM, copy, and rename positions.
RP_LOAD_DIMENSION_PKG	This is the package used by loadDimData utility.
RP_LOAD_FACT_PKG	This package is used by loadFactData utility.
RP_PARTITIONMGR_PKG	This is the package used by rmPartitionMgr utility.
RP_PATCH_RDM_PKG	This is the package used by the patchRDM utility.
RP_SCHEMA_INFO_PKG	This is a common package that has procedures to provide some information about the data mart schema and is used by various other packages internally.
RP_STATUS_PKG	This package is used by various interdependent utilities to maintain RDM status.
RP_TRUNCATE_TABLE_PKG	This package is used internally by loadHier and reconfigGlobalDomainPartitions.
RP_WORKBOOK_COMMIT_PKG	This package is used by the workbook commit process.

Access Schema Descriptions

RPAS Data Mart consists of a central data mart schema and seven access schemas. The access schemas are based on different RPAS functionality. Basically, a group of functionalities are performed through an access schema.

The access schema names and properties are configured by the installer. But RPAS refers to them with a fixed alias. The seven aliases, along with their functionality, are summarized in [Table 17–26](#).

Table 17–26 Data Mart Schema Access Schemas

Access Schema	Functionality
rpas_patch_user	RDM patch
rpas_batch_user	Batch processing
rpas_dimload_user	Loading dimension data
rpas_factload_user	Loading fact data
rpas_hiermgr_user	Hierarchy operations such as DPM
rpas_wkbk_user	Workbook operations such as workbook build and commit
rpas_etl_user	External data integration between RPAS and non-RPAS applications

Each access schema contains private synonyms for all the Data Mart schema objects it needs access to. Each schema is granted a minimum set of privileges through roles. Only required data mart schema objects are visible, and only required operations are allowed from each access schema. RPAS processes internally use the right access schema based on the functionality.

Some access schemas such as `rpas_dimload_user` or `wpas_wkbk_user` create temporary staging tables in their own schemas for the duration of the RPAS process.

Security Considerations

RPAS Data Mart's security uses a role-based security model. The seven access schemas have defined privileges granted through pre-defined roles by the RPAS code. The access schemas with their corresponding roles are summarized in [Table 17-27](#).

Table 17-27 Data Mart Access Schemas

Access Schema	Role
<code>rpas_patch_user</code>	<code>rpas_patch_role</code>
<code>rpas_batch_user</code>	<code>rpas_batch_role</code>
<code>rpas_dimload_user</code>	<code>rpas_dimload_role</code>
<code>rpas_factload_user</code>	<code>rpas_factload_role</code>
<code>rpas_hiermgr_user</code>	<code>rpas_hiermgr_role</code>
<code>rpas_wkbk_user</code>	<code>rpas_wkbk_role</code>
<code>rpas_etl_user</code>	<code>rpas_etl_role</code>

For additional details on the security model, see the *Oracle Retail Predictive Application Server and Applications Security Guide*.

External Integration

This section describes the process of data integration from non-RPAS application to RDM.

RPAS provides PL/SQL procedures to create staging tables/global temporary tables and merge staged data into the fact tables. As a onetime setup process, all staging tables or GTTs can be created together. The staging tables are created in the user schema (ETL schema) from where the RPAS API is called.

As part of the execution step, the staging tables/global temporary tables are loaded with data and merged into RDM fact tables using RPAS API. The loading and merging can be repeated as and when required.

Setup

Run RPAS API to create staging table or global temporary tables in the user schema with the list of all fact names to be integrated. The API creates one table per fact group for optimal merge performance and returns the list of table names. It is recommended to provide all the fact names in one run, so that the API can optimally group them into tables.

Global temporary tables are always created with position names and fact columns.

The staging tables can be created in two ways.

- Staging table with position names and fact values
- Staging table with mapped position IDs and fact values

	Staging With Position Names	Staging With Position IDs
Position Name-ID Mapping	Done by RPAS procedures	Done by integrator
Performance	Overall performance gets affected	Complete process performs better
Table Properties	No constraints No partitioning	Reference-partitioned just like the target fact table in the RDM

Figure 17–8 Simplified ERD for Staging Tables with Position Names





Figure 17–9 Simplified ERD for Staging Tables with Mapped Position IDs



The staging tables thus created are owned by the integrator who can script the integration process using the fixed staging table names.

Execution

Integrator loads data in the staging table or global temporary tables and then calls the RPAS API to merge staged data with the input list of staging and/or temporary table names.

The staged data can be merged in two ways.

- SQL Merge
- Partition Exchange

	SQL Merge	Partition Exchange
Transaction Protection	Transaction protected No contention	Transaction not protected
Recommended When	Low volume data	High volume bulk data

PL/SQL APIs The procedures to create staging tables and merge staged data are provided in the `rp_etl_pkg` package. The details are as follows.

Table 17–28 Staging Table Procedures

Procedure	Description
Function <code>create_staging_tables</code> (<code>fact_list_in</code> IN <code>varchar2_table</code> , <code>by_name_in</code> IN integer, <code>stg_table_list_out</code> OUT <code>varchar2_table</code>) Return number;	- Creates staging tables in the ETL schema for the fact names passed in. - Returns the name of the tables created and status (0 if successful, 1 otherwise). - One staging table is created for the facts belonging to one fact group for optimal performance. - The fact names that are not defined in rpas metadata are ignored. - The <code>by_name_in</code> parameter determines if the staging tables are created with position names or IDs. - When <code>by_name_in</code> = 1, staging table will be created with names; when <code>by_name_in</code> = 2, staging table will be created with IDs. - Naming convention 'RP_'<FactGroup>_'ST'<SequenceNumber>
Function <code>create_global_temp_tables</code> (<code>fact_list_in</code> IN <code>varchar2_table</code> , <code>ggt_list_out</code> OUT <code>varchar2_table</code>) Return number;	- Creates global temporary tables in the ETL schema for the fact names passed in. - Returns the name of the tables created and status (0 if successful, 1 otherwise). - One GTT is created for the facts belonging to one fact group for optimal performance. - The fact names that are not defined in rpas metadata are ignored. - Naming convention 'RP_'<FactGroup>_'GT'<SequenceNumber>
Function <code>merge_staged_data</code> (<code>table_list_in</code> IN <code>varchar2_table</code> , <code>merge_method_in</code> IN integer) Return number;	- The data staged in the staging tables or Global temporary tables provided, is merged in the target fact tables. - The input table list can be a mixture of staging tables and GTTs. - Only valid records are loaded into the RDM fact tables, for example, the records with position names not existing in RDM are not loaded into the RDM fact tables. - Returns a status of 0 if successful, 1 otherwise. - The <code>merge_method_in</code> parameter determines the actual merge method: SQL MERGE or Partition Exchange. - When <code>merge_method_in</code> = 1, SQL MERGE will be used; when <code>merge_method_in</code> = 2, Partition Exchange will be used. - Only SQL Merge method can be used with GTT. - If the input table is a GTT, the actual <code>merge_method_in</code> value is ignored. - The procedure does not issue commit in the end. However, when Partition-Exchange method is used or the staging table is created with IDs, there is an implicit commit in the procedure itself.

`varchar2_table` is created as

create or replace type varchar2_table force is table of varchar2(40);

Note: Note the following:

- The External Integration can be performed by the pre-defined "rpas_etl_user" access schema or any other schema as long as it meets the following requirements:
 - The schema must have "CREATE TABLE" direct privilege.
 - The schema must be granted the pre-created role "rpas_etl_role".
 - Fact NA values are stored as 'null' in the RDM fact tables
 - The Fact NA property is configurable and can be changed via RDM patching process.
-


Configuring the RPAS Data Mart

This section provides details about configuring the RPAS Data Mart.

RDM Repository

The RDM repository is a directory on the file system used to store the RDM configuration, input files, log files, and metadata. Its directory structure is very similar to that of an RPAS domain, except that no dimension or measure data is stored there.

The RDM repository is created by rdmMgr when the RDM is built as part of the pre-installation process, and it is updated when the RDM is patched. It is a centralized location for RDM properties, Oracle database credential wallets, and network configuration files. In every integrated domain, the property called "rdm_path" contains the path to this repository. [Figure 17–10](#) shows the relationship between the domain and RDM repository.

Figure 17–10 Relationship Between RPAS Domain and RDM Repository

The location to create the RDM repository is specified when running the RDM schemas preparation script during the pre-installation process.

Directory Structure

Figure 17–11 shows the directory structure of the RDM repository.



Figure 17–11 RDM Repository Directory Structure

Table 17–29 RDM Repository Directory Structure

Directory	Contents	Populated By	Processed By
config	Schema information file and integration and partition configuration files	rdmMgr copies the schema information file from the user-specified location to this location with a fixed name after creating the RDM repository. rpasInstall copies the configuration files from user-specified locations to this location with fixed names.	rdmMgr uses the files from here.
input	Shared hierarchies and measures data input files	The input dimension input files must be manually copied to the dimdata folder beneath and fact data files in the factdata folder.	loadDimData (using dimdata subdirectory) and loadFactData (using factdata subdirectory) loads data from here and moves to a processed folder beneath
data/admin.db	properties.ary to store the data mart and user schema names	rpasInstall/rdmMgr	Any operation connecting to the database
scripts	Any script output of RDM manager.	RDM Manager	DBA
plsql	Plsql packages and procedures	rpasInstall/buildSchema	rpasInstall/loadSqlPackages.ksh
output	RDM log files	RDM utilities	
wallet	Oracle database credential and SSL wallet	prepareRDM.ksh	Any operation connecting to the database
tns_admin	Oracle database network configuration (tnsnames.ora and sqlnet.ora)	RDM Manager	Any operation connecting to the database

RDM Properties in a Domain

In an integrated domain, all RDM-related properties are encapsulated in an array called `rdm_properties.ary`. [Figure 17–12](#) shows the RDM-related entities in a domain.

Figure 17–12 RDM Entities in a Domain**Table 17–30 RDM Entities in a Domain**

Directory	Contents	Populated By
data/admin.db	New array <code>rdm_properties.ary</code> to store <code>domain_name</code> and <code>rdm_path</code>	rpasInstall/rdmMgr

Schema Information

In order to implement RDM schema security, the DBA and the RPAS administrator need to work together and set up a multi-schema environment.

Table 17–31 lists all schemas and their corresponding roles.

Note: The schema names can be changed, but the role names and database aliases (DB connect strings) are constants.

Table 17–31 Schemas and Roles

Schema	Schema ID	Recommended Name	Role	DB Connect String	Notes
RDM Admin Schema	rpas_data_mart	rpas_data_mart	N/A	rpas_data_mart_conn	The main schema for all persistent data
RDM Patch User	rpas_patch_user	rpas_patch_user	rpas_patch_role	rpas_patch_conn	
RPAS Batch User	rpas_batch_user	rpas_batch_user	rpas_batch_role	rpas_batch_conn	
RDM Dim Loader	rpas_dimload_user	rpas_dimload_user	rpas_dimload_role	rpas_dimload_conn	
RDM Fact Loader	rpas_factload_user	rpas_factload_user	rpas_factload_role	rpas_factload_conn	Requires space for staging tables
RPAS Hier Manager	rpas_hiermgr_user	rpas_hiermgr_user	rpas_hiermgr_role	rpas_hiermgr_conn	
RPAS Workbook User	rpas_wkbk_user	rpas_wkbk_user	rpas_wkbk_role	rpas_wkbk_conn	Requires space for staging tables
RPAS ETL Schema	rpas_etl_user	rpas_etl_user	rpas_etl_role	rpas_etl_conn	Requires space for staging tables

The schema names and Oracle database network configuration can be configured in the schemaInfo.xml file that is passed to the RDM schemas preparation script.

The format of the schemaInfo xml is as follows:

```
<?xml version="1.0" encoding="UTF-8" ?>
<rdm_configuration>
  <schema_set>
    <schema id="rpas_data_mart">
      <name>rpas_data_mart</name>
      <tablespace>
        <name></name>
        <space_quota></space_quota>
      </tablespace>
    </schema>
    <schema id="rpas_patch_user">
      <name>rpas_patch_user</name>
      <tablespace>
        <name></name>
        <space_quota></space_quota>
      </tablespace>
    </schema>
  </schema_set>
</rdm_configuration>
```

```

        </schema>
        <schema id="rpas_batch_user">
        <name>rpas_batch_user</name>
        <tablespace>
        <name></name>
        <space_quota></space_quota>
        </tablespace>
        </schema>
        <schema id="rpas_dimload_user">
        <name>rpas_dimload_user</name>
        <tablespace>
        <name></name>
        <space_quota></space_quota>
        </tablespace>
        </schema>
        <schema id="rpas_factload_user">
        <name>rpas_factload_user</name>
        <tablespace>
        <name></name>
        <space_quota></space_quota>
        </tablespace>
        </schema>
        <schema id="rpas_hiermgr_user">
        <name>rpas_hiermgr_user</name>
        <tablespace>
        <name></name>
        <space_quota></space_quota>
        </tablespace>
        </schema>
        <schema id="rpas_wkbk_user">
        <name>rpas_wkbk_user</name>
        <tablespace>
        <name></name>
        <space_quota></space_quota>
        </tablespace>
        </schema>
        <schema id="rpas_etl_user">
        <name>rpas_etl_user</name>
        <tablespace>
        <name></name>
        <space_quota></space_quota>
        </tablespace>
        </schema>
    </schema_set>
    <tns_parameters>
        <protocol>tcp</protocol>
        <host>burr41002v.us.oracle.com</host>
        <port>1521</port>
        <server>dedicated</server>
        <service_name>rpasodb.us.oracle.com</service_name>
        <sid></sid>
    </tns_parameters>
</rdm_configuration>

```

The schema IDs must not be changed. The names can be changed. The tablespace properties, name and space_quota, are optional. If no tablespace is provided, a default tablespace of the database is used. If no space_quota is provided, no table space quota will be granted to the user. The space_quota must be accompanied by the table space name. However, the prepareSchemas.sql can be modified manually before executing.

Note that <service_name> or <sid> must be present and only one can be present.

This file must be provided while creating the RDM repository described in [RDM Repository](#).

Integration Configuration

An integration tool within RPAS ConfigTools creates the integration configuration from user inputs and one or more domain configurations. A name is specified for the integration configuration, which is saved as the name attribute of the highest level XML element. It is also used in the file name of the integration configuration. The example described in this section shows an integration configuration with a name specified as "RDM".

The language attribute is used to specify the language of all the labels in the integration configuration. It is used as the source language in the translation tables for those labels. It can be any of the 18 languages supported by RPAS.

Integration configuration consists of four sections.

Shared Hierarchies

The shared hierarchy section defines the hierarchical structure of the RDM, and it follows a similar XML structure to that of the hierarchy.xml in a domain, but much simplified.

This section consists of multiple hierarchies, each forming a dimension graph. Each hierarchy has a name, a label, and an order number that is used to define the order of multiple dimensions in an intersection. Each dimension has a name and a label, and the parent-child relationship of the XML structure defines the aggregation and spread relationship of the dimensions.

Shared Facts

A fact in RDM is the corresponding data entity of a measure in a domain. The properties for a fact are name, label, base intersection, type, fact table, NA value, description, and purge age.

If a fact is shared by the RDM, a fact group name must be specified. Otherwise, the fact is shared by domains only.

Integration Map

Each entry defines the mapping between one domain measure and one database fact.

Domain and Configuration Map

This section defines the domain and configuration map. Each domain is mapped to a domain configuration. The domain name is configured here and is by default the same as the configuration name. This section is used internally by the integration tool of the Config Tools.

Example

Here is a simple example (RDM.xml) of the integration.xml file that shows all the XML elements discussed in the above sections.

```
<?xml version="1.0" encoding="UTF-8" ?>
<rpas_hsa_configuration name="RDM" language="English">
  <shared_hierarchy_set>
    <hierarchy name="CLND">
```

```

        <label>Calendar</label>
        <order>9990</order>
        <dimension name="day">
            <label>Day</label>
            <position_format>d%Y%m%d</position_format>
            <dimension name="week">
                <label>Week</label>
                <dimension name="mnth">
                    <label>Fiscal Month</label>
                    <dimension name="qrtr">
                        <label>Fiscal Quarter</label>
                        <dimension name="year">
                            <label>Fiscal Year</label>
                        </dimension>
                    </dimension>
                </dimension>
            </dimension>
        </dimension>
        </dimension>
        </dimension>
        </hierarchy>
        <hierarchy name="PROD">
            <label>Product</label>
            <order>9980</order>
            <dimension name="sku">
                <label>Sku</label>
                <dimension name="clss">
                    <label>Class</label>
                    <dimension name="dept">
                        <label>Department</label>
                    </dimension>
                </dimension>
            </dimension>
        </hierarchy>
    </shared_hierarchy_set>
    <shared_fact_set>
        <fact name="P">
            <label>P Fact</label>
            <baseint>sku_week</baseint>
            <type>2</type>
            <fact_group>r_g_sku_week_ft</fact_group>
            <navalue>0</navalue>
            <description>PPP Fact</description>
            <purgeage>1000</purgeage>
        </fact>
        <fact name="Q">
            <label>Q Fact</label>
            <baseint>sku_week</baseint>
            <type>4</type>
            <fact_group>r_g_sku_week_ft</fact_group>
            <navalue>0</navalue>
            <description>QQQ Fact</description>
            <purgeage>1000</purgeage>
        </fact>
    </shared_fact_set>
    <integration_map>
        <entry>
            <domain>d1</domain>
            <measure>a</measure>
            <fact>P</fact>
        </entry>
        <entry>

```

```
<domain>d2</domain>
<measure>a</measure>
<fact>P</fact>
</entry>
<entry>
  <domain>d1</domain>
  <measure>c</measure>
  <fact>Q</fact>
</entry>
</integration_map>
<domain_set>
  <domain name="d1">
    <config_path>...</config_path>
  </domain>
  <domain name="d2">
    <config_path>...</config_path>
  </domain>
  <domain name="d3">
    <config_path>...</config_path>
  </domain>
</domain_set>
</rpas_hsa_configuration>
```

Partition Information

The RDM partition information XML file is modeled after the global domain configuration. It contains the partition dimension name and multiple partition definitions, each with a partition name and its corresponding list of positions.

This XML file is generated by the RDM Partition Manager from the domains. See [RDM Partitioning and RDM Partition Manager](#) for more information.

RDM SSL Connection

For best security practices in the connection between the RPAS Server utilities and the Oracle Database, we recommend enabling the optional SSL feature. The Oracle Database server natively supports SSL connections from clients, and the feature needs only to be enabled and configured for use with RPAS. Note that only one-way SSL is covered in this document. For two-way SSL setup, see the *Oracle Database Security Guide* at the following URL:

<http://docs.oracle.com/database/121/DBSEG/asossl.htm>


SSL Overview

The following diagram shows the components of the SSL configuration for RDM. The client/server connection will be protected by one-way SSL, which only authenticates the server on the client side. As a result, only the Oracle server wallet is required to have a private key. The client wallet will only need to contain the server CA (Certificate Authority) certificate.

The CA certificate can be either self-signed or issued by a third-party CA. In the case of third-party CA, more than one certificate may need to be imported into the client wallet (usually called CA certificate chain).

The server side configuration is mostly done manually by the DBA. A script is provided to create the wallet if self-signed certificate is used.

The client side configuration is done through scripts by the RPAS admin.

Figure 17-13 SSL Configuration for RDM

Set Up SSL on Oracle Server

This section contains information on setting up SSL on the Oracle server side. This should be done by a DBA who has the permission to modify the configuration files of the Oracle server.

Create Oracle Server Wallet The Oracle server wallet can be created the same way as the RPAS server wallet for one-way SSL. The process has been documented in great detail in the *Oracle Retail Predictive Application Server Administration Guide for the Classic Client* "SSL" chapter. On UNIX, the wallet must be created by the user account which starts the Oracle database processes and for security reason the wallet is only accessible by its creator.

If self-signed CA certificate is used, the shell script `createSSLWallets.sh` provided by RPAS can be used to create the root and server wallets, private keys and all related certificates.

When running the script `createSSLWallets.sh`, choose option 3 to create wallets for Oracle database server, as shown in the following figure:

Figure 17-14 Create Wallet Select Option

```

zsh.exe - Shortcut
/cygdrive/c/work/rpas/DomainUtilities/scripts>createSSLWallets.sh
Create Oracle Wallets and Self-Signed Certificates for RPAS SSL.

Please choose one of the following options(1-3):
  1. Create wallets for RPAS one-way SSL (type 1)
  2. Create wallets for RPAS two-way SSL (type 2 or 4)
  3. Create wallets for Oracle database server


Enter 1-3:
3
  
```

The script asks for the root directory where the wallets will be created, your organization name (any name such as "QA"), and passwords for the root and server

wallets. In the end of a successful run, the wallet location and the path to the CA certificate file will be displayed to the user (see [Figure 17-15](#)).

The CA certificate file can be copied to a common location and must be imported into the Oracle client wallet later on.

Figure 17-15 Create Wallet Complete



Update Oracle Server Network Configuration The network configuration file `sqlnet.ora` normally resides in `$ORACLE_HOME/network/admin` directory.

The following table lists the required settings for SSL.

Required SSL Settings in `sqlnet.ora` on Oracle server

```
SSL_CLIENT_AUTHENTICATION = FALSE
WALLET_LOCATION =
(SOURCE =
(METHOD = FILE)
(METHOD_DATA =
(DIRECTORY = wallet_location)
)
)
```

The SSL client authentication parameter must be set to `FALSE` for one-way SSL. The `wallet_location` is the absolute path to the Oracle server wallet as discussed in the previous section.

The following table lists the optional settings for SSL.

Optional SSL Settings in `sqlnet.ora` on Oracle server

```
SSL_CIPHER_SUITES= (SSL_cipher_suite1 [SSL_cipher_suite2])
SSL_VERSION=3.0
```

For more information regarding these settings, refer to *Oracle Database Security Guide* at the following URL:

<http://docs.oracle.com/database/121/DBSEG/asossl.htm>

Update Oracle Listener Configuration The listener configuration file listener.ora normally resides in \$ORACLE_HOME/network/admin directory. After modification, the Oracle listener must be restarted for the changes to take effect.

The following table shows the changes in bold.

Required SSL Settings in listener.ora on Oracle server

SSL_CLIENT_AUTHENTICATION = FALSE

WALLET_LOCATION =

(SOURCE =

(METHOD = FILE)

(METHOD_DATA =

(DIRECTORY = wallet_location)

)

)

LISTENER =

(DESCRIPTION_LIST =

(DESCRIPTION =

(ADDRESS = (PROTOCOL = IPC)(KEY = EXTPROC1521))

)

(DESCRIPTION =

(ADDRESS = (PROTOCOL = TCP)(HOST = odbsvr-host)(PORT = 1521))

)

(DESCRIPTION =

(ADDRESS = (PROTOCOL = TCPS)(HOST = odbsvr-host)(PORT = 2484))

)

)

The wallet location and SSL settings are the same as that in the network configuration file (must be in both files). A new listener endpoint with protocol TCPS must be added to the LISTENER setting. Port number 2484 is the typical port for TCPS but other number can also be used.

Set up SSL on Oracle Client This section contains information on setting up SSL on the Oracle client side for use by RPAS or RDM processes.

Update Schema Info Configuration The schema info configuration must have matching endpoint parameters to the Oracle listener. The protocol must be tcps and the port number must be the same as specified in the listener configuration.

TNS Parameters in schemaInfo.xml

```


<tns_parameters>
  <protocol>tcps</protocol>
  <host>odbserver-host</host>
  <port>2484</port>
  <server>dedicated</server>
  <service_name>myservice</service_name>
  <sid></sid>
</tns_parameters>

```

Import Server CA Certificate To import the server CA certificate:

1. The prepareRDM.ksh script in the RDM build process already creates an Oracle wallet to hold the login credentials for RDM. The same wallet should be used to hold the Oracle server CA certificate. If protocol tcps is specified in the schema info, this same script will ask for the path to the CA certificate file and import it into the wallet automatically.
2. The following figure shows the importing of CA certificate by the prepareRDM.ksh script.

Figure 17–16 Import CA Certificate



```

zsh.exe - Shortcut
rpas_hiermgr_conn hiermgr_localmark
Last command was successful.
=== Add login credential for schema wkbk_localmark to the wallet
Command: mkstore.bat -url C:/work/rpas/domains/MFP_Only/wallet -createCredential
rpas_wkbk_conn wkbk_localmark
Last command was successful.
=== Add login credential for schema etl_localmark to the wallet
Command: mkstore.bat -url C:/work/rpas/domains/MFP_Only/wallet -createCredential
rpas_etl_conn etl_localmark
Last command was successful.

-----
Protocol tcps is specified. SSL will be enabled.
Enter path to the Oracle server CA certificate file:
c:/work/RDM/root_chain.txt
Oracle server CA certificate has been imported successfully.
-----
Script prepareRDM.ksh completed successfully.

```

If third-party CA is used and there is more than one file in the CA certificate chain, the user must specify the top certificate for the script to import and manually import the rest in the order of the chain using the following command. The client wallet location is the wallet subdirectory under the RDM repository.

```
orapki wallet add -wallet {client_wallet_directory} -trusted_cert -cert {ca_cert_chain_file} -pwd {client_wallet_password}
```

Example:

```
orapki wallet add -wallet C:/wallets/client -trusted_cert -cert C:/wallets/ca_chain2.txt -pwd clientpass1
```

Update Oracle Client Network Configuration The network configuration file `sqlnet.ora` used by RPAS processes is created automatically by the RDM Manager during the RDM build process. This file resides under the `tns_admin` subdirectory of the RDM repository.

The following SSL settings are required.

Required SSL Settings in `sqlnet.ora` on Oracle Client

```
SSL_CLIENT_AUTHENTICATION = TRUE
```

```
WALLET_LOCATION =
```

```
(SOURCE =
```

```
(METHOD = FILE)
```

```
(METHOD_DATA =
```

```
(DIRECTORY = wallet_location)
```

```
)
```

```
)
```

The SSL client authentication parameter must be set to `TRUE` for one way SSL. The `wallet_location` is the wallet subdirectory of the RDM repository.

Update Oracle Net Service Names The Oracle Net Service Names configuration file `tnsnames.ora` is created automatically by the RDM Manager during the RDM build process under the `tns_admin` subdirectory of the RDM repository. All entries use the endpoint parameters specified by the schema info configuration.

Net Service Names using SSL in `tnsnames.ora` on Oracle Client

```
rpas_data_mart_conn =
```

```
(DESCRIPTION =
```

```
(ADDRESS_LIST =
```

```
(ADDRESS = (PROTOCOL = tcps)(HOST = odbsvr-host)(PORT = 2484))
```

```
)
```

```
(CONNECT_DATA =
```

```
(SERVER = dedicated)
```

```
(SERVICE_NAME = myservice)
```

```
)
```

```
)
```

```
.
```

```
.
```

```
.
```

Test and Confirm SSL Connection After SSL configuration is done on both Oracle server and client and the Oracle listener has been restarted, the user can use RDM Manager to verify the connection. (There are a small delay, sometimes a few minutes, between the listener is restarted and the SSL connection is up.)

```
rdmMgr -rdm rdmPath -testConnection
```

RPAS HSA Build Process

The RPAS HSA build process can be broken into three major steps: pre-installation, installation, and post-installation.

Pre-installation is required because the DBA's intervention is necessary to create schemas and roles. Unlike any other traditional installer, database schemas are created as part of the RPAS HSA installation process. The pre-installation process prepares the RDM repository, Oracle wallet, and SQL script to create database schemas and roles. The SQL script has to be manually executed by the DBA.

The installation process builds schema objects in the database schemas and grants privileges to each role.

The post-installation process initializes the data synchronization and registers compatible domains.

RPAS Installation Utilities

This section covers the main RPAS installer and all helper utilities and scripts.

RDM Preparation Script

A shell script is created to configure the RDM before building the database schemas. The script is interactive but can be made silent if all parameters are specified by environment variables.

Note: Note that if SSL is to be used for the RDM database connection, it must be configured first before running the script.

Script

prepareRDM.ksh [rdmPath] [schemaInfoFile]

Functions

1. Create RDM file system repository.
 - Call `rdmMgr -create` in order to create the RDM repository.
2. Generate RDM pre-build scripts and network configuration files.
 - Call `rdmMgr -genPreBuildScript` in order to create pre-installation script and network configuration files.
 - The resulted files are put in the RDM repository:
 - `prepareSchemas.sql` under `scripts`. You are prompted for a password when executing this sql file.
 - `tnsnames.ora` and `sqlnet.ora` under `tns_admin`
3. Create Oracle wallet for RDM.
 - Create an Oracle wallet `rdmPath/wallet`.
4. Add all RDM user credentials in the wallet.
 - You are prompted for a password for these schema logins.

The password for each schema in the wallet must match that of the corresponding schema created in the database.

RPAS Installer

The existing rpaInstall program supports building an HSA system.

Command Line Syntax

rpaInstall {command} {options}

rpaInstall

rpaInstall -buildRDM -genScript -integrationCfg <integration_config> {-partitionInfo <partitionInfo_file> | -noPartition} -log <log_file> -rdm <rdm_path>

rpaInstall -buildRDM -integrationCfg <integration_config> {-partitionInfo <partitionInfo_file> | -noPartition} -rdm <rdm_path> -log <log_file>

rpaInstall -buildRDM -fromScript -rdm <rdm_path> -log <log_file>

rpaInstall -rebuildRDM -integrationCfg <integration_config> {-partitionInfo <partitionInfo_file> | -noPartition} -rdm <rdm_path> -in <input_directory> -log <log_file>

rpaInstall -rebuildRDM -fromScript -rdm <rdm_path> -log <log_file>

Command Values

-buildRDM | -rebuildRDM: the type of installation to be performed for integration deployment

Option Values

Table 17–32 Option Values

Option	Description
-integrationCfg <config_file>	Option for the integration configure file
-partitionInfo <partitionInfo_file>	Option for the partition information input file
-noPartition	Option if no partition information is provided
-rdm <path>	Option for the RDM repository path
-genScript	Only generates the SQL script with RDM schema definitions, but does not deploy it. Two script files named buildDataMart.sql and buildUserSchemas.sql are generated under the RDM repository subdirectory scripts. One file script creates the Data Mart schema's schema objects and other is for the other users' schema objects.
-fromScript	Builds or patches RDM using the scripts buildDataMart.sql and buildUserSchemas.sql under the RDM repository subdirectory scripts. This option lets the database administrator customize the RDM creation statements and build.
-log <log_file>	For logging. All logging output from subprocesses must be redirected to this file. At the end of the rpaInstall run, a copy of this file must be copied to {rdm_repository}/output/rpaInstall with timestamp in the file name.
-verbose	Results in more detailed log information.

RDM Manager

The RDM Manager is a command line utility that manages operations of the RPAS Data Mart (RDM). It creates the RDM repository directory structure and the RDM

database schema from an integration configuration and optional partitioning information. It also manages domain registration and provides some other useful tools to the RDM in an integrated environment.

Configuration Files and Output Files Note that the integration configuration file and partition information file are not specified on the command lines. They must be pre-staged under the config directory of the RDM repository, usually by the `rpasInstall` utility. The scripts and network configuration files generated by the RDM manager all have pre-defined output directories and names.

Pre-staged Configuration Files

- Integration configuration
 - Copy to `rdmPath/config/integrationCfg.xml` in advance.
- Partition Information
 - Copy to `rdmPath/config/partitionInfo.xml` in advance.

Script Output Files

- All output scripts go to `rdmPath/scripts/`
- The scripts have fixed names.
 - Pre-build script: `prepareSchemas.sql`: Sql script to create users and corresponding roles with few basic privileges.
 - Build scripts: `buildDataMart.sql`: Sql script to create all data mart schema objects and `buildUserSchemas.sql`: Sql script to create private synonyms. This script must be run for each schema other than the data mart.
 - Build script: `grantPrivileges.sql`: Sql script to grant all required privileges to each access schema. This must be re-executed only when the schema objects are created for the first time or when existing schema objects are recreated.

Oracle Database Network Configuration Files

- The following files go under `rdmPath/tns_admin`
 - `tnsnames.ora`
 - `sqlnet.ora`

Command Line Syntax

Generic command syntax

`rdmMgr -rdm {rdmPath} {command} [options]`

Commands used by the RDM preparation script

`rdmMgr -rdm rdmPath -create -schemaInfo f1`

`rdmMgr -rdm rdmPath -genPreBuildScript`

Commands used by the RPAS Installer

`rdmMgr -rdm rdmPath -clearSchema`

`rdmMgr -rdm rdmPath -buildSchema`

`rdmMgr -rdm rdmPath -genBuildScript`

`rdmMgr -rdmPath -buildFromScript`

Commands that can be used by the RPAS Administrator

```
rdmMgr -rdm rdmPath -listProperties
rdmMgr -rdm rdmPath -testConnection
rdmMgr -rdm rdmPath -register -d domainPath -name domainName
rdmMgr -rdm rdmPath -unregister -name domainName
rdmMgr -rdm rdmPath -listDomains
rdmMgr -rdm rdmPath -listHistory
rdmMgr -rdm rdmPath -purgeDeletedData
rdmMgr -rdm rdmPath -updateRegistration -d domainPath -name domainName
rdmMgr -rdm rdmPath -move -dest destRdmPath
```

Commands**Table 17–33** *Commands*

Command Name	Description
-create	Creates the RDM directory structure and copies the schemaInfo file to the config directory beneath.
-listProperties	Lists the RDM properties.
-testConnection	Tests all schema connections.
-clearSchema	Drops any existing RDM schema objects in all schemas.
-buildSchema	Builds the RDM schema objects in all schemas and grant privileges to the roles created initially. Assumes integrationCfg.xml and partitionInfo.xml are present under the config directory of the RDM repository.
-genPreBuildScript	Generates the SQL script that can be used to create schemas and roles. The generated script prepareSchemas.sql is put under the scripts directory of the RDM repository.
-genBuildScript	Generates the SQL script that can be used to create the schema objects, populate some metadata, and grant privileges to the roles initially created. Assumes integrationCfg.xml and partitionInfo.xml are present under the config directory of the RDM repository. The generated scripts buildDataMart.sql, buildUserSchemas.sql, and grantPrivileges.sql are put under the scripts directory of the RDM repository.
-buildFromScript	Builds the RDM schema objects and grant privileges using the sql scripts buildDataMart.sql, buildUserSchemas.sql and grantPrivileges.sql generated by genBuildScript.
-register	Registers the Rpas domain with RDM.
-unregister	Unregisters the Rpas domain from RDM.
-listDomains	Lists the current domains and status.
-listHistory	Lists the RDM schema and RDM Repository history.
-purgeDeletedData	Physically purges all dimension data that is marked for deletion, along with all associated fact data.
-updateRegistration	Updates the existing registration of a domain. Depending on what has changed or has become out of sync, it updates domain path, RDM path or ITT tables.

Table 17–33 (Cont.) Commands

Command Name	Description
-move	Moves the RDM repository to a new location and updates all domains with the new RDM path. If any of the domains cannot be updated, a warning message will be displayed. The user can run the -updateRegistration command to update the RDM path in the domain separately.

Parameters**Table 17–34 Parameters**

Parameter Name	Description
-d	Path to the domain.
-rdm	Path to RDM.
-schemaInfo	Path to the schema information file (this file is copied to the config directory under RDM repository afterwards for future reference).
-name	The domain name as defined in the integration configuration.
-dest	The destination path for moving the RDM repository.

Create RDM Repository

The following is completed during this process:

- The RDM Repository directory structure is created.
- The schema names in the schemaInfo.xml are saved as RDM properties.

Build RDM Schema

The following RDM objects are created:

- Dimension tables
- Fact tables
- Metadata Tables
- Data flow table
- Index Translation Tables (ITT)
- Commit, logging, and dimension staging tables
- Other ODB objects (such as sequences, views, indexes, types, and pl/sql packages)
- Along with these, private synonyms for each schema user are also created.

Test ODB Connection

The connections to all ODB schemas are tested via the rdmMgr -testConnection command.

Register a Domain with RDM

The rdmMgr utility registers a domain at a time. The prerequisites to register a domain are:

1. The domain to be registered with the RDM must have compatible hierarchies.

2. The domain's all informal positions must be formalized or removed.

The following must be done during the registration process:

1. Validate the domain hierarchy's compatibility with the RDM.
2. A new RDM property, `domain_name`, is added and assigned in the domain.
3. The domain name is registered with RDM in a domain table. An integer domain ID is created, which is used internally in the RDM.
4. The ITT tables in the RDM are updated with the domain's positions.

Unregister a Domain From RDM

The `rdmMgr` utility un-registers a domain at a time. The following is done when a domain is unregistered.

- The domain name is removed from the domain table in the RDM.
- All RDM properties created during registration are removed from the departing domain.
- All entries in the ITT tables that belong to the departing domain are removed.
- The informal positions of the departing domain are formalized to keep the measure data.
- To preserve the departing domain data, it is recommended that shared measure data belonging to the domain be transferred before unregistering the domain. It can be transferred to the domain via `transferFactdata` utility.

List Properties or Domains

The output of `rdmMgr -listDomains` is as follows. There are two header lines. Two blank lines indicate the end of the output.

```
Domain Name  Path
-----
RPAS_UT      /vol.nas/rpas_se/RPAS_UT
<two blank lines>
rdmMgr ran successfully
```

Output of `rdmMgr -listProperties` is as followed. There are two header lines. Two blank lines indicate the end of output.


```
Property Name  Value
-----
rpas_data_mart:  rpas1
rpas_wkbk_user:  whatever
...
<two blank lines>
rdmMgr ran successfully
```

Pre-Installation – Preparing the Integration Environment

[Figure 17-17](#) and [Figure 17-18](#) show the pre-installation process flow and domain preparation. Since the HSA integration is a pluggable feature of RPAS, the existing


processes for building and patching individual domains must not be changed. However, in order to integrate the domains, they must be upgraded and patched to RPAS 14.1 or a later version.

Figure 17–17 Prepare RDM Repository



The RDM Admin Schema must be granted privileges using the provided script. All other schemas must be granted the corresponding roles. At this point, the roles have the privilege to create private synonyms only. They are granted other required privileges during the post-installation process.

Figure 17–18 Prepare the Domains




Installation – Building the RDM Schemas

The RDM is built and populated from integration configuration and partition information by the `rpasInstall` utility. The schema users that have been created in the pre-installation step are populated by the schema objects and metadata.


Figure 17–19 shows the process of building the RDM schemas. The RDM manager can directly build the schemas in the database, or an alternative process can be used if so preferred.

- All schema objects along with the PL/SQL packages are created in the data mart schema.
- Private synonyms pointing to the data mart schema objects are created in each user schema.
- All required privileges are granted to the roles.

Figure 17–19 Building RDM Schemas



An alternative process for building the RDM schema is showed below. Two DDL/DML SQL script can be generated first, one for the data mart schema and one for other schema users. After optimization by the DBA or database performance expert, it can then be used to populate the schema.

Figure 17–20 Building RDM Schemas - Alternative Process

Post-Installation – Initialization and Registration

After the HSA system is built, the dimension data must be loaded, the participating domains must be registered, and the shared fact data must be loaded.

Load Dimension Data

Note that the dimension data input files must be copied to the input/dimdata directory under the RDM repository prior to running loadDimData. This can be done manually or through a wrapper script.

Register Participating Domains

A domain can be registered any time after the RDM has been built, as long as its configuration has been included in the integration configuration. This allows for the incremental integration of multiple domains.


If a domain has not been included in the integration configuration, it will require a patching process to bring the domain into the HSA system. The patching process is discussed in [Patching RDM and Integrating One or More New Domains](#).

Only a valid domain can be registered. The validity of a domain is checked by its conformity with the RDM at the partition, position, and ITT level. These checks are performed by rdmMgr during registering. The checkRDM utility also has the ability to verify a domain's validity.

Load Shared Fact Data

There are two ways to populate the shared data in RDM. The loadFactData utility can load fresh data from flat input files. The transferFactData utility can transfer existing data in the domains to RDM selectively according to user input.

The flow chart in [Figure 17–21](#) shows all the post-installation steps in the correct order.

Figure 17-21 Post-Installation Steps

Operating and Maintaining the RPAS Data Mart

This section describes the utilities used to maintain the RPAS Data Mart.

Note: Note that the RDM Manager can perform maintenance functions in addition to building the RDM. Its usage is described in [RPAS Installation Utilities](#).

- upgradeRDM is used to upgrade the version of the RPAS Data Mart
- patchRDM is used to modify the content of the RPAS Data Mart
- loadDimData is used to update the hierarchy data contained in the RPAS Data Mart
- loadFactData is used to update the measure data contained in the RPAS Data Mart
- transferFactData is used to move data between the RPAS Data Mart and a domain

- checkRDM is used to verify the integrity of the RPAS Data Mart

upgradeRDM

The upgrading of RDM is a batch side operation like domain upgrading, where the shared RDM schema or repository of the existing RDM is upgraded by either modifying, adding, or removing the existing RDM that was originally built through the RpasInstall binary using the integration configuration XML file created through Configuration Tools.

The upgrading of RDM is handled by a dedicated binary called upgradeRdm, which works in a similar way to upgradeDomain, which is used to upgrade the domain.

Usage

```
upgradeRdm -rdm rdmPath -upgrade [Options]
upgradeRdm -rdm rdmPath -dryRun [Options]
upgradeRdm -rdm rdmPath -status
upgradeRdm -rdm rdmPath -listAllModules
upgradeRdm -rdm rdmPath -listUpgradeModules
```

Commands

Table 17–35 *upgradeRDM Commands*

Command	Description
-upgrade	Upgrades the RDM schema and repository to the current version of RPAS.
-dryRun	Reports the upgrades that are applied to the RDM Schema and Repository after this utility is run with the -upgrade command.
-status	Displays what actions the RDM repository or databases need in order to be updated to the current version of RPAS utilities and libraries.
-listAllModules	Lists all the upgrade modules contained in the upgrade library.
-listUpgradeModules	Lists the upgrade modules contained in the upgrade library that can be executed during the -upgrade command.

Options

Table 17–36 *upgradeRDM Options*

Option	Description
-verbose	Shows details about each upgrade change to be applied.

Notes

Note the following:

Before running the upgradeRdm utility, back up the existing RDM databases. The upgrade process is not reversible.

If only the RPAS_HOME version has changed and the RDM schema version has not changed, then running the upgradeRdm will only update the following two tables:

- `rp_g_rdm_history_md` (RDM upgrade history table)
- `rp_g_upgrade_module_md`

A domain running one version of RPAS can connect to an RDM running another version of RPAS. However, each works only with the appropriate version of RPAS. For example, a domain built on or upgraded to RPAS 14.1.1 can integrate with a RDM on version 14.1.2. However, only the RDM utilities from RPAS version 14.1.2 can work with the RDM and only domain utilities from RPAS version 14.1.1 can work with the domain.

In an integrated environment with multiple domains and a single `RPAS_HOME`, it is required to upgrade all the domains at the same time.

The history table contains the RDM schema version and the RPAS version.

EVENT_ID	SCHEMA_VERSION	RPAS_VERSION	LAST_DEV_TAG	UTILITY	ACTION	START_TIME	COMPLETE_TIME
0	1	14.1.1	2014052101	rdmMgr	Create	20-JUN-14 10:52:29 AM	20-JUN-14 10:52:29 AM

If the RDM schema version is 1 and built with an RPAS version of 14.1.1, then if you run the `upgradeRdm` binary built with RPAS version 14.1.2 on the above RDM, then it will only update the RDM history table to reflect the RPAS version 14.1.2. Once the history table has been updated, only the version 14.1.2 RDM utilities can be run on the RDM.

patchRDM

The patching of RDM is a batch-side operation like domain patching, in which the shared RDM hierarchy, fact, and integration data are patched by either modifying, adding, or removing from the existing RDM that was originally built by the `RpasInstall` binary using the integration configuration XML file created through Configuration Tools.

The following operations on the integration configuration are supported by the `patchRDM` utility:

- Add New Hierarchy
- Modify Hierarchy Label
- Modify Dimension Label
- Add new Fact
- Remove Fact
- Modify Fact Label
- Modify Fact Group of a Fact
- Modify Fact Description
- Modify Fact Purge Age
- Reassign Fact in Integration Map
- Modify Domain Identifier in Integration Map
- Modify Measure Name in Integration Map
- Add Fact Mapping to Integration Map

■ Delete Fact Mapping from Integration Map

The RDM patching is handled by a dedicated binary called `patchRDM`, which works in a similar way to `RpasInstall`, which is used to patch the domain.

Usage

```
patchRDM -rdm <RDM Path> -config <Path to the modified integration config XML file>
```

Commands

Table 17–37 *patchRDM Commands*

Command	Description
-config	Path to modified integration config XML file.

Options

Table 17–38 *patchRDM Options*

Option	Description
-log <location for the default log file>	Directory location where the log file needs to be stored. The default name of the log file is <code>patchRdmLog.<timestamp></code> . When the -log option is not used, a default log file is generated under the current working directory.
-dryRun <location for the default report file>	Generates an easy-to-read formatted report showing the differences between the old and the new integration configurations. The default name of the report is <code>configDiffReport.<timestamp></code> .

The -dryRun Option for the patchRDM Binary

This option does not trigger any patching actions on RDM. It only generates a report about the differences between the old and new integration configurations. This option helps the user to run a check before patching to confirm the changes that are to be made to the integration configuration. The output is a formatted text file named '`configDiffReport.<timestamp>`'.

Sample configDiffReport.<timestamp>

```
Patch Hierarchy Label
  Hierarchy: clnd      New Label: Calendar_patch      Current Label: Calendar
  Hierarchy: prod      New Label: Product_patch       Current Label: Product
Add New Hierarchy
Hierarchy: clnd_new    Label: Calendar_new           Order: 999
  Add new dimensions to new Hierarchy: Calendar_new:
  Dimension: day_new   Label: Day_new
  Dimension: mnth_new  Label: Month_new
  Dimension: qrtr_new  Label: Quarter_new
  Dimension: ssn_new   Label: Half_new
  Dimension: week_new  Label: Week_new
  Dimension: woy_new   Label: Week of Year_new
  Dimension: year_new  Label: Year_new
Remove Fact
  Fact: utdysrc01u
PatchFactDescription
  Fact Name: buwpslsclru
Current Fact Description: BUWpSlsClrU Description  New Fact Description  :
```


loadDimData

The loadDimData utility is used to load hierarchy information into the RPAS Data Mart. It is similar to the loadhier utility used to load hierarchy information into an RPAS domain.

Flat File Format

The utility only supports input data files in a CSV (comma separated values) format with a header line with an added extension of .dat. When multiple input files per dimension are to be loaded, all files must contain the exact same header line.

Here is a sample data file format (with header line) for a product dimension with two levels of sku and styl:

```
sku,sku_label,styl,styl_label
0000060404,a16: black: m_none,0000023199,dk shr satin ess ton
0000060405,a17: red: m_none,0000023188,nude toeless ct spri
```

All position names are converted to lower case before loading into the RDM.

Filtered Shared Levels

The loadDimData staging process loads only shared levels from the input file. However, if the lower levels in the input file are not shared, multiple duplicate records may be loaded in the staging tables.

Validation of Input Data

Before the staged data is merged to the target dimension tables, the following validations are performed:

Note: If any of the validation fails, the user is instructed to review the error logging table and then fix and reload the input data if necessary.

1. No empty position names in the input file: This is automatically enforced by NOT NULL constraints on all name columns of the staging tables. Records with this violation are rejected and logged in the .bad file. This validation occurs at staging time.
2. Completeness of the input data: The input file must contain all shared levels of the dimension. Records with missing fields are not loaded.
3. Uniqueness of records: If multiple records have the same base position name, only one record is loaded in the database. Other duplicate records are ignored with a warning message. Users should check the error table for ignored records.
4. No space in position names: Records with spaces in the position names are discarded and logged.
5. Rollup relationship consistency: All staged records are checked for rollup relationship consistency. This check identifies the records that have multi-parent conflicts. A position rolls up to a parent position in one record, but rolls up to a different parent position in another record. When such conflict is found, the process stops and no records are loaded into the database. The following example illustrates such a conflict:

```
Sku1, sku1_label, styl1, styl1_label, clss1, clss1_label
Sku2, sku2_label, styl1, styl1_label, clss2, clss2_label
```

6. Conflict in position labels: Records with this conflict are rejected silently. The following example illustrates such conflict:

```
Sku1, sku1_label, styl1, styl1_label, clss1, clss1_label
Sku2, sku2_label, styl1, another_label, clss1, clss1_label
```

7. Rollup relationship conflict between input and existing data: If `-forceInputRollups` is specified, the utility overwrites the existing rollups with the data from the input file; otherwise, the input is rejected.
8. For calendar dimension, the dates (or hours) defined in the input file must be consecutive, in general. However, the utility does not check for gaps and the user has the flexibility to leave gaps in the input data.

Usage

```
loadDimData {ARGUMENTS}{COMMANDS}[INPUT][OPTIONS]
```

```
loadDimData -rdm <rdm_path> {-load <dim_name> | -loadAll} [-file <file_path> | -dir
<dir_path>] [-retain] [-cleanStage] [-lang <language>] [- forceInputRollups]
[-purgeAge <purgeAge>]
```

```
loadDimData -rdm <rdm_path> {-load <dim_name> | -loadAll} [-file <file_path> | -dir
<dir_path>] [-retain] -stageOnly [-cleanStage]
```

```
loadDimData -rdm <rdm_path> {-load <dim_name> | -loadAll} [-retain]
-applyStagedData [-lang <language>] [- forceInputRollups] [-purgeAge <purgeAge>]
```

Arguments

Table 17–39 *loadDimData* Argument

Argument	Description
-rdm <rdm_path>	Path to the RDM repository where input files are placed in the input/dimdata subdirectory and has database logon credentials.

Commands

Table 17–40 *loadDimData* Commands

Command	Description
-load <dimName>	Loads and refreshes the specified dimension.
-loadAll	Loads all dimensions with .csv.dat files present INPUT: (Optional) If it is not specified, the input file is expected to be in the default directory \$RDM_HOME/input/dimdata.
-dir <inputDirectory>	Processes all applicable files under inputDirectory. Input files are expected to have format: [hierName].csv.dat[.extension].
-file inputFile	Processes an inputFile only.

Options

These are optional.

Table 17–41 *loadDimData Options*

Option	Description
-retain	Does not move input files to the processed subdirectory.
-lang <language>	Language for the labels; default is English.
-stageOnly	Appends the data in the staging table only. The staged data can be validated and merged to the target tables later. This option cannot be combined with any other option, except retain and cleanStage.
-cleanStage	Purges the staged data before adding new data to the staging tables.
-applyStagedLoad	Validates and then merges the staged data from the staging tables to their respective target dimension tables. No input file is needed.
-foreInputRollups	Enforces new hierarchy roll-up changes so that they dominate existing hierarchy roll-ups in case they conflict with the roll-ups specified in the input file.
-purgeAge <purgeAge>	Specifies the purgeage during loadDimData. If the purgeage specified is 0, it purges existing positions that are not in the input. If not specified, loadDimData obtains purgeage from dimension information.

loadFactData

The loadFactData utility is used to load measure data into the RPAS Data Mart. Once present in the RPAS Data Mart, this data is available to any domain that shares that measure through the RDM.

Input File Format

The flat file must have a header line that defines the columns in the rest of the file. The header line contains dimension level names and fact names. The combination of the level names defines the base intersection of all the facts specified in the header line. Only CSV files are supported.

Sample Input File

```
sku, str, week, fact_1, fact_2, fact_3
sku_1, str_1, w01_2010, 55, Some string value, true
```

The first line in the above sample is the header line that defines three facts (fact_1, fact_2, and fact_3) and their base intersection: sku_strweek. The types of the three facts are numeric, string, and Boolean, respectively.

Input Files Location

All input files for fact data must be located in the input/factdata directory under the RDM root directory, which is specified by a command line parameter -rdm of the loadFactData utility.

At the end of the successful loading, the utility moves the input files into the processed directory.

Input File Naming

No restrictions apply to the naming of the input files, except for the file name extension: All file names must have extension .ovr or .rpl, which defines the loading

behavior. The .ovr extension means overlaying the existing data with the input data; the .rpl extension means replacing the existing data with the input for the facts being loaded.

Input Data Validation

The loadFactData utility first validates the header line:

1. The dimension level names and fact names are validated against the RDM.
2. The utility also validates that all the facts have the same base intersection defined by the combination of the dimension level names. The utility ignores the facts that cannot be found in the RDM.

All position names are validated against the corresponding dimension tables in the RDM. Data records containing invalid position names are ignored and logged. Records that cannot be parsed to the specification of the header line and fact data types are also ignored and logged.

The loadFactData catches the duplicate keys (combination of the position names) in each input file. When the input file contains duplicate keys, only one record for each duplicated key is loaded. The duplicate keys are returned to the user so that the user can verify if the correct data was loaded. If not, the user can prepare a new input file with the correct data and re-load it.

Data Loading Modes

The loadFactData utility can run in two different modes:

- **Batch Mode:** In this mode, the utility is used to load large volume of data, and it needs to run within the batch window. In batch mode, the process usually runs in multiple database transactions, which means data can be partially loaded in the case of a failure.
- **Trickle Feed Mode:** In this mode, the utility is used to load small amount of data. It can run during intra-day, and each load is a single database transaction.

In both modes, when there is a failure, the cause of the failure must be resolved and the same loading process must be re-run to ensure that all desired input is loaded into the database.

Logging

The loadFactData utility supports the same logging as all other RPAS utilities. Specifically, the utility logs the following messages:

Warning

- Skipped input files with file name extensions other than .ovr or .rpl.
- Skipped records that could not be parsed.
- Duplicate keys found.
- Message that indicates bad records are logged in the RDM.

Additional Logging in RDM

Records with invalid position names are logged in a bad records table in the user's schema. The table is named records_ignored_XXXXXX, where XXXXXX is a six-digit sequence number. The user is expected to review the records in the table and drop the table when it is no longer needed.

Usage

`loadFactData -rdm <path to rdm directory> [-facts <list of facts to load>] [-file <list of input file>] [-tf] [-maxProcesses max] [-retain]`

Table 17–42 *loadFactData*

Option	Description
<code>-rdm <path to rdm directory></code>	Required. Must be a valid path.
<code>-facts <list of facts to load></code>	Optional. When present, only specified facts will be loaded from the input files; when not present, all facts in the input files will be loaded.
<code>-file <list of input files></code>	Optional. When present, only specified files will be loaded; when not present, all files in the input/factdata ddirectory are loaded.
<code>-tf</code>	Optional. When specified, the utility runs in Trickle Feed mode. The default mode of the utility is batch.
<code>-maxProcesses max</code>	When specified, the utility creates at most max processes to run some step in parallel. The <code>-tf</code> option does not support parallel processing.
<code>-retain</code>	Optional. When specified, the utility does not clean input files and metadata used after load.

transferFactData

This utility is used to transfer data between the RDM and a domain. To transfer data from multiple domains to the RDM, this utility must be executed multiple times, based on the number of domains involved. The `transferFactData` utility only transfers the shared measure data that is specified by the user in the command. Data can be transferred in either direction: domain to RDM or RDM to domain. The primary uses for `transferFactData` are in the creation of the RPAS Data Mart and as a part of a batch process within a domain.

Usage

`transferFactData -d domainPath {command} {parameter} [options]`

Command

One of the following two commands must be present to transfer data from the domain to RDM or RDM to the domain.

- `-toRDM`
- `-fromRDM`

After the command, a list of shared measures that must be transferred must be explicitly specified. Note that data transfer from the domain to RDM clears the shared measure data in the domain after the transfer. The data transferred from RDM to the domain is copied into the shared measure array in the domain.

Parameters

Table 17–43 Parameters

Parameter	Description
-measureList measures	The measureList parameter is a list of measure names separated with a comma and no additional spaces unless the entire list is quoted.
-measureListFile path	The measureListFile parameter is a file path to a text file containing any number of lines of comma separated measure names.
-allMeasures	Transfer all shared measures contained in the domain.

Options

Table 17–44 transferFactData Options

Option	Description
-replicate	This option retains the domain copy of the measure data after the transfer from the domain to RDM. By default, transfer from the domain to RDM clears the shared measure data from the domain.
-processes N	Number of processes to run in parallel. The default is 1 (single process). If specified, transferFactData is executed in parallel on each local domain. Non-HBI measures can take advantage of this parallelization.

Example:

```
transferFactData -d <domainPath> -toRDM -measureList measure1,measure2 -processes
4
transferFactData -d <domainPath> -toRDM -measureListFile saleMeasures.txt
-replicate
transferFactData -d <domainPath> -fromRDM -allMeasures -processes 4
```

checkRdm

HSA involves complex inter-communication between the RPAS domains and the RPAS data mart called RDM. Both the domain side arrays and RDM side tables are affected as part of the routine operations performed by other domain and RDM utilities.

Users may encounter issues while processing the domain or RDM and need to obtain diagnostic information to understand and fix the issues. The diagnostic information can help the user to detect any asynchronous data between the domain and the RDM.

The checkRdm binary can be used to obtain diagnostic information from both the RDM and the domain, either independently or when they are integrated.

Usage

```
checkRdm {-rdm rdmPath | -d domPath} {commands} [options]
```

Arguments

Table 17–45 *checkRdm Arguments*

Argument	Description
-d	Path to a domain.
-rdm	Path to RDM repository.
-dim	Shared dimension name.
-hier	Hierarchy name.
-name	Domain integration name.

Commands

Table 17–46 *checkRdm Commands*

Command	Description
-checkParents	Checks the rollups of all the positions.
-listNonsharedDims	Lists all non-shared dimensions of a shared hierarchy.
-listSharedHiers	Lists all sharing hierarchies between the domain and RDM.
-listSharedMeasures	Lists measures of the domain that are sharing with RDM.
-preRegValidation	Performs pre-registration validation to ensure the domain can be registered to RDM. Integration name is optional; if present, it ensures that the name can be used for registration.
-printHier	Prints the dimension tree of a hierarchy (domain or RDM).
-validateHier	Runs validation checks on the specified hierarchy to make sure the domain is integrated correctly. Dimension name is optional; if present, only the specified dimension is validated.
-validatePartitioning	Validates the partitioning between RDM and domain.
-validateRdm	Validates the sanity of the RDM.

Options

Table 17–47 *checkRdm Options*

Option	Description
-all	Checks all the dim arrays/tables from the specified dim to the leaf dim. Can only be used with the -dim option.
-listDims	Optionally lists dimensions on both the domain and RDM.

Examples

[Table 17–48](#) provides examples of the checkRdm utility on the mfprtl domain.

Table 17–48 *checkRdm Examples*

Command	Result/Output
checkrdm -d \$dompath -listNonSharedDims -hier prod	Non-shared dimensions in hierarchy 'prod': pl1, pl2
checkrdm -d \$dompath -listSharedHiers	Shared Hierarchies: clnd, prod, loc
checkrdm -d \$dompath -listSharedHiers -listDims	Dimensions of all shared hierarchies Hier Dim Domain RDM (only PROD is shown) prod skup X X skug X X scls X X clss X X dept X X pgrp X X dvsn X X cmpp X X pl1 X pl2 X
checkrdm -d \$dompath -listSharedMeasures	Measure Fact Name Fact Group ----- buwpmiscinr buwpmiscinr chnlsclsweek buwpmiscoutr buwpmiscoutr chnlsclsweek utdylfd03u utdylfd03u utlfd1
checkrdm -rdm \$rdmpath -d \$dompath -preRegValidation (domain not integrated)	PreRegValidation successful
checkrdm -rdm \$rdmpath -d \$dompath -preRegValidation (domain integrated)	checkrdm: Domain is expected to be a non-integrated domain.
checkrdm -rdm \$rdmpath -printHier -hier prod	Dimension graph of hierarchy 'prod' in RDM: skup --> skug --> scls --> clss --> dept --> pgrp --> dvsn --> cmpp
checkrdm -d \$dompath -printHier -hier prod	Dimension graph of hierarchy 'prod' in domain: skup --> skug --> scls --> clss --> dept --> pgrp --> dvsn --> cmpp \-> pl1 \-> pl2 Dimension graph of hierarchy 'prod' in RDM: skup --> skug --> scls --> clss --> dept --> pgrp --> dvsn --> cmpp
checkrdm -d \$dompath -validateHier -hier prod	ValidateHier successful

Table 17–48 (Cont.) checkRdm Examples

Command	Result/Output
checkrdm -d \$dompath -validateHier -hier prod -dim clss	ValidateHier successful
checkrdm -d \$dompath -validatePartitioning	Repartition is not needed. Check status flag ... VALID
checkrdm -rdm \$RDMHOME -validateRdm	Connected to data_mart_<RDM_USER> Connected to patch_<RDM_USER> Connected to batch_<RDM_USER> Connected to dimload_<RDM_USER> Connected to factload_<RDM_USER> Connected to hiermgr_<RDM_USER> Connected to wkbk_<RDM_USER> Connected to etl_<RDM_USER> The RDM sanity test passed! checkrdm completed successfully

RDM Partitioning and RDM Partition Manager

RPAS global domains are typically divided into local domains at a certain partition level according to some business requirements such that different business units can operate on their local domains separately. RDM partitioning, however, is solely for performance purposes. The data inside RDM is shared by multiple integrated domains, and multiple RPAS processes need to access the shared data simultaneously. Proper partitioning of RDM is vital to the overall performance of an HSA system. If performance is not an issue or the amount of data in the RDM is very small, RDM partitioning can be optional.


Determining the RDM Partition Level

In order to support partitioning in RDM, all integrated domains are required to be partitioned on the same branch of a hierarchy (for example, the product hierarchy).

When all domains are partitioned at the same level, the RDM should be partitioned at the same level as well.

When the domains are partitioned at different levels, the RDM should be partitioned at the highest common level that is lower than or equal to the lowest domain partition level.

In the example shown in [Figure 17–22](#), domain 1 is partitioned at PGRP and domain 2 at DEPT. But since domain 1 does not have DEPT level, the RDM partition level is assigned to SCLS, which is the highest common level below all domain partition levels.

Figure 17–22 Determining the RDM Partition Level**Example: Determining RDM Partition Level**

Note that the RDM partition level determined by the above method is the recommended level that keeps the number of positions in each partition to a minimum. The user can choose another level *lower* than that as long as it is shared by all domains and RDM for other reasons, such as anticipating a future domain partitioned on that level.

Mapping Local Domains to RDM Partitions


Many, if not most, domain operations, such as workbook build and commit, are conducted inside the individual local domain. If shared data from RDM is needed for those operations, that data will be transferred to the workbook or local domain. When the data is committed, it then will be transferred back to RDM.

Reading data from RDM is relatively fast. However, writing data back to RDM can be a performance bottleneck since the database needs to maintain transaction logs to keep track of all the changes in case a rollback is needed.

A technique called partition exchange can be used to speed up this database write-back operation. A temporary data set can be prepared with updated partitioned data, and the partitions can then be swapped very quickly as a whole instead of going through each individual record.

To support partition exchange, local domains must be mapped to the RDM partition in a one-to-many relationship. In other words, one local domain can be mapped to one or multiple RDM partitions, but two local domains cannot be mapped to the same RDM partition. Ideally, each local domain is mapped to only one RDM partition, but this is sometimes impossible when multiple integrated global domains have different partition levels or partition configurations. As long as one RDM partition is mapped to only one local domain and not shared by different local domains, the partitioning benefits can be achieved.

Figure 17–23 shows the relationship between local domains and RDM partitions.

Figure 17-23 RDM and Domain Partition Mapping


The partition mapping is configured initially by the RDM Partition Manager and maintained by the same utility through the repartitioning process, when the mapping is broken.

RDM Schema Partitioning

A few tables in the RDM contain information to support the partitioning. The concept of partition ID is introduced to help organize the partitions. The ID value starts at 0 and is incremented by 1 sequentially. The partition level dimension table has a PARTITION_ID column and is partitioned by a list of partition ID value(s). Only one value per list is supported. In other words, each partition ID is mapped to one partition. In fact, the required name of the partition is constructed by using this pattern: "P" + {Partition ID}, such as "P0", "P1", and so on. Partition P0 is designated as the default partition (see [The Format of RDM Partition Configuration](#)).

The rest of the child dimension tables below the partition level and the fact tables that have a foreign key relationship to any of the partitioned dimension tables are all partitioned by reference.

The partition mapping metadata table RP_G_PARTITION_MD contains the mappings from partition level positions to the partition ID. The partition level ITT (index translation table) table contains additional information about local domain mapping that is useful in the repartitioning operation.

Figure 17–24 RP_G_PARTITION_MD

RDM Partition Manager

A utility called RDM Partition Manager is used to manage partition configuration and carry out the repartitioning operation for the RDM. The use of this utility is described below.

This utility is used to manage Oracle database partitions for the RPAS Data Mart (RDM). It can create, update, and export partition configurations and carry out the repartitioning operation for RDM.

Usage

```
rdmPartitionMgr -createConfig -d domPath -name domIntName -partitionDim
dimName -output output.xml
```

```
rdmPartitionMgr -updateConfig -d domPath -name domIntName -input input.xml
-output output.xml
```

```
rdmPartitionMgr -exportConfig -rdm rdmPath -output output.xml
```

```
rdmPartitionMgr -repartition -rdm rdmPath [-input input.xml]
```

```
rdmPartitionMgr -report -rdm rdmPath
```

Commands

Table 17–49 RDM Partition Manager Commands

Command	Description
-createConfig	Create an initial partition configuration based on the first domain. The XML file is marked to use full update mode in subsequent updates.
-updateConfig	Update an existing configuration with information from an additional domain.
-exportConfig	Export the current RDM partition configuration to an XML file. The XML file is marked to be updated incrementally.
-repartition	Repartition the RDM. If an input file is provided, it will repartition RDM to match the partition configuration in the file. Otherwise, it will repartition RDM based on internal dimension tracking information.

Table 17–49 (Cont.) RDM Partition Manager Commands

Command	Description
-report	Print out a report on the current RDM partition configuration.

Parameters

Table 17–50 RDM Partition Manager Parameters

Parameter	Description
-d	Path to the domain.
-rdm	Path to the RDM repository.
-partitionDim	RDM partition level.
-input	Input XML file path.
-output	Output XML file path.
-name	Integration name of the domain.

Note the following:

Full update mode regroups positions based on the domain partitioning of all domains. It generates a smaller number of partitions but may shuffle positions into different groups for every update. It removes all empty partitions from the list except the default partition. It also removes positions that do not belong to any domains. Full update mode is best used for initial partitioning.

Incremental update mode maintains existing position locations as long as it is possible. It does not remove empty partitions or positions not used in any current domains. Incremental update mode can reduce the time for repartitioning and is best used for adding a new domain to an existing integration environment.

The Format of RDM Partition Configuration

The partition configuration file is generally referred to as partition info XML file since it is in XML format. It contains the following four sections:

- <partitiondim>: the partition level
- <update-mode>: full or incremental
- <domains>: record keeping for all the domains that have been analyzed by the RDM Partition Manager. It is used internally only by the RDM Partition Manager.
- <rdm-partitions>: definitions of partitions in RDM. Each partition has a partition ID and a list of assigned positions. Note that there is an empty partition with ID 0. It is the default partition that should not contain any positions initially and is used during RDM hierarchy operations later on to temporarily hold positions for which a proper partition has not yet been determined.

Sample RDM Partition Configuration partitionInfo.xml

```
<?xml version="1.0" ?>
<rpas>
  <partitiondim>scls</partitiondim>
  <update-mode>full</update-mode>
  <domains>
    <domain name="mfprt1">
      <subdomain name="ldom01">30000, 30001, 30002, 30003, 30004, 30005,
```

```

30006, 30007, 30008, 30009, 30048, 30049, 30050, 30051, 30052, 30053, 30054,
30055, 30056, 30057, 30058, 30059, 30082</subdomain>
    <subdomain name="ldom02">30010, 30011, 30012, 30013, 30014, 30015,
30016, 30017, 30060, 30061, 30062, 30063, 30064, 30065</subdomain>
    <subdomain name="ldom03">30018, 30019, 30020, 30021, 30022, 30033,
30034, 30035, 30036, 30037, 30038, 30066, 30067, 30068, 30069, 30070, 30071,
30072, 30073, 30074, 30075, 30076</subdomain>
    <subdomain name="ldom04">30039, 30040, 30041, 30042, 30043, 30044,
30045, 30046, 30047, 30077, 30078, 30079, 30080, 30081</subdomain>
  </domain>
  <domain name="mfprtl_ut">
    <subdomain name="ldom01">30003, 30004, 30005, 30006, 30007, 30008,
30009, 30013, 30014, 30015, 30016, 30017, 30056, 30057, 30058, 30059</subdomain>
    <subdomain name="ldom02">30010, 30011, 30012, 30060, 30061, 30062,
30063, 30064, 30065</subdomain>
    <subdomain name="ldom03">30018, 30019, 30020, 30021, 30022, 30033,
30034, 30035, 30036, 30037, 30038, 30071, 30072, 30073, 30074, 30075,
30076</subdomain>
    <subdomain name="ldom04">30039, 30040, 30041, 30042, 30043, 30044,
30045, 30046, 30047</subdomain>
    <subdomain name="ldom05">30000, 30001, 30002, 30048, 30049, 30050,
30051, 30052, 30053, 30054, 30055, 30082</subdomain>
  </domain>
</domains>
<rdm-partitions>
  <partition id="0"></partition>
  <partition id="1">30003, 30004, 30005, 30006, 30007, 30008, 30009, 30056,
30057, 30058, 30059</partition>
  <partition id="2">30000, 30001, 30002, 30048, 30049, 30050, 30051, 30052,
30053, 30054, 30055, 30082</partition>
  <partition id="3">30013, 30014, 30015, 30016, 30017</partition>
  <partition id="4">30010, 30011, 30012, 30060, 30061, 30062, 30063, 30064,
30065</partition>
  <partition id="5">30066, 30067, 30068, 30069, 30070</partition>
  <partition id="6">30018, 30019, 30020, 30021, 30022, 30033, 30034, 30035,
30036, 30037, 30038, 30071, 30072, 30073, 30074, 30075, 30076</partition>
  <partition id="7">30077, 30078, 30079, 30080, 30081</partition>
  <partition id="8">30039, 30040, 30041, 30042, 30043, 30044, 30045, 30046,
30047</partition>
</rdm-partitions>
Initial RDM Partitioning
</rpas>

```

Initial RDM Partitioning


If it is determined that RDM partitioning is not needed, the RPAS installation process (rpasInstall) can take the -noPartition option and no other information is needed. Otherwise, an initial RDM partition info XML file is required in order to build an RDM with partitioning. The partition configuration for a domain is usually created manually. For RDM, this file should be generated by the RDM Partition Manager. It is difficult to create the RDM partition configuration manually since it is required to maintain the one-to-many relationship from all local domains in all integrated domains to the RDM partitions.

All domains that are to be integrated must be built or upgraded first. The RDM Partition Manager then goes through all the domains one by one, gathering and analyzing local domain information and calculating the optimal RDM partition configuration conformant to all domains.

Figure 17–25 shows the process of creating the initial partition configuration for RDM.

Figure 17–25 Creating Initial Partitioning Configuration for RDM

Process: Creating Initial Partitioning Configuration for RDM




Patching RDM and Integrating One or More New Domains

When one or more domains are added to an existing HSA system with RDM partitioning, it is generally required to repartition the RDM so that its partitions are compatible with the local domains of the new domains. RDM Partition Manager should export existing partition configuration from the RDM and then update it by going through each of the new domains. After the new partition configuration is finalized, the RDM Partition Manager can repartition the RDM accordingly.

Figure 17–26 shows the repartitioning process for adding new domains. This must be done before registering any of the new domains.

Figure 17–26 RDM Repartition for Integrating New Domains

Process: RDM repartition for integrating new domains



Hierarchy Operations Affecting RDM Partitioning

Any hierarchy operations on the RDM partition level in the RDM or any of the integrated domains may alter the one-to-many mapping of local domains to RDM partitions. When that happens, repartition is needed to make the mapping conformant.

The operations affecting RDM partitioning include: adding new positions on the RDM partition level and moving positions from one local domain to another (reclass). Deleting or renaming positions on the RDM partition level only causes them to be removed from the partition mapping table and, as a result, no repartition is needed.

Dimension Data Loading in RDM and Domain

When loadDimData loads new dimension data in the RDM that brings in new partition level positions, it should load them into appropriate partitions. These are determined by analyzing the local domain mappings in the ITT tables. However, if loadHier has not been run in any of the domains and, as a result, those new positions

do not yet have corresponding entries in the ITT table, loadDimData cannot determine which partitions to put the new partition positions. In that case, it will put them into the default partition P0.

When loadHier loads new dimension data in the domain that affects the RDM partition (new partition level positions or reclass), it should set the Repartition recommended status flag (see [Coordinating RDM Partitioning Using the RDM Status Table](#)). The user needs to manually run the repartitioning operation through the RDM Partition Manager to move the corresponding data in the RDM to appropriate partitions.

Domain Reclassification

Domain reclassification, as a result of following utilities, should be followed by RDM repartitioning.

- loadHier
- reconfigGlobalDomainPartitions

These utilities set the Repartition recommended flag after a reclass operation. The repartitioning process for RDM must be started manually.

Dynamic Position Maintenance

Dynamic Position Maintenance (DPM) operations, including both online DPM and batch DPM (informalPositionMgr), are carried out from local domains and no reclass is involved. DPM can determine which RDM partitions to put into the new informal positions when they are created. No repartitioning is needed.

Coordinating RDM Partitioning Using the RDM Status Table

In the RPAS HSA, hierarchy data is stored in RDM and all integrated domains. The positions in RDM for a given dimension must be a superset of the positions in any of the integrated domains. Since hierarchy data is loaded and updated by separate processes for RDM and domains, it is possible that they can become out of sync, either temporarily or accidentally.

Shared measures are stored in the RDM only. When the domain has more positions than the RDM, operations on the shared measures can cause data loss since the data slices for those extra positions cannot be saved to RDM. This is a serious problem for HSA during day-to-day operations. In order to prevent this from happening, a method must be devised to maintain the hierarchy consistency across RDM and domains.

One more reason to have this is that the database should not do repartitioning when some other process is running on the database and vice versa. This is not needed when operating on only one domain. However, in an HSA integration environment, multiple domains can exist in different time zones and on different servers, and multiple utility processes running on them can impact the RDM at the same time. A method is required to coordinate these processes when they make changes to the RDM.

The RDM Status Table can be used to help control various RPAS/RDM processes. The processes can set, unset, or check flags in the table so as to coordinate their operations.

Table 17–51 RDM Status Table: RP_G_RDM_STATUS_MD

SEQ_ID	STATUS_CODE	DOMAIN_ID	INFO	CREATED_TIME
2	2002	0	PROD	
3	1001	1	LOC	

Table 17–51 (Cont.) RDM Status Table: RP_G_RDM_STATUS_MD

SEQ_ID	STATUS_CODE	DOMAIN_ID	INFO	CREATED_TIME
4	1002	NULL	PROD	
6	1004	0	SCLS	

Table 17–52 RDM Status Message Table: RP_G_STATUS_MSG_MD

STATUS_CODE	STATUS_MSG
1001	loadHier in progress
1002	loadDimData in progress
1003	reconfigGlobalDomainPartitions in progress
1004	Repartition in progress
2002	Repartition recommended

Status: loadHier in Progress

When loadHier is in progress, repartitioning is not allowed, since the positions in the domain can still be removed, added, or moved.

RDM Status	loadHier in Progress
Set By	loadHier
Resolved By	loadHier
Restricted Utilities	rdmPartitionMgr

Status: reconfigGlobalDomainPartitions in Progress

Similar to loadHier, when reconfigGlobalDomainPartitions is in progress, repartitioning is not allowed, since the positions in the domain can still be removed, added, or moved across local domains.

RDM Status	reconfigGlobalDomainPartitions in Progress
Set By	reconfigGlobalDomainPartitions
Resolved By	reconfigGlobalDomainPartitions
Restricted Utilities	rdmPartitionMgr

Status: loadDimData in Progress

Database changes should not occur when a data load is in progress.

RDM Status	loadDimData in Progress
Set By	loadDimData
Resolved By	loadDimData
Restricted Utilities	rdmPartitionMgr Patch and Upgrade

Status: Repartition in Progress

The loadHier utility updates the ITT table with local domain mapping. This information is used to calculate the correct partitions for repartitioning. Because of this, local domain mapping changes must not be allowed while repartitioning is going on.

Any RDM utility must be restricted, since there could be errors of the table being locked during its redefinition. In addition, this can cause inconsistent results.

RDM Status	Repartition in Progress
Set By	rdmPartitionMgr
Resolved By	rdmPartitionMgr
Restricted Utilities	loadHier loadDimData renamePositions reconfigGDP Mace loadFactData transferFactData workbook operations – wbbatch and rpasWbServer Patch and Upgrade

Status: Repartition Recommended

With this flag on, the RDM is not optimally partitioned. As a result, performance can be affected. However, it does not prevent any utility from executing. This flag can be reset by rdmPartitionMgr after repartitioning or by loadHier/reconfigGDP when the related positions have been modified in the domain and repartitioning is no longer needed in RDM.

RDM Status	Repartition recommended
Set By	loadHier reconfigGDP
Resolved By	rdmPartitionMgr loShIER RECONFIGgdp
Restricted Utilities	N/A

Coordination of Domain Processes in an Integrated Environment

In an integrated environment, HSA can have multiple domains registered. Though registered with RDM, one domain may not be aware of what process is running in any other domain, since multiple domains can exist in different time zones and on different servers. Apart from RDM utilities, some domain utilities can have an impact on the RDM as well and thus an indirect impact on all other registered domains. Hence, it becomes important to sequence the domain and RDM utilities in the right manner.

HSA recommends sequencing for batch execution of domain and RDM utilities. Interdependent utilities must follow the sequence. However, there are some utilities that can be executed in parallel.


Apart from recommendations, HSA also provides a way to store current RDM status that is used by interdependent utilities to coordinate better between and among themselves. The status is stored in a RDM table, and a utility is provided to view them.

Recommended Process Sequencing

As with non-HSA RPAS, batch processes follow a sequence in which hierarchies are loaded before measure data and measure data is loaded before mace is executed. Similarly, in an HSA environment, the sequence consists of loading the hierarchy and dimension data, followed by loading measure and fact data, and eventually followed by mace. But now the sequencing is much more important because of multiple domains.

A typical execution of batch processes is illustrated in [Figure 17-27](#).

Figure 17-27 Typical Batch Process Execution



The above sequence should be the best sequence with respect to integrity and performance for most of the cases. However, some variations in the above sequence are also acceptable. For example, loadHier can complete first and then loadDimData start next in parallel to loadMeasure. In this way, rdmPartitionMgr execution is not necessary because loadDimData manages loading the positions in right partitions because loadHier has already recorded local domain mappings for each position through loadHier.

Similarly, the execution of batch processes can vary slightly from the above diagram only if they follow the matrix explained in [Parallel Execution of Processes](#). This matrix documents the utilities that can and cannot be executed together even if they are executed in different domains.

Parallel Execution of Processes

[Table 17-53](#) lists the interdependent utilities. The cross signifies that these utilities cannot be executed together, even on different domains. Note that rdmMgr is only for domain registration and un-registration here.

Table 17–53 Interdependent Utilities

	loadHi er	reconfi gGDP	Mace	Patch/ Upgra de	Renam e Positio ns	loadDi mData	loadTr anslati on	loadFa ctData	Transf erFact Data	rdmMg r*	rpasW bServ er	rdmPa rtition Mgr
loadHi er				X						X		X
reconfi gGDP				X						X		X
Mace				X								X
Patch/ Upgra de	X	X	X	X	X	X	X	X	X	X	X	X
Renam e Positio ns										X		
loadDi mData				X						X		X
loadTr anslati on				X						X		X
loadFa ctData										X		X
transfe rFactD ata				X						X		X
rdmM gr*	X	X		X	X	X	X	X	X	X	X	X
rpasW bServe r												
rdmPa rtition Mgr	X	X		X	X	X	X	X	X	X	X	X

The intersections for the utilities that do not have a cross can be executed in parallel.

A domain utility must not be re-executed on the same domain while its one instance is already in progress. However, RDM utilities may have different behavior, as explained in [Table 17–54](#).

Table 17–54 RDM Utilities

Utility	Second Instance When One is Already in Progress
loadDimData	Allowed
loadTranslation	Allowed
loadFactData	Allowed only if mutually exclusive fact groups are loaded in different calls.
transferFactData	Transfer from domain to RDM is allowed only for mutually exclusive fact groups. Transfer of RDM to domain is allowed.

Table 17–54 (Cont.) RDM Utilities

Utility	Second Instance When One is Already in Progress
rdmMgr	Not allowed. Domain registration and un-registration locks the whole domain.
rdmPartitionMgr	Not allowed

RDM Status Tables

Two tables in the RDM schema are available for RDM status.

RP_G_RDM_ST_MSG_MD is a static table with status codes and their corresponding messages. This table is populated at the installation time with a fixed number of codes and messages.

RP_G_RDM_STATUS_MD is the table that is created at the installation time but populated at the runtime by various utilities. The utilities set and clear status messages in this table. Each status is identified by a unique sequence ID. The domain utility's status has a domain ID recorded. Sometimes, additional information is provided by the 'Info' column, such as the dimension name that loadDimData is currently loading.

Figure 17–28 Status Table

DATA_MART_HARSHP.RP_G_RDM_ST_MSG_MD	
P * STATUS_CODE	NUMBER (12)
* STATUS_MSG	VARCHAR2 (256 BYTE)
RP_G_RDM_ST_MSG_MD_PK(STATUS_CODE)	

DATA_MART_HARSHP.RP_G_RDM_STATUS_MD	
P * SEQ_ID	NUMBER (12)
* STATUS_CODE	NUMBER (12)
DOMAIN_ID	NUMBER (12)
INFO	VARCHAR2 (32 BYTE)
* CREATED_TIME	TIMESTAMP
RP_G_RDM_STATUS_MD_PK(SEQ_ID)	

RDM Status Messages

HSA provides twelve RDM status messages. Out of these, ten are in progress status messages and the other two are information status messages. Based on each status message, some of the utilities may be restricted. The details are presented in [Table 17–55](#).

While in progress messages are provided to ensure that a restricted utility is not executed when an interdependent utility is already in progress, the information status messages are provided to avoid data corruption and poor performance.

Table 17–55 RDM Status Messages

RDM Status	Set By	Resolved By	Restricted Utilities	Comments
loadHier in Progress	loadHier	loadHier	rdmPartitionMgr rdmPatch rdmUpgrade	loadHier updates the ITT tables.
reconfigGDP in Progress	reconfigGDP	reconfigGDP	rdmPartitionMgr rdmPatch rdmUpgrade	reconfigGDP may update the ITT tables.
Mace in Progress	Mace (Begin Batch)	Mace (End or Abort Batch)	rdmPartitionMgr rdmPatch rdmUpgrade loadFactData	
loadDimData in Progress	loadDimData	loadDimData	rdmPartitionMgr rdmPatch rdmUpgrade Mace	
loadTranslation in Progress	loadTranslation	loadTranslation	rdmPartitionMgr rdmPatch rdmUpgrade	
Domain Registration in Progress	rdmMgr	rdmMgr	rdmPartitionMgr rdmPatch rdmUpgrade	
Domain Un-registration in Progress	rdmMgr	rdmMgr	rdmPartitionMgr rdmPatch rdmUpgrade	
loadFactData in Progress	loadFactData	loadFactData	rdmPartitionMgr rdmPatch rdmUpgrade mace	
transferFactData in progress	transferFactData	transferFactData	rdmPartitionMgr rdmPatch rdmUpgrade	

Table 17–55 (Cont.) RDM Status Messages

RDM Status	Set By	Resolved By	Restricted Utilities	Comments
Repartition in Progress	rdmPartitionMgr	rdmPartitionMgr	rdmPartitionMgr rdmPatch rdmUpgrade loadHier loadDimData renamePositions Mace loadFactData transferFactData workbook opera-tions loadTranslation	Any RDM utility or domain utility that involves database write operations must be restricted since errors may exist that the table is locked during its redefinition. Otherwise, there could be inconsistent results.
RDM is missing positions*	loadHier reconfigGDP	loadHier loadDimData reconfigGDP	rdmPartitionMgr Mace loadFactData transferFactData (for respective domain) workbook operations (for respective domain)	Domain utilities on only those domains for which positions are missing in the RDM are restricted. The status table has a domain ID for which RDM is missing positions. More details are explained below.
Repartition Recommended*	loadHier reconfigGDP	rdmPartitionMgr	NA	The RDM is not optimally partitioned. Hence, performance can be affected. But, it does not stop any utility from executing. The utilities log a warning message saying since the RDM needs re-partitioning, performance may be affected. More details are explained below.

RDM is Missing Positions

Hierarchy data is stored in RDM as well as in all integrated domains. The positions in RDM for a given dimension must be a superset of the positions in any of the integrated domains. Since hierarchy data is loaded and updated by separate processes for RDM and domains, it is possible that they can become out of sync either temporarily or accidentally.

Shared measures are stored in the RDM only. When the domain has more positions than the RDM, operations on the shared measures can cause data loss since the data slices for those extra positions cannot be saved to RDM. This is a problem for HSA during day-to-day operations. "RDM is missing positions" status message is the way to solve this problem.

Repartition Recommended

For optimal database performance, a domain's sub-domains must maintain a one-to-many mapping to RDM partitions. However, due to various operations, this mapping may be disturbed. At such times, this status message will be set. This is a warning status that RDM must be re-partitioned as soon as possible to avoid poor performance.

RDM Status Utility

A utility is provided to view and clear the RDM status messages. The utility can display:

All current RDM messages.

All static status codes and corresponding messages.

A specific status code, if it is currently active.

The utility can also clear a specific code, which can be identified by its unique sequence ID. This is not recommended unless there is an unrecoverable erroneous scenario (for example, a core dump).

Usage of `rdmStatus` Utility

```
rdmStatus -rdm rdmPath -list [-listFormat]
```

```
rdmStatus -rdm rdmPath -listCodes
```

```
rdmStatus -rdm rdmPath -checkStatus code [-listFormat]
```

```
rdmStatus -rdm rdmPath -clearStatus seq_id
```

Commands

Table 17-56 *Commands*

Command Name	Description
-list	Lists all current RDM status.
-listCodes	Lists all available RDM status codes and corresponding messages.
-checkStatus	Checks the status for the specified code. The code is displayed by the listCodes command.
-clearStatus seq_id	Clear the status specified by the sequence ID. The sequence ID is displayed by -list command.

Options

Table 17-57 *Options*

Option Name	Description
-listFormat	Changes format of output from being a table to being a list.

RPAS HSA Translation

The RDM stores shared hierarchies (dimensions and levels in database terms) and shared measures (facts in database terms) which can be accessed by database reporting tools. To support internationalization of these reports, the translations for the following items are also stored in the RDM:

- Dimension labels
- Level labels
- Fact labels and descriptions
- Position labels

The RDM supports translations for the 18 languages as listed in the Internationalization chapter. The source files are generated by the RPAS Installer in csv format for dimension and level labels, and fact labels and descriptions and put under the input directory of the RDM repository during the RDM build process. The sources file for position labels need to be provided by the customer as flat files.

A utility named loadTranslation is used to load the translations into RDM.

Flat File Requirements

This section describes flat file requirements.

Dimension Label Flat File

The file name must be dimension_labels.csv. A language extension can be optionally added when the file is sent to Translation Team, for example, the file sent to Translation Team can have name dimension_labels.csv.english.

The file must contain exactly three columns with fixed order: dimension name, language name, and label.

Level Label Flat File

The file name must be level_labels.csv. A language extension can be optionally added when the file is sent to Translation Team, for example, the file sent to Translation Team can have name level_labels.csv.english.

The file must contain exactly three columns with fixed order: level name, language name, and label.

Fact Label and Description Flat Files

The fact labels and descriptions are in separate files, with file names fact_labels.csv and fact_descriptions.csv, respectively. Again, optional file name extensions can be added when the files are sent for translation.

The files will have three columns with fixed order: fact name, language name, and labels or descriptions.

Position Label Flat Files

Position label files are not generated from RPAS Install. They are generated from external sources or prepared by the dimension loaders. Since the Translation Team returns the translated file in the same base name and format as the original file, RPAS 14.1 requires the original files for RDM to have base file name <levelName>_labels and to be in csv format, for example, the label file for sku will have base name sku_labels.

Source Language

The rdmMgr, invoked by the RPAS Installer, loads the labels and fact descriptions into RDM at the time it builds the schema from the integration configuration, where the "source" language is specified as a root level attribute. That attribute is used as the

source language attribute for all labels and descriptions in the integration configuration. It is also saved as the language property of the RDM.

loadTranslation

RDM utility loadTranslation is used to load the translations into RDM. The input to this utility are the csv files returned from the Translation Team.

Input Files

The input files should be utf-8 encoded (this is the same as in the prior versions of RPAS).

Dimension Label Translation The base name of the file must be dimension_labels. The file must contain three columns with fixed order: dimension name, language name, and dimension label. The dimension names and language names are case insensitive. The labels will be loaded exactly as provided.

Level Label Translation The base name of the file must be level_labels. The file must contain three columns with fixed order: level name, language name, level label. The level names and language names are case insensitive.

Fact Label and Description Translation The fact label and description translations are in separate files, with base name of the files being fact_labels and fact_descriptions, respectively. The files will have three columns with fixed order: fact name, language name, and label/description. Again, the first two columns are case insensitive. The labels/descriptions will be loaded exactly as provided. The labels will be loaded exactly as provided.

Position Label Translation One translation file is required for each level. The base name of the file must be <levelName>_labels, for example, the translation file for SKU label must have base name sku_labels. The position label translation files will have three columns with fixed order: position name, language name, and label translation. Again, the first two columns are case insensitive, and the translation shall be loaded exactly as provided.

loadTranslation Utility

This section describes the loadTranslation utility.

Command Line Syntax

```
loadTranslation -rdm <rdmRoot> [-sourceLang <langName>] [-file  
<fullPathOfInputFile>]
```

-rdm specifies the path to the root directory of the RDM repository.

-sourceLang specifies the source language of the translation, it defaults to the language property of the RDM if not provided.

-file specifies the input file to load. When this parameter is not present, the utility loads all input files in <rdmPath>/input/translation. After successfully loading the files, the utility moves them to <rdmPath>/input/translation/processed.

Validations Input files whose file name does not meet the requirement specified in the above section are skipped.

The utility also validates each record of the input. It should log and skip invalid records. The utility validates the first two fields of each record: the first field must be a valid name - dimension name, level name, fact name, or position name, depending on the input type. The second field must be a valid language name (one of the 18 pre-loaded into RDM). The values in the first two fields are not case sensitive.

Logging The utility supports all RPAS log levels. The invalid records are logged at warning level.

Appendix: Curve Administration Guide

This appendix provides details about the utilities used in Curve administration.

curvevalidate

The curvevalidate utility automatically executes during the domain installation, and it can also be run at any time against a master domain or one subdomain. If run against the master domain, it checks the master and all subdomains. If run against a subdomain, it checks the master domain and only the subdomain (not all other subdomains). This function verifies that:

- Profile and Source intersections and source data are properly defined
- Profile intersections respect the partition dimension

curvevalidate Usage

```
curvevalidate -d domainpath [-s]
```

[Table A-1](#) provides descriptions of the arguments used by the curvevalidate utility.

Table A-1 Arguments Used by the curvevalidate Utility

Argument	Description
-d <i>domainpath</i>	Indicates the path to the domain.
-s	Sets defaults.
-debug	Causes temporary measures to be retained for debugging purposes.
-h	Displays information about curvebatch in the terminal screen.
-version	Displays version information.
-loglevel	Sets the logger verbosity level. Available verbosity levels are as follows: ■ all ■ profile ■ information ■ warning ■ error ■ none
-noheader	Disables the use of timestamp header.

1. Each Profile must have at least one Source Level.
2. For each Profile:
 - a. For global domains, all intersections {Data Intersection, Profile Intersection, Stored Intersection, Aggregation Intersection, and Approval Intersection} must be below the partition (not HBI).
 - b. Data Intersection (if a data source is specified) must conform to X in {Profile Intersection, Stored Intersection, Aggregation Intersection, and Approval Intersection}.
 - c. Profile Intersection must conform to the Stored Intersection.
 - d. Aggregation Intersection must conform to the Approval Intersection.
 - e. Aggregation Intersection must not be below the Approval Intersection.
 - f. Aggregation Intersection must be above the Data Intersection (if data source specified).
 - g. If the Aggregation Intersection conforms to Profile Intersection:
 - The Profile Type must not be diff(8).
 - The Aggregation Intersection must be above the Profile Intersection.
 - The Aggregation Intersection must be above the Stored Intersection.
 - h. If Aggregation Intersection does not conform to Profile Intersection:
 - The Profile Type must be diff (8).
 - At least one common hierarchy must exist between the Aggregation Intersection and X in {Profile Intersection, Stored Intersection}.
 - For each common non-PROD hierarchy H of Aggregation Intersection and X in {Profile Intersection, Stored Intersection}: Aggregation Intersection's H dimension must not be below X's H dimension.
3. For each Source Level:
 - a. For global domains, all intersections {Profile Intersection, Stored Intersection, and Aggregation Intersection} must be below the partition (not HBI).
 - b. Parent Profile's Data Intersection (if data source specified) must conform to X in {Profile Intersection, Stored Intersection, and Aggregation Intersection}.
 - c. Profile Intersection must conform to Stored Intersection.
 - d. Aggregation Intersection must be above parent Profile's Data Intersection (if data source specified).
 - e. If Aggregation Intersection conforms to Profile Intersection:
 - The Profile Type must not be diff(8).
 - The Aggregation Intersection must be above the Profile Intersection.
 - The Aggregation Intersection must be above the Stored Intersection.
 - f. If Aggregation Intersection does not conform to Profile Intersection:
 - The Parent Profile Type must be diff (8).
 - There must be at least one common hierarchy between the Aggregation Intersection and X in {Profile Intersection and Stored Intersection}.

- For each common non-PROD hierarchy H of Aggregation Intersection and X in {Profile Intersection and Stored Intersection}: Aggregation Intersection's H dimension must not be below X's H dimension.

curvebatch

This section provides details about the curvebatch utility.

curvebatch Usage

```
curvebatch -d domainpath [-level # ] [-debug] | -h | -version
```

[Table A-2](#) provides descriptions of the arguments used by the curvebatch utility.

Table A-2 Arguments Used by the curvebatch Utility

Argument	Description
-d <i>domainpath</i>	Indicates the path to the domain.
-level #	The # signifies the profile ID. When used, a valid profile ID must be provided.
-debug	Causes temporary measures to be retained for debugging purposes.
-h	Displays information about curvebatch in the terminal screen.
-version	Displays version information.
-loglevel	Sets the logger verbosity level. Available verbosity levels are as follows: ■ all ■ profile ■ information ■ warning ■ error ■ none

Appendix: RPAS Test Driver

This appendix describes the RPAS Test Driver (RTD).

Support

Support for the RTD library is limited to its basic understanding and usage of the published XML-based APIs. This library is available on specific versions of the Windows platform where the RPAS release is officially supported.

Oracle does not provide support for compilation or linkage of this library. Also, Oracle does not support LoadRunner scripts or any other scripts used to run the rtd.dll library.

Introduction

This appendix covers the use of rtd.dll provided for the purpose of performance testing of the Oracle Retail RPAS server.

Audience

This DLL allows a user to create, open, alter, save, commit, and refresh workbooks by sending commands directly to the RPAS server through an XML structure.

Knowledge of the usage of the Oracle Retail RPAS client is assumed. A user must have a working domain that can be connected to with a client. A user must also know the internal names of dimensions, positions, workbook templates, worksheets, and windows.

Although the interface provided by this library could be used by any testing tool, it has most frequently been used with HP LoadRunner. The use of RTD with LoadRunner is outlined in later sections.

Finding Required Information

The majority of interactions with RTD assume knowledge of names rather than labels. In order to easily find dimensions and position names, the RPAS client (foundation.exe) can be made to display the internal names rather than labels.

This is accomplished by adding "ShowPositionNames=1" under the [Options] section in %windir%\foundation.ini, for example, C:\Windows\foundation.ini. (The foundation.ini file may be in another location, as described step 3 in "[Translation Administration](#)".) Then, when building and displaying a workbook, the client displays position names instead of labels.

To find template, worksheet, and window names, make a search under your domain's "repos" directory for files named "tpl.cfg" with the label (displayed name) of your template. The template name, worksheet, and window names are contained within this file.

Compatibility

The version of RTD must be compatible with the RPAS server version. If an RPAS hot fix or upgrade requires a client update, then it is likely a matching RTD build will be required. Contact support for further information.

Interface

The C interface provided by the library is simple to interact with, the majority of interactions being with `rtd_send()` and `rtd_getResponse()` functions. The complexity arises in the messages sent and received by these functions, detailed in the next section.

To log on to the server:

```
int rtd_logon(const char* user, const char* password, const char* connection,
const char* domain, const char* fcfPath, const char* language, const char*
logLevel)
```

Or,

```
int rtd_logon_ex(const char* user, const char* password, const char* connection,
const char* domain, const char* fcfPath, const char* language, const char*
logLevel, const char* domainRootDirectory)
```

[Table B-1](#) provides descriptions of the arguments used by `int rtd_logon`.

Table B-1 Arguments Used by `int rtd_logon`

Argument	Description
<code>const char* user</code>	User name.
<code>const char* password</code>	User password.
<code>const char* connection</code>	Connection defined in the .fcf file.
<code>const char* domain</code>	Domain defined in the .fcf file.
<code>const char* fcfPath</code>	Path to the .fcf file.
<code>const char* language</code>	Language of the domain, usually "english". (Currently unused.)
<code>const char* logLevel</code>	Log level of the RPAS server, defaults to "debug" if blank.
<code>const char* domainRootDirectory</code>	The domain path in foundation.fcf that is prepended with this string. This allows the foundation.fcf to specify a relative (or empty) domain path.

To log off the server:

```
void rtd_logoff()
```

To send an XML request to the server:

```
int rtd_send(const char* request)
```

To get the XML response from the server request:

```
char* rtd_getResponse()
```

To get a particular property from the response XML tag:

```
char* rtd_getResponseProperty(const char* property)
```

To free the char buffer that was allocated by `rtd_getResponse()` or `rtd_getResponseProperty()`:


```
void rtd_freeCharPtr(char*)
```

Note: The char buffer must be free in order to prevent memory leaks.

The functions `rtd_logon`, `rtd_logon_ex`, and `rtd_send` indicate success by returning 1 and failure by returning 0.

Figure B-1 shows an example of the expected use of the library. RTD function calls are indicated in the shaded blocks.

Figure B-1 RTD Example Workflow



XML Requests

This section describes requests that can be sent as XML strings to the server. All tags and attributes are case-sensitive.

Connection Messages

There are two connection messages.

<Logon>

Description:

Log a user into a domain, alternative to `rtd_logon()`.

Properties:

user, password, connection, domain, fcfPath, logLevel, and language.

<Logoff>

Description:

Log the current user off, alternative to `rtd_logoff()`.

Wizard Messages

The messages below appear in the order in which they are called. They are repeated as necessary for the wizard pages. The wizard pages can be determined by creating the same workbook in the client.

<WizardInitialize>

Description:

Start wizard to create a new workbook. Some wizards do not have any wizard pages. After calling this, call `rtd_getResponseProperty("skipDisplay")` will be true if this wizard does not have any wizard pages.

Properties:

templateName: Internal name of the workbook template, as defined in `tmpl.cfg`.

<GetNextWizardDialog>

Description:

Move to the next page in a wizard. This must be called after `<WizardInitialize>` to navigate to the first wizard page. After calling this, call `rtd_getResponseProperty("nextDialog")` to return the name of the new wizard page.

Properties:

templateName: Internal name of the workbook template, as defined in `tmpl.cfg`.

currentDialog [optional]: Not currently used.

forward: Direction to move the wizard. Values are "true" to go forward or "false" to go backward, but this should almost always be true.

<GetWizardLayout>

Description:

Returns the layout of a wizard page. The output is a tag called `<wizardPage>`. Inside that tag is a collection of `<control>` tags. Each of the `<control>` tags has attributes for name, class, windowText, styles, xStart, yStart, width, height, enabled, data, option, and associations of a control on the wizard page.

Properties:

templateName: Internal name of workbook template, as defined in `tmpl.cfg`.

currentDialog: Name of wizard page as defined in `tmpl.cfg` or returned by `<GetNextWizardDialog>`.

<SelectPositions>

Description:

Select which positions should be included in the workbook. The `<position>` child tags are sent in the same message as the `<SelectPositions>` tag. `<GetNextWizardDialog>` must be called to navigate to the correct wizard page before calling `<SelectPositions>`. This is typically used on two-tree wizard pages.

Properties:

templateName: Internal name of workbook template, as defined in `tmpl.cfg`.

currentDialog: Name of wizard page defined by the template configuration file.

Children:

`<position>`

Properties:

dimension: Dimension that contains the position.

name: Name of the position.

<SetWizardData>

Description:

This allows explicit control over the UI elements within each page of the wizard. This enables operations beyond simple workbook builds. This is typically used on custom wizard pages.

Properties:

templateName: Internal name of workbook template, as defined in `tmpl.cfg`.

currentDialog: Name of wizard page defined by the template configuration file.

name: UI control name.

class: UI control class.

The control classes include Edit, DropDownList, Listbox, Groupbox, Checkbox, RadioButton, Text, SingleHierSelect, Tree, SuperTree, DatePicker.

The string "true" or "false" is used set radio buttons and check boxes.

The following example illustrates this use in the context of AutoWorkbook builds.

```
<!-- Create the workbook -->
<WizardInitialize templateName="AutoWb" />

<!-- AutoWorkbookPage0 - select action -->
<GetNextWizardDialog templateName="AutoWb" forward="true" />
<SetWizardData templateName="AutoWb" currentDialog="AutoWorkbookPage0"
```

```
name="addWb" class="RadioButton">true</SetWizardData>

<!-- AutoWorkbookPage1 - select template -->
<GetNextWizardDialog templateName="AutoWb" forward="true" />
<SetWizardData templateName="AutoWb" currentDialog="AutoWorkbookPage1"
name="ListBox" class="Listbox">R_TEST</SetWizardData>

<!-- AutoWorkbookPage2 - select owner -->
<GetNextWizardDialog templateName="AutoWb" forward="true" />
<SetWizardData templateName="AutoWb" currentDialog="AutoWorkbookPage2"
name="ListBox" class="Listbox">adm</SetWizardData>

<!-- AutoWorkbookPage3 - enter scheduling data -->
<GetNextWizardDialog templateName="AutoWb" forward="true" />
<SetWizardData templateName="AutoWb" currentDialog="AutoWorkbookPage3"
name="BuildLabel" class="Edit"><value>RTD Auto Build WB</value></SetWizardData>
<SetWizardData templateName="AutoWb" currentDialog="AutoWorkbookPage3"
name="BuildFrequency" class="Edit">1000</SetWizardData>
<SetWizardData templateName="AutoWb" currentDialog="AutoWorkbookPage3"
name="NextBuildDate" class="Edit">05/05/2010</SetWizardData>
<SetWizardData templateName="AutoWb" currentDialog="AutoWorkbookPage3"
name="GroupAccess" class="RadioButton">true</SetWizardData>
<SetWizardData templateName="AutoWb" currentDialog="AutoWorkbookPage3"
name="groupNames" class="Listbox">adm</SetWizardData>

<!-- ... workbook build details omitted ... -->

<WizardFinish templateName="AutoWb" />
```

<WizardFinish>

Description:

Finish the wizard and build the workbook. After calling this, call `rtd_getResponseProperty("workbookName")` to get the internal name of the newly created workbook. Wizard-only templates do not have a `workbookName` and often provide an INFO message. If the workbook cannot be built, an ERROR message is returned. Use `rtd_getResponseProperty("ERROR")` to check the error message and `rtd_getResponseProperty("INFO")` to check the info message.

Properties:

`templateName`: Internal name of workbook template, as defined in `tmpl.cfg`.

Workbook Messages

After a workbook has been created through the wizard process, it must be opened and manipulated. These functions do this for any created workbook.

<OpenWorkbook>

Description:

Open an existing workbook.

Properties:

`workbookName`: Internal name of a stored workbook of the form "txxx" where the x's represent a variable number of digits. If the workbook was just created by the wizard process, use the value returned by `rtd_getResponseProperty()`. See <WizardFinish> above.

access [optional]: Currently unused.

<Fetch>

Description:

Fetch data from the workbook.

Properties:

workbookName: Internal name of a stored workbook. See <OpenWorkbook> for more information.

windowName: Internal name of the window on which this measure is located as defined by `tmpl.cfg`.

Children:

<WindowContext>: Tag that defines what hierarchies and dimensions would be visible if there was a screen. At a minimum, the dimensions of the current cell need to be specified. There should only be one tag per hierarchy.

Properties:

hierarchy: Name of hierarchy.

Children:

<dimension>: Tag that represents the name of a dimension.

Properties:

value: Name of the dimension.

<position>: Tag that represents a position of the cell.

Properties:

dimension: Dimension that contains the position.

name: Name of the position.

hier: Hierarchy that contains the dimension.

axis: Location of this position in the grid if this was on the screen. Values are "row", "column" or "page".

<SetValues>

Description:

Set the values of cells in the workbook.

Properties:

worksheet: Name of sheet as defined in `tmpl.cfg`.

measure: Measure whose value should be set.

spreadMeasure [optional] - Currently unused.

spreadMethod: Method to use when spreading this cell to lower level cells. See the *RPAS Configuration Tools User Guide* for more information.

value: Desired value of the cell.

type: Type of cell. More accurately, this is the type of the measure. Locate this by examining the configuration of the domain with the Oracle RPAS Configuration Tools.

Children:

<WindowContext>: Tag that defines what hierarchies and dimensions would be visible if there was a screen. At a minimum, the dimensions of the current cell need to be specified. There should only be one tag per hierarchy.

Properties:

hierarchy: Name of hierarchy.

Children:

<dimension>: Tag that represents the name of a dimension.

Properties:

value: Name of the dimension.

<position>: Tag that represents a position of the cell.

Properties:

dimension: Dimension that contains the position.

name: Name of the position.

<CommitWorkbook>

Description:

Transfer the data in the workbook back to the domain.

Properties:

synchronous [optional]: If "true", the workbook commits the data now. Otherwise, it is the same as "commit later".

asap [optional]: If "true", the workbook is added to the server's commit ASAP queue.

If the 'synchronous' option is "true", then the 'asap' option is ignored. If both 'synchronous' and 'asap' are omitted or blank, then the default is to commit data now.

<RefreshWorkbook>

Description:

Overwrite the data in the workbook with data from the domain.

Properties:

refreshGroup [optional]: If a workbook has several refresh groups, you may specify the one to use here. If omitted or blank, the default group is used.

<MenuEvent>

Description:

Execute a custom menu event in the current workbook.

Properties:

id: The numeric ID of the custom menu event. The IDs start at 1 and are in the same order as they appear in the Dynamic Template configuration file.

expectedCaption [optional]: Specifies the caption that should be returned. Success is false if the returned caption is different. The return caption is not validated if this is blank.

expectedText [optional]: Specified the text that should be returned. Success is false if the returned text is different. The returned text is not validated if this is blank.

<CreatePOPMeasure>**Description:**

Create a new percent of parent measure (also known as percent participation).

Properties:

window: The name of the window that the new measure is added to.

measure: The name of the "parent" measure for the new percent of parent measure.

hier [optional]: The hierarchy name for a relative percent of parent measure.

hiers [optional]: All hierarchies for an absolute percent of parent measure. Hierarchy names are delimited by a space.

dims [optional]: All dimensions for an absolute percent of parent measure. Dimension names are delimited by a space.

Note: Either 'hier' or both 'hiers' and 'dims' must be specified. Both may not be specified. Specifying 'hier' indicates that the percent of parent measure is relative. Specifying both 'hiers' and 'dims' indicates that the percent of parent measure is absolute. 'hiers' and 'dims' are parallel arrays and their order must match.

<SaveWorkbook>**Description:**

Save the workbook to work on later.

Properties:

workbookLabel: Label (displayed name) for the workbook

workbookName: Internal name of a stored workbook. See <OpenWorkbook> for more information.

user: Access for the current user. Values are "NONE", "READ" or "WRITE".

group: Access for the named group. Values are "NONE", "READ" or "WRITE".

groupName: Name of the group for the above access value.

world: Access for the rest of the world. Values are "NONE", "READ" or "WRITE".

close: Close workbook. Values are "true" or "false".

<CloseWorkbook>**Description:**

Close the currently open workbook.

Properties:

method: If "forward", save and close the workbook. If "backward", close but do not save the workbook.

RTD Use with HP LoadRunner

In this section, familiarity with HP LoadRunner is assumed. For further information on this topic, consult the HP LoadRunner documentation.

The RTD library is not a LoadRunner extension and as such does not provide support for record and replay functionality. Scripts must be hand-coded to suit the RPAS workflow under test and only then can they be replayed using the message passing functions outlined in the previous section.

Prerequisites

RTD functions can be called from any C-based virtual user (vuser). The script must reference a header that contains the RTD function prototypes. Additionally, certain LoadRunner replay settings must be set in order to assure correct behavior.

Header File

This file contains the functions exported by the RTD library; these are in a format usable by the LoadRunner C runtime engine.

lr_rtd.h

```
#ifndef _RTD_RTD_H_
#define _RTD_RTD_H_
int    rtd_logon(const char* user,
                const char* password,
                const char* connection,
                const char* domain,
                const char* fcfPath,
                const char* language,
                const char* logLevel);
int    rtd_logon_ex(const char* user,
                  const char* password,
                  const char* connection,
                  const char* domain,
                  const char* fcfPath,
                  const char* language,
                  const char* logLevel,
                  const char* domainRootDirectory);
void   rtd_logoff();
int    rtd_send(const char* request);
char*  rtd_getResponse();
char*  rtd_getResponseProperty(const char* property);
void   rtd_freeCharPtr(char* ptr);
#endif // _RTD_RTD_H_
```

LoadRunner Scripting

This section contains the details of LoadRunner scripting.

Script Settings

RTD is not thread-safe and so the virtual user script must be run as a process rather than a thread:

- Open run-time settings
- Miscellaneous section
- Select "Run Vuser as a process"

Note: Failure to do this will result in unpredictable and often inexplicable failures at runtime.

Script Sections

All LoadRunner C-based scripts consist of three sections:

- vuser_init
- Actions
- vuser_end

It is recommended that the vuser_init section is used for initialization and the body of the test placed in the Actions section. The vuser_end section is typically left unused.

Script Implementation

This section walks through two examples that validate that LoadRunner and RTD are working correctly together.

Script Section: vuser_init

The vuser_init section, which can be common to all RTD scripts, deals with the DLL load. In order for it to load correctly, the directory containing the library must be in your system PATH. The DLL load should work without the PATH setting, but results can vary from machine to machine. Explicitly setting the PATH guarantees success.

```
#include lr_rtd.h
#define RTD_LIBRARY "c:\\path\\to\\rtd.dll"
vuser_init {

    int rc = lr_load_dll(RTD_LIBRARY);
    if (rc!=0) {
        lr_error_message("rtd library did not load (rc=%d). Not in your %PATH%?", rc);
        lr_abort();
    }
}
```

This section references the lr_rtd.h interface file defined previously. If it is not in the script directory, it must be prefaced by the full path.

Script Section: Action (simple test)

The Actions section contains user actions; these vary on a case-by-case basis. This example uses a simple script to check that the RTD and LR integration is working correctly.

```
Action {
    rtd_logon("user", "pwd", "connect", "dom", "c:\\path\\to\\Foundation.fcf",
"english", "debug");
    lr_think_time(60);

    rtd_logoff();
}
```

This example causes the virtual user to log on, sleep 60 seconds, and then log off. The login details should be altered to suit the RPAS environment under test. If this succeeds, then RTD is correctly installed.

The following is a more realistic example.

Script Section: Action (real-world)

In this more complicated example, a workbook is created and the name retrieved. It is then opened and closed. Additionally, rather than having multiple `rtd_send()` calls in-line, the XML is placed within a LoadRunner parameter file. This is then iterated over within the LoadRunner script.

The parameter file must consist of one complete XML entity per line, single column called "command" with the following contents.

```
<Logon ... />
<WizardInitialize templateName="perfctest"/>
<GetNextWizardDialog templateName="perfctest" forward="true"/>
<SelectPositions templateName="perfctest" currentDialog="wiz_PerfTestphsd"
  hierarchy="PHSH" dimension="PHSD"><position dimension="PHSD"
  name="PHSDP001"/></SelectPositions>
<GetNextWizardDialog templateName="perfctest" forward="true"/>
<SelectPositions templateName="perfctest" currentDialog="wiz_PerfTestweek"
  hierarchy="CLND" dimension="WEEK"><position dimension="WEEK"
  name="WEEK200801"/><position dimension="WEEK"
  name="WEEK200802"/></SelectPositions>
<GetNextWizardDialog templateName="perfctest" forward="true"/>
<SelectPositions templateName="perfctest" currentDialog="wiz_PerfTestitem"
  hierarchy="PROD" dimension="ITEM"><position dimension="ITEM"
  name="ITM000"/><position dimension="ITEM" name="ITM001"/></SelectPositions>
<WizardFinish templateName="perfctest"/>
get_wb_name
<OpenWorkbook workbookName="{_wb}" access="write"/>
<CloseWorkbook method="backwards"/>
<Logoff/>
```

The `vuser_init` section must be identical to the prior example.

The Actions section must contain the following:

```
Action {
char *wb_name, *xml, *response;
  int success;
  while(...) {

if (0==strcmp("get_wb_name", script.param)) {
wb_name=rtd_getResponseProperty("workbookName");
  lr_save_string(wb_name, "_wb");
  rtd_freeCharPtr(wb_name);
} else {
rc=rtd_send(lr_eval_string(xml));
if (!rc) {
  response=rtd_getResponse();
  lr_error_message("error - response was: %s", response);
  rtd_freeCharPtr(response);
  lr_abort();
}
}
}
```

The above section accomplishes a number of things:

- Uses parameters instead of hard-coded calls
- Builds a workbook through the wizard

- Obtains the workbook name
- Saves the workbook name to the `_wb` parameter name
- Replaces subsequent instances of the `_wb` parameter with the workbook name, through `lr_eval_string()`

Considerations for Scenario Use

- When RTD is used in a scenario orchestrated by the LoadRunner controller, it is necessary to have RTD on each of the load generators (injectors). RTD must be in the same path, referenced in the LoadRunner script, on each generator.
- Check the script runtime settings for the scenario. Match those used within the virtual user generator (vugen) environment.
- Ensure that if RTD is upgraded, it is upgraded on each generator.

Appendix: RPAS Test Automation

This appendix describes an RPAS utility called rpac (RPAS Pluggable Automation Component). It can be used to perform various tests on an existing domain, based on a set of test cases contained in XML files.

This appendix contains the following sections:

- [Introduction](#)
- [Usage](#)
- [Writing Test Cases](#)
- [Example](#)
- [Output](#)
- [Schema](#)

Introduction

The rpac utility is used to create highly customizable automated test suites for the RPAS server. The XML-based framework allows rapid development of tests and ensures that test cases can be quickly copied across locations. The flexible XML schema allows the user to test a large number of RPAS features, such as writing data to a domain, building a workbook, running custom menu options, and verifying the results.

The XML framework also introduces consistency, accuracy, and reproducibility in the test results. To reproduce a test case, all that is needed is the XML file and a copy of the domain.

Note that rpac operates directly against the domain itself and therefore does not test any Client/Server interactions.

Usage

To run an rpac test, run the following from the command line:

```
rpac -d <path to domain> -title <title> -input <path to xml input> [commands]  
[options]
```

The output of this command is written to unittest.xml in the local directory. The results of this file contain only the tests that have failed. To display both passing and failing tests, use the -coverage option. This generates a coverage.xml in the local directory that contains both passing and failing results.

Required Parameters

Table C–1 Required Parameters

Option	Parameter	Description
-d	Path to domain	Specifies the path to the domain that the test runs against.
-title	Title	Specifies the name that this test run is given in the output files. This is used when combining the results of multiple test runs.
-input	path to XML script	Specifies the path to the XML script containing the test suite.

Commands

Table C–2 Commands

Command	Description
<none>	If you do not specify a command, then rpac will run the test script provided.
-list	This command switch provides a listing of all of the tests contain in the XML files.
-validate	This command switch validates all the tags and attributes contain in the XML files.
-printTagDoc	Switch to enable just printing the tags documentation to the console.
-writeTagDoc	Switch to enable writing an xml file 'tagdoc.xml' containing the tag documentation.

Options

Table C–3 Options

Option	Parameter	Description
-suite	pattern	A regular expression pattern to specify which test suites should be run. Any test suite whose name does not match the pattern is skipped. Example: If you have the following test suites (test1, test2, test3, test4, test5) and you only want to run suites 1, 2, and 4, specify: -suite test[124]
-coverage	<none>	When specified, rpac writes all passing and failing results to coverage.xml.
-trace		Log all tags executed.
-test		The pattern string of the tests to be run. All tests with a name matching the pattern are run; the rest of the tests are skipped.
-error		File name of the directory to contain the XML results file 'unittest.xml' and the optional coverage file 'coverage.xml'. Uses the current directory if not provided.

Writing Test Cases

The structure of an rpac test script is defined by the following tags.

Table C-4 Test Case Tags

Tag	Description
<testscript>	The outermost tag for the test script. This tag contains one or more <testsuite> tags.
<testsuite>	This tag defines a suite of tests that are designed to run together. Each suite contains its own <setup>, <teardown>, <testcase>, and <workbook-operations> tags. This tag has a name attribute that can be used with the -suite command line attribute to specify which suites are executed in this run.
<once>	The once tag contains tag elements used to set up the environment once before the first test case is executed. It is a tag like setup, but only happens once at the start of a test suite. The once tag has been added to allow a set of measures to be backed up once during the run of the test suite.
<finalize>	The finalize tag contains tag elements used to restore the environment after the final test case is executed.
<setup>	This tag sets the preconditions that are common to all test cases. Operations inside this tag are run before every test case is executed. Example operations: ■ Set RpasToday so the test date is consistent for every test run. ■ Pre-set measure values so tests operate in a known state. ■ Check any preconditions to make sure this test run is valid.
<teardown>	This tag contains cleanup code that is run after all test cases. Operations inside this tag are run after every test case is executed. This is used to restore a domain or workbook to its default state to remove dependencies between independent test cases. Example operations: ■ Set domain measure values back to their initial state. ■ Run batch scripts to remove any test-specific files that were generated.
<workbook-operations>	This tag allows the test script to build, refresh, commit, and close workbooks. Workbook operations may exist outside of test cases, but you must encapsulate them wherever possible inside a testsuite tag for logging purposes.
<testcase>	This tag is where the code to perform each test goes. It should perform operations, then use the assert tags to verify the results. The name of this tag is used in the output to log timing results and success/failure. Note: Code in the <setup> tag is executed prior to each test case, and code in the <teardown> tag is executed after each test case.

The optional attribute role for tags testsuite, once, setup and teardown may be required for an rpac application being written for an RDM enable domain. The role assigned to the testsuite tag affects all test cases, and if no role is assigned to once, setup, and teardown, the role assigned to the testsuite tag will affect these tags as well. A role assigned to the once tag is applied to the tags contained in once, and if no role is assigned to setup and teardown, then this role will apply to these tags as well. A role assigned to the setup tag is applied to the tags contained in setup, and if no role is assigned to teardown, then this role will apply to teardown as well. A role assigned to the teardown tag only applies to the tags contain in teardown.

It is recommended that the teardown section restore the domain to its previous state. This removes dependencies between test cases and ensures consistency from one run to the next.

The tags once, setup, and teardown are unique. Only the first occurrence of each of these tags is used. These tags can appear in any order under the testsuite tag, and all of these tags are optional. The code contained in these tags is executed at a set point

during the run of unit test code. However, for the tag testcase, the order of occurrence in the testsuite is the order of execution.

Example

This section provides a sample XML file that contains some rpac tests that run against an MFP domain. This test illustrates how the setup and teardown methods work, as well as how to use the various workbook operations.

The first test case demonstrates building the workbook. The first step in this operation is to shell out to the wbmgr utility to clear out all existing workbooks. This step prevents the domain from growing after repeated operations, but you should only use this step in cases where you know that existing workbooks are not needed.

After that, the test verifies the preconditions of the test case. In this example, the test verifies the preconditions performed in the setup code by checking that the domain measure value is "true". After that, the test makes the wizard selections and builds the workbook. Once the workbook is built, an assertion verifies that the load rules correctly set the workbook value.

Once this test case is complete, the teardown code is executed, along with the setup code for the next test case. The test verifies that the teardown code correctly set the workbook's measure value to "false" and then issues a commit. Once the workbook is committed, the domain value is checked to see if it was properly updated. Once again, the teardown code is executed and the setup code is run at the beginning of the next test case.

For the third test case, the setup code should have initialized the domain value back to "true", and the workbook value back to "false". The test case verifies these operations, then performs a workbook refresh and verifies that the workbook value has been correctly set to "true".

After that, the workbook is closed and the test suite is complete.

Note: Because the teardown code modifies workbook data, the workbook close step cannot be encapsulated inside of a test case. If it was, the teardown code would try to modify a workbook that was no longer there and produce a failing result.

```
<?xml version="1.0" encoding="UTF-8"?>
<!--
  Copyright (c) 2004, 2012, Oracle and/or its affiliates. All rights reserved.
  File: wctest.xml
  TestCase Description:
  - Setup:
    Set the RPAS Today value
    Set the domain's measure value to "true"
  - Teardown:
    Set the workbook's measure value to "false"
  - Build: Build a workbook and validate the load rule
  - Commit: Validate the commit rule
  - Refresh: Validate the refresh rule
  - Close the workbook
-->
<testscript>
  <testsuite name="MFP RPAC Test 11">
    <!--Setup executed at the beginning of every testcase-->
    <setup>
```

```

    <rpas-today>20070101</rpas-today>
    <edit-domain-measure-value name="ipopappenbb" namedkeys="h1_2007:100:1"
value="true"/>
  </setup>

  <testcase name="Enable OP Approval Workbook: Build">
    <!--Clear all existing workbooks-->
    <shell>wbmgr -d [DOMAIN] -remove -all</shell>

    <!--Verify pre-conditions are true-->
    <assert-domain-measure-value-eq name="ipopappenbb" namedkeys="h1_2007:100:1"
value="true"/>

    <workbook-operations>
      <build template-name="EnableOp" label="Enable OP Approval" user="adm">
        <!--Make positions selection for the wizard page-->
        <wizard-page-settings id="wiz_EnableOpssn">
          <set-tree-selections control-name="tree1" dim-name="ssn"
rollup-name="ssn">
            h1_2007
          </set-tree-selections>
        </wizard-page-settings>
        <wizard-page-settings id="wiz_EnableOpdept">
          <set-tree-selections control-name="tree1" dim-name="dept"
rollup-name="dept">
            100
          </set-tree-selections>
        </wizard-page-settings>
        <wizard-page-settings id="wiz_EnableOpchnl">
          <set-tree-selections control-name="tree1" dim-name="chnl"
rollup-name="chnl" >
            1
          </set-tree-selections>
        </wizard-page-settings>
      </build>
    </workbook-operations>

    <!--Validate the load rule worked correctly->
    <assert-workbook-measure-value-eq name="ipopappenbb" keys="0:0:0"
value="true"/>
  </testcase>

  <!--Commit the workbook and assert the measure values in domain-->
  <testcase name="Enable OP Approval Workbook: Commit">
    <assert-domain-measure-value-eq name="ipopappenbb" namedkeys="h1_2007:100:1"
value="true"/>
    <assert-workbook-measure-value-eq name="ipopappenbb" namedkeys="h1_
2007:100:1 value="false"/>
    <workbook-operations>
      <commit type="now" template-name="EnableOp" label="Enable OP Approval"
user="adm"/>
    </workbook-operations>

    <!--Validate the commit rule worked correctly->
    <assert-domain-measure-value-eq name="ipopappenbb" namedkeys="h1_2007:100:1"
value="false"/>
  </testcase>

  <!--Refresh the workbook and assert measure values in workbook and domain-->
  <testcase name="Enable OP Approval Workbook: Refresh">

```

```

        <assert-domain-measure-value-eq name="ipopappenbb" namedkeys="h1_2007:100:1"
value="true"/>
        <assert-workbook-measure-value-eq name="ipopappenbb" namedkeys="h1_
2007:100:1" value="false"/>
        <workbook-operations>
            <refresh template-name="EnableOp" label="Enable OP Approval" user="adm"/>
        </workbook-operations>

        <!--Validate the refresh rule worked correctly-->
        <assert-workbook-measure-value-eq name="ipopappenbb" namedkeys="h1_
2007:100:1" value="true"/>
    </testcase>

    <!--Close the workbook-->
    <workbook-operations>
        <close template-name="EnableOp" label="Enable OP Approval" user="adm"/>
    </workbook-operations>

    <!--Teardown executed at the end of every testcase-->
    <teardown>
        <edit-workbook-measure name="ipopappenbb" keys="0:0:0" value="false"/>
    </teardown>
</testsuite>
</testscript>

```

Output

When rpac executes, it prints the status of the current execution to the console. Once all the test cases are complete, the summary of test cases is stored in unittest.xml under the current directory. These results contain the failing test case, as well as any error messages that were captured. If -coverage was specified, then a coverage.xml file is generated containing the passing and failing results.

Example:

```
[rpasHome] rpac -d testDomain -input wbTest.xml -title "'MFP OpApproval RPAC Test"
```

```
Running rpac tests for RPAC : MFP Tests...
```

```
Now executing wbmgr in domain.
Deleting all workbooks in the domain...
Destroying workbook 't0000000000' in domain.
wbmgr completed successfully
5.292541
0.308167
0.244300
```

```
Ran 3 tests: 0 errors, 0 failures in 5.85 seconds
```

If something had gone wrong, the output would look like this:

Example:

```
[rpasHome] rpac -d testDomain -input wbTest.xml -title "'MFP OpApproval RPAC
Test"
```

```
Now executing wbmgr in domain
Deleting all workbooks in the domain...
Destroying workbook 't0000000000' in domain
wbmgr completed successfully
```

```
Failure
0.192497
Failure
```

```
Ran 3 tests: 0 errors, 2 failures in 6.42 seconds
```

Contents of unittest.xml

After execution, rpac creates a unittest.xml file containing the test results. Here is an example from the successful execution above:

```
<unittest name='MFP OpApproval RPAC Test' runcount='3' timestamp='2012-10-17
10:06:49' elapsedTime='13.72 seconds'>
</unittest>
```

Here is the unittest.xml output from the failing run:

```
<unittest name='MFP OpApproval RPAC Test' runcount='3' timestamp='2012-10-17
10:07:38' elapsedTime='6.42 seconds'>
  <failurelist>
    <failure>
      <testname> </testname>
      <message>Got Error: Expected measure 'ipopappenbb' at ArrayKey(0,
0, 0, 0) is false;
                                false == true</message>
      <file>OpApproval.xml</file>
      <line>52</line>
    </failure>
    <failure>
      <testname> </testname>
      <message>Got Error: Expected measure 'ipopappenbb' at ArrayKey(0,
0, 0, 0) is false;
                                false == true</message>
      <file>OpApproval.xml</file>
      <line>60</line>
    </failure>
  </failurelist>
</unittest>
```

Schema

This section provides schema details.

1. Top-Level Parent Tag

This section provides details about the top-level parent tag.

1.1 testscript

The testscript is the top-level tag, which contains all the rpac tags.

1.2 testsuite

The testsuite contains setup, teardown, and testcase tag elements.

1.3 testcase

The testcase contains a list of tag elements to be executed.

1.4 once

The once tag contains tag elements used to set up the environment once before the first test case is executed. It is a tag like setup, but only happens once at the start of a test suite. The once tag has been added to allow a set of measures to be backed up once during the run of the test suite.

1.5 finalize

The finalize tag contains tag elements used to restore the environment after the final test case is executed.

1.6 setup

The setup tag is used to set any preconditions that must be true prior to running each test case. These can include initializing measure values, setting the date, or running batch scripts to prepare the domain.

1.7 teardown

The teardown tag contains tag elements used to restore the environment to a default state after a test case is executed. This can include returning measures to their default values or running batch scripts to clean up the domain.

1.8 log

The log tag logs a message to the console. It has an optional attribute, msg, containing the message to be logged.

Sample Usage

Here is the sample implementation of the above tags

```
<testscript>
<testsuite name="MFP RPAC Test">
<setup>
<!-- setup code -->
</setup>
<testcase name="Measure Analysis Workbook : build">
<!--testcase code -->
</testcase >
<teardown>
<!-- teardown code -->
</teardown>
<finalize>
    <log msg="Do Finalize" />
</finalize>
</testsuite>
</testscript>
```

2. Generic Tags

This section provides details about generic tags.

2.1 Attribute

The attribute tag specifies the attributes of a tag in tagdoc.xml.

2.2 Description

The description describes the functionality of each tag and its attributes in tagdoc.xml.

2.3 Fail

Fail creates a failure message.

Attribute	Description	Category
msg	Message to be used in the failure text	Optional

2.4 key

Sets the key specs value of the named key.

Attribute	Description	Category
name	Key name	Required
keyspec	Key value spec	Required

2.5 rpas-today

Sets the RpasToday environment variable to the value of this tag. RPAS_TODAY tells RPAS what day that it should think today is.

Sample usage:

```
<rpas-today>20070101</rpas-today>
```

Where: the date is given in yyyyymmdd format.

2.6 shell

The shell tag starts a shell process with the command provided as value of this tag.

Attribute	Description	Category
working-directory	Directory to use for shell; if not provided then the current working directory is used.	Optional
check-status	Boolean to determine if a non-zero return status throws an exception.	Optional
is-shell	Boolean to indicate that the process is a shell.	Optional
assert-status-equal	Assert return status equals provide value.	Optional
assert-status-not-equal	Assert return status does not equals provide value.	Optional
output-file	File to which shell command redirects its output.	Optional

Sample usage:

```
<shell check-status="true" is-shell="false">wbmgr -d [DOMAIN] -remove -all</shell>
```

2.7 register-measure

Registers a measure with the provided properties.

Attribute	Description	Category
name	Measure name	Required
Intx	Base intersection	Required
type	Type of measure	Required
naval	NA value of the measure	Required
defagg	Default aggregation	Required

Sample usage:

```
<register-measure name="MeasureX" intx="WEEK_STR_DAY_" type="int" defagg="total"
naval="0" />
```

2.8 unregister-measure

Unregisters a measure registered by rpac.

Attribute	Description	Category
name	Measure name	Required

Sample usage:

```
<unregister-measure name="MeasureX" />
```

3. Domain Operation Tags

This section provides details about domain operation tags.

3.1 Assertions

This section provides details about assertions.

3.1.1 assert-dimension-contain Assert a dimension contains or does not contain a position name.

Attribute	Description	Category
name	Dimension name	Required
posname	List of comma-separated position names	Required
containflag	If flag is true, then position names are checked to see if they are contained in the dimension.	Required

Sample usage:

```
<assert-dimension-contain name="DEPT" posname="dept1" containflag="true" />
```

3.1.2 assert-dimension-exists Assert a dimension name exists.

Attribute	Description	Category
name	Dimension name	Required

Sample usage:

```
<assert-dimension-exists name="DEPT" />
```

3.1.3 assert-dimension-size Assert a dimension size is between a minimum and a maximum value.

Attribute	Description	Category
name	Dimension name	Required
min	Minimum value	Required
max	Maximum value	Required

Sample usage:

```
<assert-dimension-size name="DEPT" min="8" max="32" />
```

3.1.4 assert-domain-measure-exists Assert a domain measure exists.

Attribute	Description	Category
name	Dimension name	Required

Sample usage:

```
<assert-domain-measure-exists name="r_ex_test" />
```

3.1.5 assert-domain-measure-intx Assert a domain measure base intersection matches an expected value.

Attribute	Description	Category
name	Dimension name	Required
intx	Intersection string to compare with	Required

Sample usage:

```
<assert-domain-measure-intx name="r_ex_test" intx="str_sku_week" />
```

3.1.6 assert-domain-measure-value-eq Assert a domain measure value is equal to an expected value. This tag can be run in a number of ways:

Check against a loadmeasure file:

By specifying <path>, <start>, and <width>, you can validate that the contents of a measure are equal to the contents of a fixed-width loadmeasure file.

Check against specified positions:

By providing <keys> or <namedkeys>, you can verify that all of the positions in your specification are equal to the provided value. You can specify <keyint> to provide positions at an alternate rollup.

Check entire measure:

If you do not supply the <keys>, <namedkeys>, or <path> tag, then rpac will assert that every logical cell in the measure is equal to the value provided.

Attribute	Description	Category
name	Measure name	Required
keys	List of comma-separated position names	Optional
namedkeys	List of comma-separated position names	Optional
path	Path to load measure type file to compare with	Optional
keyint	Intersection string	Optional
value	Expected value	Optional
start	Starting column of measure field	Optional
width	Width of measure field	Optional

Sample usage (using keys):

```
<assert-domain-measure-value-eq name="bulylagtx" keyint="year" namedkeys=" 2007" value="Test1" />
```

Sample usage (using path):

```
<assert-domain-measure-value-eq name="bulylagtx" path="loadmeas.dat" start="100" width="20" />
```

Sample usage (check all):

```
<assert-domain-measure-value-eq name="bulylagtx" value="Test1" />
```

3.1.7 assert-domain-measure-value-ge Assert a domain measure value is greater than or equal to an expected value. This tag can be run in a number of ways:

Check against a loadmeasure file:

By specifying <path>, <start>, and <width>, you can validate that the contents of a measure are equal to the contents of a fixed-width loadmeasure file.

Check against specified positions:

By providing <keys> or <namedkeys>, you can verify that all of the positions in your specification are equal to the provided value. You can specify <keyint> to provide positions at an alternate rollup.

Check entire measure:

If you do not supply the <keys>, <namedkeys>, or <path> tag, then rpac will assert that every logical cell in the measure is equal to the value provided.

Attribute	Description	Category
name	Measure name	Required
keys	List of comma-separated position names	Optional
namedkeys	List of comma-separated position names	Optional
path	Path to load measure type file to compare with	Optional

Attribute	Description	Category
keyint	Intersection string	Optional
value	Expected value	Optional
start	Starting column of measure field	Optional
width	Width of measure field	Optional

Sample usage (using keys):

```
<assert-domain-measure-value-ge name="bulylagtx" keyint="year" namedkeys=" 2007" value="Test1"/>
```

Sample usage (using path):

```
<assert-domain-measure-value-ge name="bulylagtx" path="loadmeas.dat" start="100" width="20" />
```

Sample usage (check all):

```
<assert-domain-measure-value-ge name="bulylagtx" value="Test1" />
```

3.1.8 assert-domain-measure-value-gt Assert a domain measure value is greater than a an expected value. This tag can be run in a number of ways:

Check against a loadmeasure file:

By specifying <path>, <start>, and <width>, you can validate that the contents of a measure are equal to the contents of a fixed-width loadmeasure file.

Check against specified positions:

By providing <keys> or <namedkeys>, you can verify that all of the positions in your specification are equal to the provided value. You can specify <keyint> to provide positions at an alternate rollup.

Check entire measure:

If you do not supply the <keys>, <namedkeys>, or <path> tag, then rpac will assert that every logical cell in the measure is equal to the value provided.

Attribute	Description	Category
name	Measure name	Required
keys	List of comma-separated position names	Optional
namedkeys	List of comma-separated position names	Optional
path	Path to load measure type file to compare with	Optional
keyint	Intersection string	Optional
value	Expected value	Optional
start	Starting column of measure field	Optional
width	Width of measure field	Optional

Sample usage (using keys):

```
<assert-domain-measure-value-gt name="bulylagtx" keyint="year" namedkeys=" 2007"
```

```
value="Test1"/>
```

Sample usage (using path):

```
<assert-domain-measure-value-gt name="bulylagtx" path="loadmeas.dat" start="100"
width="20" />
```

Sample usage (check all):

```
<assert-domain-measure-value-gt name="bulylagtx" value="Test1" />
```

3.1.9 assert-domain-measure-value-le Assert a domain measure value is less than or equal to an expected value. This tag can be run in a number of ways:

Check against a loadmeasure file:

By specifying <path>, <start>, and <width>, you can validate that the contents of a measure are equal to the contents of a fixed-width loadmeasure file.

Check against specified positions:

By providing <keys> or <namedkeys>, you can verify that all of the positions in your specification are equal to the provided value. You can specify <keyint> to provide positions at an alternate rollup.

Check entire measure:

If you do not supply the <keys>, <namedkeys>, or <path> tag, then rpac will assert that every logical cell in the measure is equal to the value provided.

Attribute	Description	Category
name	Measure name	Required
keys	List of comma-separated position names	Optional
namedkeys	List of comma-separated position names	Optional
path	Path to load measure type file to compare with	Optional
keyint	Intersection string	Optional
value	Expected value	Optional
start	Starting column of measure field	Optional
width	Width of measure field	Optional

Sample usage (using keys):

```
<assert-domain-measure-value-le name="bulylagtx" keyint="year" namedkeys=" 2007"
value="Test1"/>
```

Sample usage (using path):

```
<assert-domain-measure-value-le name="bulylagtx" path="loadmeas.dat" start="100"
width="20" />
```

Sample usage (check all):

```
<assert-domain-measure-value-le name="bulylagtx" value="Test1" />
```

3.1.10 assert-domain-measure-value-lt Assert a domain measure value is less than an expected value. This tag can be run in a number of ways:

Check against a loadmeasure file:

By specifying <path>, <start>, and <width>, you can validate that the contents of a measure are equal to the contents of a fixed-width loadmeasure file.

Check against specified positions:

By providing <keys> or <namedkeys>, you can verify that all of the positions in your specification are equal to the provided value. You can specify <keyint> to provide positions at an alternate rollup.

Check entire measure:

If you do not supply the <keys>, <namedkeys>, or <path> tag, then rpac will assert that every logical cell in the measure is equal to the value provided.

Attribute	Description	Category
name	Measure name	Required
keys	List of comma-separated position names	Optional
namedkeys	List of comma-separated position names	Optional
path	Path to load measure type file to compare with	Optional
keyint	Intersection string	Optional
value	Expected value	Optional
start	Starting column of measure field	Optional
width	Width of measure field	Optional

Sample usage (using keys):

```
<assert-domain-measure-value-lt name="bulylagtx" keyint="year" namedkeys=" 2007" value="Test1"/>
```

Sample usage (using path):

```
<assert-domain-measure-value-lt name="bulylagtx" path="loadmeas.dat" start="100" width="20" />
```

Sample usage (check all):

```
<assert-domain-measure-value-lt name="bulylagtx" value="Test1" />
```

3.1.11 assert-domain-measure-value-ne Assert a domain measure value is not equal to an expected value. This tag can be run in a number of ways:

Check against a loadmeasure file:

By specifying <path>, <start>, and <width>, you can validate that the contents of a measure are equal to the contents of a fixed-width loadmeasure file.

Check against specified positions:

By providing <keys> or <namedkeys>, you can verify that all of the positions in your

specification are equal to the provided value. You can specify <keyint> to provide positions at an alternate rollup.

Check entire measure:

If you do not supply the <keys>, <namedkeys>, or <path> tag, then rpac will assert that every logical cell in the measure is equal to the value provided.

Attribute	Description	Category
name	Measure name	Required
keys	List of comma-separated position names	Optional
namedkeys	List of comma-separated position names	Optional
path	Path to load measure type file to compare with	Optional
keyint	Intersection string	Optional
value	Expected value	Optional
start	Starting column of measure field	Optional
width	Width of measure field	Optional

Sample usage (using keys):

```
<assert-domain-measure-value-ne name="bulylagtx" keyint="year" namedkeys=" 2007" value="Test1" />
```

Sample usage (using path):

```
<assert-domain-measure-value-ne name="bulylagtx" path="loadmeas.dat" start="100" width="20" />
```

Sample usage (check all):

```
<assert-domain-measure-value-ne name="bulylagtx" value="Test1" />
```

3.1.12 assert-domain-popcount-eq Assert a domain measure popcount matches an expected value.

Attribute	Description	Category
name	Measure name	Required
pop-count	Expected popcount	Required

Sample usage:

```
<assert-domain-popcount-eq name=" Ipopappenbb" pop-count="100" />
```

3.1.13 assert-domain-popcount-eq-zero Assert a domain measure is not populated.

Attribute	Description	Category
name	Measure name	Required

Sample usage:

```
<assert-domain-popcount-eq-zero name=" Ipopappenbb" />
```

3.1.14 assert-domain-popcount-ne-zero Assert a domain measure is populated.

Attribute	Description	Category
name	Measure name	Required

Sample usage:

```
<assert-domain-popcount-ne-zero name=" Ipopappenbb" />
```

3.1.15 assert-bitsize-value-eq Assert the dimInfo bitsize or dimregistry bitsize value of a dimension.

Attribute	Description	Category
name	List of comma-separated dimension names	Required
category	Bit Size Category (dimInfo or dimRegistry)	Required
value	comma-separated Bit Size Values	Required

Sample usage:

```
<assert-bitsize-value-eq name="week,skup," category="dimInfo" value="12,14" />
```

3.1.16 assert-dim-registry-version Assert the value for dimension registry version.

Attribute	Description	Category
value	Registry version	Required

Sample usage:

```
<assert-dim-registry-version value="1" />
```

3.1.17 assert-reindex-threshold Assert the threshold value of the dimension.

Attribute	Description	Category
name	List of comma-separated dimension names	Required
value	List of comma-separated threshold values	Required

Sample usage:

```
<assert-reindex-threshold name="scls,week" value="25,103" />
```

3.1.18 assert-reindex-required Assert whether Diminfo Bitsize is equal to DimRegistry Bitsize.

Attribute	Description	Category
bitsize-flag	Returns true if a Bitsize mismatch occurs	Required
name	Dimension name	Required

Sample usage:

```
<assert-reindex-required name="week" bitsize-flag="false" />
```

3.1.19 assert-reindex-in-progress Assert the status of reindex with the user input.

Attribute	Description	Category
flag	Flag to indicate whether reindexing Domain is in progress.	Required

Sample usage:

```
<assert-reindex-in-progress flag="false" />
```

3.1.20 set-current-domain Set the current domain to either the master or a sub-domain.

Attribute	Description	Category
index	Domain index	Required

Sample usage:

```
< set-current-domain index="0" />
```

3.1.21 backup-domain-measures Back up a list of measures.

Attribute	Description	Category
names	List of measure names	Required

Sample usage:

```
<backup-domain-measures names="UTDyDst03U,UTDyDst05U,tdwpunelapbopb,bucpgmpp" />
```

3.1.22 restore-domain-measures Restore a list of measures.

Attribute	Description	Category
names	List of measure names	Required

Sample usage:

```
<restore-domain-measures names="UTDyDst03U,UTDyDst05U,tdwpunelapbopb,bucpgmpp" />
```

3.1.23 assert-domain-hier-exists Assert a hierarchy exists in domain or not.

Attribute	Description	Category
name	Hierarchy name	Required

Attribute	Description	Category
exists	If flag is true, then check if hierarchy exists in domain.	Required

Sample usage:

```
<assert-domain-hier-exists name="clnd" exists="true"/>
```

3.2 Domain Data Edit Operations

This section provides details about domain data edit operations.

3.2.1 clear-all-measures Clears all measures used in the following tags: `<edit-workbook-measure>`, `<set-workbook-measure>`, `<set-domain-measure>`.

This sets all the contents of these measures to the NA value.

3.2.2 clear-domain-measure Clear a domain measure of all data. This sets all the contents of this measure to its NA value.

Attribute	Description	Category
name	Measure name	Required

Sample usage:

```
<clear-domain-measure name="r_ex_pick_real" />
```

3.2.3 clear-domain-measure-list Clear a list of domain measures of all data.

Attribute	Description	Category
path	Path to a file containing a new-line delimited list of measures	Required

Sample usage:

```
<clear-domain-measure-list path="measlist.txt" />
```

3.2.4 execute-expression Execute an expression.

Attribute	Description	Category
expression	The expression to execute/calculate. This is a full expression string,	Optional
incremental	True if this is to be an incremental evaluation.	Optional

Sample usage:

```
<execute-expression expression="a = b+c" incremental="false" />
```

3.2.5 set-domain-measure Set the value of a domain measure based on provided attributes.

Attribute	Description	Category
name	Measure name	Required
keys	List of comma-separated position names	Optional
namedkeys	List of comma-separated position names	Optional
path	Path to load measure type file to compare with	Optional
keyint	Intersection string	Optional
value	Expected value	Optional
start	Starting column of measure field	Optional
width	Width of measure field	Optional
incremental	True if evaluation is to be incremental	Optional

Note: Note that for Date type measures, RPAC supports setting the array cell using either of the following two Date & Time formats - "%m/%d/%Y" or "%m/%d/%Y-%H:%M:%S" where m = 2 digit month, d = 2 digit date, Y = 4 digit year, H = 2 digit hour, M = 2 digit min, and S = 2 digit sec.

Sample usage:

```
<set-domain-measure name="r_ex_pick_real" keyint="dept" keys="1" value="150" type="real" />
```

Sample usage (using keys):

```
<set-domain-measure name="meas1" keys="0-2:5" value="0.0" />
```

Sample usage (using keyint):

```
<set-domain-measure-value-ne name="meas2" namedkeys="chnl1:week1" value="02/24/2012" />
```

Sample usage (load-measure file):

```
<set-domain-measure name="meas1" path="loadMeas1.dat" start="100" width="20"/>
```

Sample usage (set all):

```
<set-domain-measure name="meas1" value="2"
```

3.2.6 dump-domain-measure Dump a measure to a text file.

Attribute	Description	Category
name	Measure name	Required
path	File path to text file to be used	Required
noValues	Flag to disable dumping of cell values.	Optional

Sample usage:

```
<dump-domain-measure name="utdydst03u" path="test3.txt" noValues="true"/>
```

3.2.7 load-domain-measure Load a measure from a text file.

Attribute	Description	Category
name	Measure name	Required
path	File path to text file to be used	Required

Sample usage:

```
<load-domain-measure name="utdydst03u" path="test3.txt" />
```

4. Workbook Operation Tags

This section provides details about workbook operation tags.

4.1 Workbook Wizard Operations

This section provides details about workbook wizard operations.

4.1.1 build Builds workbook based on provided tag data.

Attribute	Description	Category
template-name	Template name	Required
label	Label for workbook	Required
user	User used to access workbook	Required

Sample usage:

```
<build template-name="MeasureAnalysis" label="MEAS AN" user="adm">
```

4.1.2 wizard-page-settings Contains the settings for a wizard page.

Attribute	Description	Category
id	ID use to match up settings with wizard pages. For configured templates, this follows the pattern: wiz_<wbt_name><base_dim> Example: For the EnableOp workbook, the wizard page that sets CHNL is: Fwiz_EnableOpchnl Note: For custom templates, obtain the name from the administrator.	Required

Sample usage:

```
<wizard-page-settings id="wiz_EnableOpchnl">
```

4.1.3 set-hier-selection Sets the selection of a SingleHier Select control of a custom wizard page.

Attribute	Description	Category
control-name	Name of the control for a wizard page. This can be obtained from the template.cfg file.	Required
dim-name	Name of the dimension containing the selection.	Required
selection	Name of the position selected.	Required

Sample usage:

```
<set-hier-selection control-name="loc_select" dim-name="chnl" selection="1" >
```

4.1.4 set-selected Sets the value for a Boolean control on a custom wizard page.

Attribute	Description	Category
control-name	Name of the control for a wizard page. This can be obtained from the configuration	Required
selected	Set this to TRUE or FALSE	Required

Sample usage:

```
<set-selected control-name="UseExtraMeasures" value="TRUE">
```

4.1.5 set-selection Sets the selection of a drop-down list control on a custom wizard page.

Attribute	Description	Category
control-name	Name of the control for a wizard page. This can be obtained from the configuration.	Required
selection	Text for the selection operation.	Required

Sample usage:

```
<set-selection control-name="ForecastLevel" value="2">
```

4.1.6 set-selections Sets the selections of a list box or tree control on a custom wizard page. This tag should contain a space-delimited list of the selections.

Attribute	Description	Category
control-name	Name of the control for a wizard page. This can be obtained from the configuration.	Required

Sample usage:

```
<set-selections control-name="ExtraMeasures">
meas1 meas2 meas3 </set selections>
```

4.1.7 set-text Sets the contents of an edit control on a custom wizard page.

Attribute	Description	Category
control-name	Name of the control for a wizard page	Required
test	Text for the operation	Required

Sample usage:

```
<set-text control-name="MyText" value="abc">
```

4.1.8 set-tree-selections Sets the selections for a two-tree control on either a standard or custom wizard page.

Attribute	Description	Category
control-name	Name of the control. For two-tree controls, this should always be set to "tree1"	Required
rollup-name	The highest level of the wizard page.	Required
dim-name	Name of the dimension on which the selections are made.	Required

Sample usage:

```
<set-tree-selections control-name="tree1" rollup-name="chn1" dim-name="str">
str1,str2,str3
</set-tree-selections>
```

4.2 Workbook Operations

This section provides details about workbook operations.

4.2.1 Workbook-operations Contains the workbook operations to be executed.

4.2.2 Calc Does workbook calculations.

Attribute	Description	Category
template-name	Template name	Required
label	Label for workbook	Required
user	User used to access workbook	Required

Sample usage:

```
<calc template-name="RptAdmin" label="Reporting Administration" user="adm"/
```

4.2.3 custom-menu Checks for the existence of, or executes, a custom menu.

Attribute	Description	Category
template-name	Template name	Required
label	Label for workbook	Required
user	User used to access workbook	Required

Attribute	Description	Category
menu-label	Menu label for this custom menu. This can be obtained from the configuration	Required
menu-exist	Set this to "true" to see if the menu item exists, or leave it blank to execute the custom menu.	Optional

Sample usage:

```
<custom-menu template-name="BottomUp" label="Botup" user="adm"
menu-label="menuItem83" menu-exist="true"/>
```

4.2.4 refresh Does workbook refresh.

Attribute	Description	Category
template-name	Template name	Required
label	Label for workbook	Required
user	User used to access workbook	Required
rule-group	Set this to use a rule group other than the default refresh group.	Optional

Sample usage:

```
<refresh template-name="BottomUp" label="Botup" user="adm" rule-group="refresh_
alt"/>
```

4.2.5 Close Does workbook close. Arguments may be provided for debug purposes.

Attribute	Description	Category
template-name	Template name	Optional
label	Label for workbook	Optional
user	User used to access workbook	Optional

Sample usage:

```
<close template-name="BottomUp" label="Botup" user="adm"/>
```

4.2.6 Does workbook commit Does workbook commit.

Attribute	Description	Category
type	Commit type (now/later)	Required
template-name	Template name	Required
label	Label for workbook	Required
user	User used to access workbook	Required

Sample usage:

```
<commit type="now" template-name="BottomUp" label="Bottom Up MFP Refresh"
user="adm" />
```

4.3 Workbook Assertions

This section provides details about workbook assertions.

4.3.1 assert-window-contain-measure Asserts that the measures provided are visible (or not visible) on the window.

Attribute	Description	Category
name	Name of the window in the workbook configuration. User-generated windows cannot be used.	Required
measure	Comma-separated list of measures names.	Optional
containflag	If true, will check that measures provided are visible. If false, will check that the measures provided are not visible.	Optional

Sample usage:

```
<assert-window-contain-measure name="chnldeptssn__W" measures="Ipopappenbb"/>
```

4.3.2 assert-window-exists Asserts if a window exists in the current workbook.

Attribute	Description	Category
name	Window name of current workbook	Required

Sample usage:

```
<assert-window-exists name="chnldeptssn__W" />
```

4.3.3 assert-window-intersection Asserts window intersection matches provided intersection.

Attribute	Description	Category
name	Window name of current workbook.	Required
baseint	Intersection string to compare. Leave blank if checking for scalar.	Optional

Sample usage:

```
<assert-window-intersection name="chnldeptssn__W" baseint="chnldeptssn_" />
```

4.3.4 assert-window-not-exists Asserts if a window does not exist in the current workbook.

Attribute	Description	Category
name	Window name of current workbook	Required

Sample usage:

```
<assert-window-not-exists name="chnldeptssn__W" />
```

4.3.5 assert-window-measure-exists Asserts a measure exists.

Attribute	Description	Category
name	Measure name	Required

Sample usage:

```
<assert-measure-exists name=" Ipopappenbb" />
```

4.3.6 assert-workbook-measure-intx Asserts a workbook measure base intersection matches an expected value.

Attribute	Description	Category
name	Measure name	Required
intx	Intersection string to compare with	Required

Sample usage:

```
<assert-measure-intersection name=" Ipopappenbb" baseint"chnldeptssn_" />
```

4.3.7 assert-workbook-measure-value-eq Asserts that a workbook measure at the specified positions is equal to the provided value.

Parent tag(s): testcase

Child tag(s): None

Attribute	Description	Category
keys	The keys to be used to compare measure values. Can be either indices or position names. Each dimension spec in the keys-spec is delimited by a ":". Multiple position ranges for a dimension can be specified by using "," as separators. When using indices, a range of positions can also be used. Ex - "0-2:4-7,11-13:4" implies indices 0,1,2 for the first dimension, 4,5,6,7,11,12,13 for the second, and 4 for the third dimension of the measure. Note: Because indices are based on the workbook selections, they may not be consistent across all test cases. Care must be exercised to ensure that you are using the correct indices for your specific test case.	Required
namedkeys	Position names used to get keys of the measure array. Using indices instead of position names or positions that do not exist in the MeasureStore (Workbook) will be flagged as an error. Similar to "keys" above, but using position names only.	Optional

Attribute	Description	Category
path	Path of the flat file used for comparison against the Measure array. This looks like pointing to a pre-set environment variable. "path" cannot be used along with "keys" or "namedkeys". "path" is used only to compare the measure data against a pre-existing flat file. If none of "path", "keys" and "namedkeys" are specified, the comparison occurs for all the keys in the measure array of the Measure Store (Workbook).	Optional
start	The field start value. If this does not exist in the XML, the value will be taken from the measure's properties. Used in flat-file comparison.	Optional
width	The field start value. If this does not exist in the XML, the value will be taken from the measure's properties. Used in flat-file comparison.	Optional
value	The expected value of the measure. Comparison on equality.	Optional

Sample usage: (indices)

```
<assert-workbook-measure-value-eq name="drwppthrscisp" keys="0-2" value="0.0" />
```

Sample usage (names):

```
<assert-workbook-measure-value-eq name="drwppthrscisp" namedkeys="pos1,pos2,pos3" value="0.0" />
```

4.3.8 assert-workbook-measure-value-ge Asserts a workbook measure value at the specified positions is greater than or equal to the provided value.

Attribute	Description	Category
keys	The keys to be used to compare measure values. Can be either indices or position names. Each dimension spec in the keys-spec is delimited by a ":". Multiple position ranges for a dimension can be specified by using "," as separators. When using indices, a range of positions can also be used. Ex - "0-2:4-7,11-13:4" implies indices 0,1,2 for the first dimension, 4,5,6,7,11,12,13 for the second, and 4 for the third dimension of the measure. Note: Because indices are based on the workbook selections, they may not be consistent across all test cases. Care must be exercised to ensure that you are using the correct indices for your specific test case.	Required

Attribute	Description	Category
namedkeys	Position names used to get keys of the measure array. Using indices instead of position names or positions that do not exist in the MeasureStore (Workbook) will be flagged as an error. Similar to "keys" above, but using position names only.	Optional
path	Path of the flat file used for comparison against the Measure array. This looks like pointing to a pre-set environment variable. "path" cannot be used along with "keys" or "namedkeys". "path" is used only to compare the measure data against a pre-existing flat file. If none of "path", "keys" and "namedkeys" are specified, the comparison occurs for all the keys in the measure array of the Measure Store (Workbook).	Optional
start	The field start value. If this does not exist in the XML, the value is taken from the measure's properties. Used in flat-file comparison.	Optional
width	The field start value. If this does not exist in the XML, the value is taken from the measure's properties. Used in a flat-file comparison.	Optional
value	The expected value of the measure. Comparison on equality.	Optional

Sample usage (indices):

```
<assert-workbook-measure-value-ge name="drwppthrsc1sp" keys="0-2" value="0.0" />
```

Sample usage (names):

```
<assert-workbook-measure-value-ge name="drwppthrsc1sp" namedkeys="pos1,pos2,pos3" value="0.0" />
```

4.3.9 assert-workbook-measure-value-gt Asserts a workbook measure value is greater than a provided value.

Attribute	Description	Category
keys	The keys to be used to compare measure values. Can be either indices, or position names. Each dimension spec in the keys-spec is de-limited by a ":". Multiple position ranges for a dimension can be specified by using "," as separators. When using indices, a range of positions can also be used. Ex - "0-2:4-7,11-13:4" implies indices 0,1,2 for the first dimension, 4,5,6,7,11,12,13 for the second and 4 for the third dimension of the measure. Note: Because indices are based on the workbook selections, they may not be consistent across all test cases. Care must be exercised to ensure that you are using the correct indices for your specific test case.	Required
namedkeys	Position names used to get keys of the measure array. Using indices instead of position names or positions that do not exist in the MeasureStore (Workbook) will be flagged as an error. Similar to "keys" above, but using position names only.	Optional
path	Path of the flat file used for comparison against the Measure array. This looks like pointing to a pre-set environment variable. "path" cannot be used along with "keys" or "namedkeys". "path" is used only to compare the measure data against a pre-existing Flat file. If none of "path","keys" and "namedkeys" are specified, the comparison occurs for all the keys in the measure array of the Measure Store (Workbook).	Optional
start	The field start value. If this does not exist in the XML, the value is taken from the measure's properties. Used in a flat-file comparison.	Optional
width	The field start value. If this does not exist in the XML, the value is taken from the measure's properties. Used in flat-file comparison.	Optional
value	The expected value of the measure. Comparison on equality.	Optional

Sample usage (indices):

```
<assert-workbook-measure-value-gt name="drwppthrsc1sp" keys="0-2" value="0.0" />
```

Sample usage (names):

```
<assert-workbook-measure-value-gt name="drwppthrsc1sp" namedkeys="pos1,pos2,pos3" value="0.0" />
```

4.3.10 assert-workbook-measure-value-le Asserts a workbook measure value is greater than or equal to a provided value.

Attribute	Description	Category
keys	The keys to be used to compare measure values. Can be either indices, or position names. Each dimension spec in the keys-spec is delimited by a ":". Multiple position ranges for a dimension can be specified by using "," as separators. When using indices, a range of positions can also be used. Ex - "0-2:4-7,11-13:4" implies indices 0,1,2 for the first dimension, 4,5,6,7,11,12,13 for the second, and 4 for the third dimension of the measure. Note: Because indices are based on the workbook selections, they may not be consistent across all test cases. Care must be exercised to ensure that you are using the correct indices for your specific test case.	Required
namedkeys	Position names used to get keys of the measure array. Using indices instead of position names, or positions that do not exist in the MeasureStore (Workbook) will be flagged as an error. Similar to "keys" above, but using position names only.	Optional
path	Path of the flat file used for comparison against the Measure array. This looks like pointing to a pre-set environment variable. "path" cannot be used along with "keys" or "namedkeys". "path" is used only when we want to compare the measure data against a pre-existing flat file. If none of "path", "keys" and "namedkeys" are specified, the comparison occurs for all the keys in the measure array of the Measure Store (Workbook).	Optional
start	The field start value. If this does not exist in the XML, the value is taken from the measure's properties. Used in a flat-file comparison.	Optional
width	The field start value. If this does not exist in the XML, the value is taken from the measure's properties. Used in flat-file comparison.	Optional
value	The expected value of the measure. Comparison on equality.	Optional

Sample usage (indices):

```
<assert-workbook-measure-value-le name="drwppthrscisp" keys="0-2" value="0.0" />
```

Sample usage (names):

```
<assert-workbook-measure-value-le name="drwppthrsc1sp" namedkeys="pos1,pos2,pos3"
value="0.0" />
```

4.3.11 assert-workbook-measure-value-lt Asserts a workbook measure value is less than a provided value.

Attribute	Description	Category
keys	The keys to be used to compare measure values. Can be either indices, or position names. Each dimension spec in the keys-spec is de-limited by a ":". Multiple position ranges for a dimension can be specified by using "," as separators. When using indices, a range of positions can also be used. Ex - "0-2:4-7,11-13:4" implies indices 0,1,2 for the first dimension, 4,5,6,7,11,12,13 for the second, and 4 for the third dimension of the measure. Note: Because indices are based on the workbook selections, they may not be consistent across all test cases. Care must be exercised to ensure that you are using the correct indices for your specific test case.	Required
namedkeys	Position names used to get keys of the measure array. Using indices instead of position names, or positions that do not exist in the MeasureStore (Workbook) will be flagged as an error. Similar to "keys" above, but using position names only.	Optional
path	Path of the flat file used for comparison against the Measure array. This looks like pointing to a pre-set environment variable. "path" cannot be used along with "keys" or "namedkeys". "path" is used only when we want to compare the measure data against a pre-existing flat file. If none of "path", "keys" and "namedkeys" are specified, the comparison occurs for all the keys in the measure array of the Measure Store (Workbook).	Optional
start	The field start value. If this does not exist in the XML, the value is taken from the measure's pProperties. Used in flat-file comparison.	Optional
width	The field start value. If this does not exist in the XML, the value is taken from the measure's properties. Used in a flat-file comparison.	Optional
value	The expected value of the measure. Comparison on equality.	Optional

Sample usage (indices):

```
<assert-workbook-measure-value-lt name="drwppthrsclsp" keys="0-2" value="0.0" />
```

Sample usage (names):

```
<assert-workbook-measure-value-lt name="drwppthrsclsp" namedkeys="pos1,pos2,pos3" value="0.0" />
```

4.3.12 assert-workbook-measure-value-ne Asserts a workbook measure value is not equal to a provided value.

Attribute	Description	Category
keys	The keys to be used to compare measure values. Can be either indices, or position names. Each dimension spec in the keys-spec is de-limited by a ":". Multiple position ranges for a dimension can be specified by using "," as separators. When using indices, a range of positions can also be used. Ex - "0-2:4-7,11-13:4" implies indices 0,1,2 for the first dimension, 4,5,6,7,11,12,13 for the second, and 4 for the third dimension of the measure. Note: Because indices are based on the workbook selections, they may not be consistent across all test cases. Care must be exercised to ensure that you are using the correct indices for your specific test case.	Required
namedkeys	Position names used to get keys of the measure array. Using indices instead of position names, or positions that do not exist in the MeasureStore (Workbook) will be flagged as an error. Similar to "keys" above, but using position names only.	Optional
path	Path of the flat file used for comparison against the Measure array. This looks like pointing to a pre-set environment variable. "path" cannot be used along with "keys" or "namedkeys". "path" is used only to compare the measure data against a pre-existing flat file. If none of "path","keys" and "namedkeys" are specified, the comparison occurs for all the keys in the measure array of the Measure Store (Workbook).	Optional
start	The field start value. If this does not exist in the XML, the value is taken from the measure's properties. Used in a flat-file comparison.	Optional
width	The field start value. If this does not exist in the XML, the value is taken from the measure's properties. Used in flat-file comparison.	Optional

Attribute	Description	Category
value	The expected value of the measure. Comparison on equality.	Optional

Sample usage (indices):

```
<assert-workbook-measure-value-ne name="drwppthrsc1sp" keys="0-2" value="0.0" />
```

Sample usage (names):

```
<assert-workbook-measure-value-ne name="drwppthrsc1sp" namedkeys="pos1,pos2,pos3" value="0.0" />
```

4.3.13 assert-workbook-popcount-eq Asserts a workbook measure popcount matches an expected value.

Attribute	Description	Category
name	Measure name	Required
pop-count	Expected popcount	Required

Sample usage:

```
<assert-workbook-popcount-eq name=" Ipopappenbb" pop-count="100" />
```

4.3.14 assert-workbook-popcount-eq-zero Asserts a workbook measure is not populated.

Attribute	Description	Category
name	Measure name	Required

Sample usage:

```
<assert-workbook-popcount-eq-zero name=" Ipopappenbb" />
```

4.3.15 assert-measure-popcount-ne-zero Asserts a workbook measure is populated.

Attribute	Description	Category
name	Measure name	Required

Sample usage:

```
<assert-workbook-popcount-ne-zero name=" Ipopappenbb" />
```

4.4 Workbook Data Edit Operations

This section provides details about workbook data edit operations.

4.4.1 clear-workbook-list Clears a list of workbook measures of all data.

Attribute	Description	Category
path	Path to a file containing a new-line delimited list of measures	Required

Sample usage:

```
<clear-workbook-list path="my_list.txt"/>
```

4.4.2 clear-workbook-measure Clears a workbook measure of all data.

Attribute	Description	Category
name	Measure name	Required

Sample usage:

```
<clear-workbook-measure name=" Ipopappenbb"/>
```

4.4.3 edit-workbook-measure Edits the value of a workbook measure based on provided attributes. This tag closely models the action of editing a cell from the Client, as opposed to set-workbook-measure, which does a direct update of the workbook arrays.

This tag can be used in a variety of ways:

Set by loadmeasure file:

By specifying <path>, <start>, and <width>, you can set the measure values according to a supplied loadmeasure file. This file must be in fixed-width format.

Set by positions:

You can set individual positions by providing a <keys> tag. If you want, you can provide the list of positions at an aggregate level by specifying the intersection using the <keyint> tag.

Set everything:

If you do not supply a <path> or <keys> tag, the entire measure will be set to the supplied value.

Attribute	Description	Category
name	Measure name	Required
keys	List of comma-separated position names	Optional
namedkeys	List of comma-separated position names	Optional
path	Path to load measure type file to compare with	Optional
keyint	Intersection string	Optional
value	Expected value	Optional
start	Starting column of measure field	Optional
width	Width of measure field	Optional

Note: For Date type measures, RPAC supports setting the array cell using either of the following two Date & Time formats - "%m/%d/%Y" or "%m/%d/%Y-%H:%M:%S" where m = 2 digit month, d = 2 digit date, Y = 4 digit year, H = 2 digit hour, M = 2 digit min, and S = 2 digit sec.

Sample usage (using keys):

```
<edit-workbook-measure name="meas1" keys="0-2:5" value="0.0" />
```

Sample usage (using keyint):

```
<assert-workbook-measure-value name="meas2" namedkeys="chnl1:week1"
value="02/24/2012" />
```

Sample usage (load-measure file):

```
<edit-workbook-measure name="meas1" path="loadMeas1.dat" start="100" width="20"/>
```

Sample usage (set all):

```
<edit-workbook-measure name="meas1" value="2"
```

4.4.4 set-workbook-measure Sets the value of a workbook measure based on provided attributes. Unlike edit-workbook-measure, this performs a direct change of the workbook data and does not update the change arrays unless told to do so. This can have an impact on the calculation cycle since bypassing the change arrays results in a full evaluation rather than incremental one.

This tag can be used in a variety of ways:

Set by loadmeasure file:

By specifying <path>, <start>, and <width>, you can set the measure values according to a supplied loadmeasure file. This file must be in fixed-width format.

Set by positions:

You can set individual positions by providing a <keys> tag. If you want, you can provide the list of positions at an aggregate level by specifying the intersection using the <keyint> tag.

Set everything:

If you do not supply a <path> or <keys> tag, the entire measure will be set to the supplied value.

Attribute	Description	Category
name	Measure name	Required
keys	List of comma-separated position names	Optional
namedkeys	List of comma-separated position names	Optional
path	Path to load measure type file to compare with	Optional
keyint	Intersection string	Optional
value	Expected value	Optional
start	Starting column of measure field	Optional
width	Width of measure field	Optional
incremental	True if evaluation is to be incremental	Optional

Note: For Date type measures, RPAC supports setting the array cell using either of the following two Date & Time formats - "%m/%d/%Y" or "%m/%d/%Y-%H:%M:%S" where m = 2 digit month, d = 2 digit date, Y = 4 digit year, H = 2 digit hour, M = 2 digit min, and S = 2 digit sec.

Sample usage (using keys):

```
<set-workbook-measure name="meas1" keys="0-2:5" value="0.0" incremental="true"/>
```

Sample usage (using keyint):

```
<set-workbook-measure-value-ne name="meas2" namedkeys="chnl1:week1"
value="02/24/2012" />
```

Sample usage (load-measure file):

```
<set-workbook-measure name="meas1" path="loadMeas1.dat" start="100" width="20"/>
```

Sample usage (set all):

```
<set-workbook-measure name="meas1" value="2"
```

5. Dynamic Position Maintenance (DPM) Tags

This section provides details about dynamic position maintenance tags.

5.1 Assertions

This section provides details about assertions.

5.1.1 assert-position-label Assert if a position does not exist.

Attribute	Description	Category
levelname	Level name	Required
position	Position name	Required
label	Position label	Required
lang	Language	Required

Sample usage:

```
<assert-position-label levelname="clss" position="4000" label="4000 Casual"
lang="english" />
```

5.1.2 assert-position-not-exists Assert if a position does not exist.

Attribute	Description	Category
levelname	Level name	Required
position	Position name	Required

Sample usage:

```
<assert-position-not-exists levelname="clss" position="4000"/>
```

5.1.3 assert-position-rollup

Assert if the given rollup is valid or not.

Attribute	Description	Category
rollup	The expected format for rollup is rollup =childLevel:childPosition, parentLevel:parentPosition	Required

Sample usage:

```
<assert-position-rollup rollup="clss:4000,dept:101" />
```

5.1.4 assert-position-status

Assert the status of a position. Status is 1 for formal, 2 for informal, 0 for placeholder, and -1 for deleted.

Attribute	Description	Category
levelname	Level name	Required
position	Position name	Required
status	Status	Required

Sample usage:

```
<assert-position-status levelname="str" position="str_dpm0000020" status="2"/>
```

5.2 Update Informal Position Operation

This section contains details about update informal position operations.

5.2.1 update-informal-position

Contains the dpm tags to add or modify dpm positions.

5.2.2 dpmposinfo

tag is the child tag of update-informal-position.

Attribute	Description	Category
lang	Language	Required

Sample usage:

```
<dpmposinfo lang="english"> </dpmposinfo>
```

5.2.3 dpmposentry

This is the child tag of dpmposinfo.

Attribute	Description	Category
levelname	Level name	Required
position	Position name	Required
isnewpos	Boolean value: True for new position and False for modify.	Required
label	Position label	Optional
oldposition	Old position name	Optional

Attribute	Description	Category
type	Integer value. Can be one of the following 0, REGULAR: one single position 1, MULTIPLE_PER_PARENT: multiple positions per parent dim 2, ONE_PER_CHILD: one position on rollup dim per child - 3, ONE_FOR_ALL_CHILDREN: one pos on rollup dim for all children 4, GEN_NAME_ONLY: only generating names/labels without applying changes	Optional
count	position count for MULTIPLE_PER_PARENT type	Optional

Sample usage:

```
<update-informal-positions>
  <dpmposinfo lang="english">
    <dpmposentry levelname="dstr" isnewpos="false" position="801" label="Pos
label" oldposition="801" type="0" count="1"/>
  </dpmposinfo>
</update-informal-positions>
```

5.3 Delete Informal Position Operation

This section provides details about the delete informal position operation.

5.3.1 delete-informal-position This tag deletes an informal position.

Attribute	Description	Category
levelname	Level name	Required
position	Position name	Required

Sample usage:

```
<delete-informal-position levelname="str" position="strdpm00000005"/>
```

6. HSA Tags

This section provides details about HSA tags.

6.1 assert-fact-group

Assert a fact exists or not in the fact group.

Attribute	Description	Category
name	Measure name	Required
factgroup	Fact group name	Required
factexists	If flag is true, then fact is checked to see if it exists in the fact group.	Required

Sample usage:

```
<assert-fact-group name="UtDyScSrc01U" factgroup="utscalar" factexists="true"/>
```

6.2 assert-fact-table

Assert a fact table exists or not.

Attribute	Description	Category
factgroup	Fact group name	Required
exists	If flag is true, then check if the fact table exists.	Required

Sample usage:

```
<assert-fact-table factgroup="utscalar" exists="true"/>
```

6.3 assert-fact-value-eq

Assert a fact value is equal to an expected value.

Attribute	Description	Category
name	Measure name	Required
keys	List of comma-separated position names	Optional
namedkeys	List of comma-separated position names	Optional
path	Path to load measure type file to compare with	Optional
value	Expected value	Optional
start	Starting column of measure field	Optional
width	Width of measure field	Optional
precision	The number of decimal places to be considered for assertion	Optional

Sample usage:

```
<assert-fact-value-eqname="UtDyDstLoU" namedkeys="20070121:30000:5" value="10"/>
```

6.4 assert-fact-value-ge

Assert a fact value is greater than or equal to an expected value.

Attribute	Description	Category
name	Measure name	Required
keys	List of comma-separated position names	Optional
namedkeys	List of comma-separated position names	Optional
path	Path to load measure type file to compare with	Optional

Attribute	Description	Category
value	Expected value	Optional
start	Starting column of measure field	Optional
width	Width of measure field	Optional
precision	The number of decimal places to be considered for assertion	Optional

Sample usage:

```
<assert-fact-value-ge name="UtDyDstLoU" namedkeys="20070121:30000:5" value="10"/>
```

6.5 assert-fact-value-gt

Assert a fact value is greater than an expected value.

Attribute	Description	Category
name	Measure name	Required
keys	List of comma-separated position names	Optional
namedkeys	List of comma-separated position names	Optional
path	Path to load measure type file to compare with	Optional
value	Expected value	Optional
start	Starting column of measure field	Optional
width	Width of measure field	Optional
precision	The number of decimal places to be considered for assertion	Optional

Sample usage:

```
<assert-fact-value-gt name="UtDyDstLoU" namedkeys="20070121:30000:5" value="10"/>
```

6.6 assert-fact-value-le

Assert a fact value is less than or equal to an expected value.

Attribute	Description	Category
name	Measure name	Required
keys	List of comma-separated position names	Optional
namedkeys	List of comma-separated position names	Optional
path	Path to load measure type file to compare with	Optional
value	Expected value	Optional
start	Starting column of measure field	Optional
width	Width of measure field	Optional

Attribute	Description	Category
precision	The number of decimal places to be considered for assertion	Optional

Sample usage:

```
<assert-fact-value-le name="UtDyDstLoU" namedkeys="20070121:30000:5" value="10"/>
```

6.7 assert-fact-value-lt

Assert a fact value is less than an expected value.

Attribute	Description	Category
name	Measure name	Required
keys	List of comma-separated position names	Optional
namedkeys	List of comma-separated position names	Optional
path	Path to load measure type file to compare with	Optional
value	Expected value	Optional
start	Starting column of measure field	Optional
width	Width of measure field	Optional
precision	The number of decimal places to be considered for assertion	Optional

Sample usage:

```
<assert-fact-value-lt name="UtDyDstLoU" namedkeys="20070121:30000:5" value="10"/>
```

6.8 assert-fact-value-ne

Assert a fact value is not equal to an expected value.

Attribute	Description	Category
name	Measure name	Required
keys	List of comma-separated position names	Optional
namedkeys	List of comma-separated position names	Optional
path	Path to load measure type file to compare with	Optional
value	Expected value	Optional
start	Starting column of measure field	Optional
width	Width of measure field	Optional
precision	The number of decimal places to be considered for assertion	Optional

Sample usage:

```
<assert-fact-value-ne name="UtDyDstLoU" namedkeys="20070121:30000:5" value="10"/>
```

6.9 assert-level-contain

Assert a level contains or does not contain a position name.

Attribute	Description	Category
name	Level name	Required
positions	List of comma-separated position names	Required
exists	If flag is true, then position names are checked to see if they are contained in the level.	Required

Sample usage:

```
<assert-level-contain name="skup" positions="6200006" exists="T"/>
```

6.10 clear-rdm-status

Clears the RDM status code.

Attribute	Description	Category
status_code	RDM status code	Required
domain_id	Domain ID	Required
info	Information	Required

Sample usage:

```
<clear-rdm-status status_code="1004" domain_id="" info=""/>
```

6.11 set-rdm-status

Sets the RDM status code.

Attribute	Description	Category
status_code	RDM status code	Required
domain_id	Domain ID	Required
info	Information	Required

Sample usage:

```
<set-rdm-status status_code="2001" domain_id="0" info="scls"/>
```

6.12 set-fact-value

Set the value of a fact based on provided attributes.

Attribute	Description	Category
name	Measure name	Required
keys	List of comma-separated position names	Optional

Attribute	Description	Category
namedkeys	List of comma-separated position names	Optional
path	Path to load measure type file to compare with	Optional
keyint	Intersection string	Optional
value	Value	Optional
start	Starting column of measure field	Optional
Width	Width of measure field	Optional

Sample usage:

```
<set-fact-value name="UtDyDst02U" namedkeys="w01_2007:30018:5" value="0"/>
```

6.13 reset-seq-generator

Reset a sequence generator to a specific value.

Attribute	Description	Category
name	Name of the generator.	Required
central	Indicates whether the reset is for centralized generator. True/False.	Optional
value	Reset to this value. Any integer ≥ 0 . Default 0.	Optional

Sample usage:

```
<reset-seq-generator name="dpm_str" value="1000"/>
```

6.14 assert-rdm-hier-exists

Assert a hierarchy exists in RDM or not.

Attribute	Description	Category
name	Hierarchy name	Required
exists	If flag is true, then check if hierarchy exist in RDM.	Required

Sample usage:

```
<assert-rdm-hier-exists name="clnd" exists="true"/>
```

6.15 assert-level-exists

Assert a level exists in RDM or not.

Attribute	Description	Category
name	Level name	Required
exists	If flag is true, then check if hierarchy exist in RDM.	Required

Sample usage:

```
<assert-level-exists name="day" exists="true"/>
```

Appendix: Environment Variables

RPAS includes a number of environment variables that are set at the system level in UNIX. At the system level, the variables are applicable to all RPAS Servers (DomainDaemons) that are run on the system.

The common syntax for setting these variables is as follows:

```
export ENVIRONMENT_VARIABLE=XXXXXX
```

ENVIRONMENT_VARIABLE is a defined variable that is recognized by RPAS. XXXXXX is an appropriate value for the variable, which could be a string, Boolean value, numeric value, or date and time. If the value represents time, this number normally represents time in milliseconds.

Note: The DomainDaemon must be restarted after setting any environment variables. An example of how this process is completed is as follows:

```
DomainDaemon -port 55123 -start -debug &
```

This appendix describes the following environment variables used for RPAS:

- [Required Settings](#)
- [Database Settings:](#)
 - [RPAS_ARRAY_TAG](#)
 - [RPAS_CACHE](#)
 - [RPAS_PAGE_SIZE](#)
 - [RPAS_PAGE_SPLIT_PERCENTAGE](#)
 - [RPAS_LOCK_TIMEOUT](#)
 - [RPAS_USER_MODE](#)
- [Java Special Expression Settings:](#)
 - [RPAS_JAVA_CLASSPATH](#)
 - [RPAS_JAVA_DEBUG_PORT](#)
 - [RPAS_JAVA_MAXHEAPSIZE](#)
- [Log Settings:](#)
 - [RPAS_CACHE_LM_MULTIPILER / RPAS_DISABLE_LM_DIRECT](#)
 - [RPAS_CLIENT_TIMEOUT / RPAS_SERVER_TIMEOUT](#)

-
- RPAS_COMMIT_FAILED_MAX / RPAS_COMMIT_SUCCESS_MAX
 - RPAS_JAVA_MAXHEAPSIZE
 - RPAS_LOCK_DETAIL_PATH
 - RPAS_LOCK_LOGGING
 - RPAS_LOCK_LOG_PATH
 - RPAS_LOCK_TRACK_DATABASES
 - RPAS_LOCK_ACCESS_LIMIT
 - RPAS_LOG_BACKUPS
 - RPAS_LOG_FD
 - RPAS_LOG_FLUSH
 - RPAS_LOG_LEVEL
 - RPAS_LOG_PATH
 - RPAS_REQUEST_TIMEOUT
 - RPAS_SHELL
 - RPAS_SPLIT_BUF_MB
 - RPAS_STATS_PATH
 - RPAS_TCPNODELAY
 - RPAS_TRACE_SSL
 - TMPDIR
 - TODAY_RETEK
 - Parallelization Settings:
 - RPAS_EXPORT_MEASURES_MAX
 - RPAS_MIN_MEAS_PER_PROC
 - RPAS_PROCESSES
 - Profiling Logging:
 - RPAS_PROFILING_ENABLE
 - RPAS_PROFILING_PATH
 - Date and Time Setting:
 - RPAS_TODAY
 - RPAS_TODAY_STATIC
 - Numeric Precision:
 - RPAS_INCAGGPRC
 - epsilon
 - RPAS_SSL_WALLET
 - LDR_CNTRL=NAMEDSHLIB=RPASZONE
 - Nagle's Algorithm

Required Settings

The following environment variables are required.

- RPAS_HOME
- RIDE_HOME
- RIDE_OPTIONS
- JAVA_HOME
- RPAS_JAVA_CLASSPATH
- PATH

Since these environment variables are required at the time of installation, and since they vary according to the operating system that RPAS is running on, these environment variables are described in the "Installing on UNIX and Linux Environments" chapter of the *Oracle Retail Predictive Application Server Installation Guide*. See that chapter for detailed instructions for those environment variables.

Database Settings

These variables are used for RPAS B-tree storage performance.

RPAS_ARRAY_TAG

RPAS BTree arrays contain a field named array tag. Whenever a BTree array is written, the value of the RPAS_ARRAY_TAG environment property, if it is set, is written as the value of the array tag field. By using this property, users can set a record in an array to provide information about the last time the array was modified.

RPAS_CACHE

RPAS_CACHE determines the number of pages of a BTree array file that are contain in memory. This variable can be set to have values between 16 and 512 (including 16 and 512). The larger you set the cache value, the more memory usage and less disk access you have.

If you set the cache size from 4 to 64, the performance of the BTree arrays should level out around 16 pages. In tests, more pages did not provide better performance. However, for other access patterns, more cache pages may improve performance. Therefore, a larger cache and page size may increase performance due to less disk access. However, having a larger cache and page size increases the memory image of an open BTree file. It is likely that if the number of cache elements is 512 and the page size is 256K, then each open BTree file could use up to 131,072K of memory.

```
export RPAS_CACHE=16
```

Values: 4, 8, 16, 32, 64, 128, 256, 512

RPAS_PAGE_SIZE

RPAS_PAGE_SIZE determines the size of a single BTree array page if the logical size of the BTree array is greater than 200,000 cells. A BTree array is compose of a number of pages and page types. Each page of a BTree array file is the same size. In the trunk, there are five pages types: header, branch, leaf, data, and free. A larger page size implies more memory usage, less fragmentation, and less disk access.

The page size also affects the disk overhead that each BTree file requires. A large page size may provide better disk usage if the array is densely packed. However, large page size with a loosely packed array uses more disk space. For example, consider the effect of having 1,000,000 cells with only one cell populated on a 256K page versus one cell populated on a 16K page. The first case requires 256,000,000K of disk space, while the second requires 16,000,000K of disk space.

Therefore, `RPAS_PAGE_SIZE` and `RPAS_CACHE` both affect performance based on the access pattern. `RPAS_PAGE_SIZE` affects disk space and memory usage, and `RPAS_CACHE` affects the memory usage.

```
export RPAS_PAGE_SIZE=32K
```

Values: 4K, 8K, 16K, 32K, 64K, 128K, 256K

RPAS_PAGE_SPLIT_PERCENTAGE

This variable sets the page split percentage. The page split percentage determines the percentage of low order keys to keep in the low order page of the B-Tree when a full page is being split to accommodate data for a new key. It is expressed as a number between 1 and 100. RPAS stores non-NA values only. If an operation (load or calculation) causes a value to be non-NA, RPAS tries to store it on a page and, if the page is full, it causes it to split. Page splitting is an expensive operation; that is, it takes a significant amount of processing time and if it happens too often, it can significantly slow down the load or calculation process causing the splits. For efficient storage and to prevent pages from being split very often, it is recommended that the value be kept between 50 and 90. If this environment variable is not set, RPAS uses a page split percentage of 90.

RPAS_LOCK_TIMEOUT

This environment variable sets the maximum number of milliseconds to wait for a database lock. The default value, 60,000, forces a database lock to wait for a maximum of one minute before throwing a database lock exception.

```
export RPAS_LOCK_TIMEOUT=60000
```

Values: 60000, 90000, 120000, and so on

When performing certain operations, it is possible for two or more users to be contending for access to the same database, which happens most commonly when two users attempt to simultaneously commit and save the same data back to the domain. This can also occur when running RPAS utilities in parallel to prevent database locking during batch operations. By default, RPAS is set up to wait one minute before returning a lock contention error when this situation occurs.

If desired, an administrator can override this default value by setting the `RPAS_LOCK_TIMEOUT` environment variable. This variable is set to the number of milliseconds to wait for a file lock before returning a lock contention error. As with any environment variable, the variable must be set prior to starting the process that uses that variable. The variable was introduced for use with the RPAS database server, which means that the variable is set for the `DomainDaemon`.

For example, the line below indicates how an administrator would tell RPAS to wait two minutes before returning a lock contention error with the `RpasDbServer` after launching the client and logging in. Any client that connects to that domain daemon would see lock contention after a two minute delay:

```
Export RPAS_LOCK_TIMEOUT=120000
```

On the client side, the exception sent from the server appears as a warning dialog box only after some new action is initiated. With the client, additional warning dialog boxes are displayed to the user.

RPAS_REQUEST_TIMEOUT

The RPAS_REQUEST_TIMEOUT environment variable is also used to handle issues with firewalls and the RpasDbServer. The RpasDbServer checks RPAS_REQUEST_TIMEOUT to determine what should happen when it has been idle for a period of time. Like any environment variable that RpasDbServer uses, this environment variable must be set prior to starting the DomainDaemon.

The environment variable RPAS_REQUEST_TIMEOUT should be set to a value that is the number of seconds of idle time that should pass before the RpasDbServer sends a “Server has timeout waiting for a request” exception to the client and exits. In this case, idle time is the time waiting for a request. If this variable is not present or is set to zero, then the RpasDbServer will never time out.

```
export RPAS_REQUEST_TIMEOUT=600
```

Values: 600, 900, 1200, and so on

RPAS_USER_MODE

This environment variable determines how the DbCloser class behaves. It controls whether the DbCloser aggressively closes databases. The “single” value puts the DbCloser in single user mode. This is useful for running batch processes when no users are logged in, because in such a situation there is no need to rapidly close databases.

```
export RPAS_USER_MODE="single"
```

Value: single

Java Special Expression Settings

The following environment variables can be used for Java special expression settings.

RPAS_JAVA_CLASSPATH

When using the Java Special Expression functionality, the RPAS_JAVA_CLASSPATH environment property contains the location of the Java expressions. The value of RPAS_JAVA_CLASSPATH should be a set of paths to the jar files that contain the compiled Java expressions. The list of paths must also contain the path to the jar file oracleRpasUtils.jar located within the lib subdirectory of RPAS_HOME.

RPAS_JAVA_DEBUG_PORT

The RPAS_JAVA_DEBUG_PORT environment variable can be used to specify to RPAS what port the JVM, created to execute a Java Special Expression, should listen to in order for a debugger to attach to the JVM. It can be used to debug Java Special Expressions during their development. The value of the RPAS_JAVA_DEBUG_PORT property should be an integer within the range of available port addresses.

RPAS_JAVA_MAXHEAPSIZE

The RPAS_JAVA_MAXHEAPSIZE environment property can be used to specify to RPAS that the JVM, created to execute a Java Special Expression, should request more memory from the OS than the default. The value of RPAS_JAVA_MAXHEAPSIZE should be the number of megabytes of memory to request followed by the letter 'm' (for example, 1024m).

Log Settings

The following environment variables can be used for log settings.

RPAS_CACHE_LM_MULTIPILER / RPAS_DISABLE_LM_DIRECT

The RPAS loadmeasure application uses these environmental variables. The RPAS_DISABLE_LM_DIRECT environmental variable may be "yes" or "1" to force the loadmeasure application to run in legacy mode which disables the use of the RPAS_CACHE_LM_MULTIPILER variable.

The RPAS_CACHE_LM_MULTIPILER variable must be between 2 and 10 and is used to compute the amount of data the loadmeasure application will cache base on its value and the value of the RPAS_CACHE variable.

RPAS_CLIENT_TIMEOUT / RPAS_SERVER_TIMEOUT

These environmental variables determine the amount of time a TCP/IP read or write call waits for data to be processed. If no data has been processed in the amount of time provided, RPAS throws an exception.

RPAS_COMMIT_FAILED_MAX / RPAS_COMMIT_SUCCESS_MAX

These environmental variables determine the number of failed or success entries that are kept in the success or failed directories for commit ASAP.

RPAS_JAVA_MAXHEAPSIZE

This environment parameter, if set, applies only to the RPAS Java Special Expression functionality, and is passed along as the -Xmx argument to the Java JVM instance.

RPAS_LOCK_DETAIL_PATH

This environment property gives a directory path to which detailed file lock history data will be written, in binary form, using one log file per RPAS process. The lock detail files, which will not be human readable, are constructed to be extremely compact so that a large amount of file history can be written efficiently. These details consist only of the file's inode (internal operating system identifier for the file), the operation (open for read, open for write, close), and a date/time code. The utilities for analyzing this information are not currently a part of the official RPAS release, so the data would need to be sent to the RPAS Development staff for analysis. Because of this, the environment variable should only be set when requested by RPAS Development or Support in cases of an ongoing support need related to file lock contention.

RPAS_LOCK_LOGGING

This environment property causes RPAS binaries to log details about the locking and unlocking of files. The value for RPAS_LOCK_LOGGING should be set to the logger level (error, warning, debug, profile, information) at which the file lock details will be output. So, for instance, with RPAS_LOCK_LOGGING=debug, if you run an RPAS binary with -loglevel debug, then you will see the file lock details, but if you run the binary with -loglevel warning, you will not see them.

RPAS_LOCK_LOG_PATH

This environment property gives a file path to which details of file locks will be written, in human-readable text form. The details match those which are output to the screen when RPAS_LOCK_LOGGING is set to an appropriate level.

RPAS_LOCK_TRACK_DATABASES

This environmental variable will filter the logged file lock information in order to limit it to one or more databases of interest. It should be set to a comma-separated list of database names, for example "admin,measdata". This would cause lock information to be logged only for files in the admin.db and measdata.db databases.

This setting works in conjunction with these other environment variable settings:

- RPAS_LOCK_LOGGING
- RPAS_LOCK_ACCESS_LIMIT
- RPAS_LOCK_LOG_PATH
- RPAS_LOCK_TRACK_PATH

RPAS_LOCK_ACCESS_LIMIT

Set this environmental variable to the maximum number of seconds a lock file can be locked before the name of the lock file is logged. The logging occurs at the level provided by the environmental variable RPAS_LOCK_LOGGING. This allows you to filter out brief lock times from the log, showing only those files which are locked for a longer time.

RPAS_LOG_BACKUPS

The RPAS_LOG_BACKUPS environment variable sets the maximum number of log files to keep.

```
export RPAS_LOG_BACKUPS=20
```

Values: 2, 3, 4, 10, 20, 30, and so on

The RPAS_LOG_BACKUPS environment variables allow an administrator to define the number of log file backups to retain for a given user. A log file is created each time for each session that a user has with the RPAS Client.

The environment variable is set by executing the following command:

```
Export RPAS_LOG_BACKUPS=X
```

X is an integer value that represents the number of backup log files to keep for each user.

RPAS_LOG_FD

This environment property causes RPAS binaries to log the opening and closing of all file handles. The value for RPAS_LOG_FD should be set to the logger level (error, warning, debug, profile, information) at which the file open/close details will be output. So, for instance, with RPAS_LOG_FD=debug, if you run an RPAS binary with -loglevel debug, then you will see the file lock details, but if you run the binary with -loglevel warning, you will not see them.

RPAS_LOG_FLUSH

Setting this environmental variable to anything other than empty or the string "false" (case-insensitive) causes the RPAS logger to flush the output after each line of output. This may cause a tiny performance impact, but ensures that all logger output has been sent before any error or crash condition. It should generally only be needed in cases of debugging.

RPAS_LOG_LEVEL

The RPAS_LOG_LEVEL environment variable sets the default logging level to use when an application starts.

```
export RPAS_LOG_LEVEL="debug"
```

Values: all, profile, debug, audit, information, warning, error, none

RPAS_LOG_PATH

The RPAS_LOG_PATH environment variable contains the path of a directory where RPAS daemon applications creates log files. The default is the current working directory where the application was started from.

```
export RPAS_LOG_PATH=" C:/RPAS/Domains/Temp"
```

Value: The absolute path to the directory where RPAS daemon applications creates log files

RPAS_REQUEST_TIMEOUT

This environmental variable specifies an amount of time (in seconds) that a TCP/IP read or write call will wait for data. This functions similarly to the RPAS_SERVER_TIMEOUT environment setting, however, this variable only applies to the RpasDbServer application. The intent of this environmental variable is to handle the issue of fire walls. A firewall may be setup to kill any TCP/IP connection that is idle for 60 minutes. If someone using the client has connected to an RpasDbServer behind a fire wall with this setup, they could run into this issue. If the RPAS Client user is connected, but leaves the connection idle for 60 minutes (for example, over lunch hour), the fire wall would then kill the connection. The Client appears connected, but will not be functional. If the user attempts to reconnect to the server, the RPAS DomainDaemon will not allow the connection since the earlier RpasDbServer process is still running. Setting the RPAS_REQUEST_TIMEOUT to 30 minutes (or 1800 seconds) using the export command before starting the DomainDaemon:

```
export RPAS_REQUEST_TIMEOUT=1800
```

Now the RpasDbServer will timeout after 30 minutes of idle TCP/IP time. The Client user will then be able to reconnect to the server since the previous session will have properly disconnected.

RPAS_SHELL

This environment setting applies only to running the RPAS Server on Windows in the Cygwin environment, and is needed only if your Cygwin installation is in a non-standard directory location. RPAS Server code will expect to find the Cygwin BASH shell executable at the path C:/cygwin/bin/bash.exe. If your Cygwin BASH shell is located at a different path, set the RPAS_SHELL environment variable to the correct location.

RPAS_SPLIT_BUF_MB

This environment parameter optionally adjusts the size (in megabytes) of the file buffer used in the internal measureFileSplitter utility. If needed, it should be set to a positive integer value; if not present, the default of 8 is used.

RPAS_STATS_PATH

This environmental variable is used to determine where to write the stats information gathered during a workbook build.

RPAS_TCPNODELAY

The RPAS Client/Server communications use a TCP/IP interface. Built in to most TCP/IP implementations is an efficiency improvement known as Nagle's Algorithm. A complete discussion is beyond the scope of this document (see http://en.wikipedia.org/wiki/Nagle's_algorithm), but it relates to added efficiency when sending many small packets of data by grouping them. However, this can cause delays in certain communication patterns, an issue that some customers detect when using RPAS. To disable Nagle's Algorithm, you may set environment variable RPAS_TCPNODELAY to "on". Conversely, to enable Nagle's Algorithm, the value of RPAS_TCPNODELAY may be set to "off". If the environment variable is unset, or set to anything besides on or off, the TCP/IP communication behavior is determined by your operating system settings.

RPAS_TRACE_SSL

This environmental variable is used to create a log file of the SSL handshake. This variable must be set to two values with a comma separator. The first value must be one of these strings: "fatal", "error", "warn", "user", "debug", "ssl", or "memory". The second value can be either an RPAS loglevel string such as "basic", "none", "error", "warning", "information", "audit", "profile", "debug", or "all"; or an absolute path to a log file.

To set the ssl trace to "user" level and log to the file "/home/larry/ssltrace.log" use:

```
Export RPAS_TRACE_SSL=user,/home/larry/sslTrace.log
```

To set the ssl trace to "debug" and log the output to the RPAS logging level "debug" use:

```
Export RPAS_TRACE_SSL=debug,debug
```

TMPDIR

The RPAS Server sometimes needs to create temporary files. By default, these are placed into the /tmp directory. If the RPAS administrator wishes to set a different location for temporary files, the TMPDIR environment variable should be set to the full path to this directory.

TODAY_RETEK

See [RPAS_TODAY](#) in the Date and Time Setting section.

Parallelization Settings

The following environment variables can be used for parallelization settings.

RPAS_EXPORT_MEASURES_MAX

The RPAS command-line utility "exportMeasure" runs sub-processes in order to support greater efficiency by parallel processing. A positive integer value assigned to RPAS_EXPORT_MEASURES_MAX sets an upper limit on the number of measures to be handled by each invocation of the sub-process. If the environment variable is not set, the default value of 100 is used.

RPAS_MIN_MEAS_PER_PROC

The RPAS command-line utility "loadmeasure" runs sub-processes in order to support greater efficiency by parallel processing. A positive integer value assigned to RPAS_MIN_MEAS_PER_PROC sets a lower limit on the number of measures being loaded in each sub-process. If the environment variable is not set, the default value of 100 is used.

RPAS_PROCESSES

This variable is used when performing batch measure registration. It sets the maximum number of processes used to stage data. The default is 1, which causes all processes to be run serially.

If a user specifies a value less than 1, it will override to 1. If a user specifies a value greater than the number of subdomains, it will be override to the number of subdomains in the global domain. If this variable is set to a number greater than 1 when doing a measure registration in a global domain, RPAS spawns that number of parallel processes to register or unregister measures in each local domain.

```
export RPAS_PROCESSES=4
```

Values: 1, 2, 3, 4, and so on

Profiling Logging

The following two environment variables may be set up to control profiling logging:

RPAS_PROFILING_ENABLE

The RPAS_PROFILING_ENABLE environment variable, when set to true, allows profiling data to be written to the profiling log file. This flag does not affect writing to a general RPAS log file, which is controlled solely by loglevel.

RPAS_PROFILING_PATH

The RPAS_PROFILING_PATH environment variable defaults to rpaProfile.log if not present. This variable specifies the profiling log file name. It can be overridden programmatically in the constructor of the profiling timer.

Date and Time Setting

The following environment variables can be used to set the date and time.

RPAS_TODAY

RPAS_TODAY tells RPAS what day that it should think today is.

```
RPAS_TODAY=20090530
```

Note: TODAY_RETEK is an older, deprecated synonym. It is allowed for historical reasons, but RPAS_TODAY is preferred.

RPAS_TODAY_STATIC

The RPAS_TODAY_STATIC environment variable affects how date and time and the RPAS_TODAY environment variable are handled. The setting of RPAS_TODAY_STATIC affects the date and time that is returned by the `DateTime::now()` function.

- If RPAS_TODAY_STATIC is set to true, all calls to `DateTime::now()` return the same date and time as the first call made to the `DateTime::now()` function. The same date and time is returned no matter how many times the `DateTime::now()` function is called.

For example, RPAS_TODAY_STATIC is set to true and RPAS_TODAY=20090530. If a call is made to `DateTime::now()` at 10:30 pm, it returns a date and time of 5/30/2009 10:30:00 pm. If another call is made 10 minutes later to `DateTime::now()`, a date and time of 5/30/2009 10:30:00 pm is also returned.

- Set the HHMMSS portion of RPAS_TODAY (YYYYMMDDHHMMSS) when setting the RPAS_TODAY_STATIC to TRUE. The function `DateTime::now()` or keyword 'now' also include the time of the day in its value. If RPAS_TODAY_STATIC is set to TRUE and if YYYYMMDDHHMMSS format is not used but YYYYMMDD format is used while setting RPAS_TODAY, then the day portion of 'now' will remain the same but the HHMMSS portion of 'now' changes.
- If RPAS_TODAY_STATIC is set to false, a call to `DateTime::now()` returns the current date and time.

For example, RPAS_TODAY_STATIC is set to false and RPAS_TODAY=20090530. If a call is made to `DateTime::now()` at 10:30 pm, it returns a date and time of 5/30/2009 10:30:00 pm. If another call is made 10 minutes later to `DateTime::now()`, a date and time of 5/30/2009 10:40:00 pm is returned.

- User level environment variable TZ (time zone) needs to be set for the user who starts the RPAS server processes (DomainDaemon, RPAS ODBC Server) or runs RPAS utilities (like `printMeasure`). A missing TZ causes the system to misinterpret the date value stored in RPAS. For example, on a Linux system, the default DATE type measure NA value of 0001/01/01 would be interpreted as 7295/12/31 23:00:00.000 if TZ were not set.

If the user is in time zone EST5EDT, use the command `export TZ=EST5EDT` to set TZ. It is suggested that the command be added to user's login script (for example, .profile).

Numeric Precision

Two variables affect the level of numeric precision that is displayed in the RPAS clients:

- `RPAS_INCAGGPRC`: An environment variable.
- `epsilon`: The smallest difference that is allowed between any two numbers.

RPAS_INCAGGPRC

Use the `RPAS_INCAGGPRC` environment variable to set how a number's precision is displayed. Set this variable to the precision level you want displayed. The default value of this variable is 1e-06. Unless this variable is altered, very small values, such as 0.00000001, are displayed as 0.

Following is an example of setting this variable to a value smaller than the default:

```
export RPAS_INCAGGPRC=0.0000000001
```

epsilon

The `epsilon` variable also affects the level of precision that is displayed. This variable is returned as a value by a public static function in the `MathUtilities.h` file. It is hardcoded to 0.0000000001 (1e-09).

RPAS uses the `epsilon` value to consider any two numbers as different if the difference is greater than or equal to `epsilon`. If the difference between any two numbers is less than `epsilon`, then those numbers are considered as equal. Indirectly, the `epsilon` value affects the display as it sets a limit to the number precision.

Note: Reducing the `epsilon` value may impact performance.

RPAS_SSL_WALLET

This variable may be used as an alternative to the `-wallet` argument to `DomainDaemon`. If the `RPAS_SSL_WALLET` environment variable is defined, its value can be used by `DomainDaemon` in place of specifying the information as a parameter to the `-wallet` argument.

Please refer to the "DomainDaemon" section of the guide for more information on the use of the `-wallet` argument.

LDR_CNTRL=NAMEDSHLIB=RPASZONE

This environment setting performs two functions:

- It creates a special named zone in the memory to load shared libraries for RPAS processes to avoid version conflict.

By default, AIX loads shared libraries in a single global zone. This could potentially cause conflict between the libraries on the system and RPAS versions. For example, RPAS has its own version of `libz`. If a non-RPAS program has loaded a different version of `libz` from the system into the global zone, RPAS will use that one instead of loading its own.
- It may also reduce the memory footprint for shared libraries in the private memory of RPAS processes.

The global zone can be filled up if the system is very busy. In that case, the shared libraries are loaded into private space of RPAS processes, which would reduce the amount of memory available for RPAS data. A named zone is difficult to fill up since it is used only by RPAS processes.

```
export LDR_CNTRL=NAMEDSHLIB=RPASZONE
```

Nagle's Algorithm

Nagle's algorithm is used to improve the efficiency of TCP/IP networks. The algorithm combines a group of small outgoing messages and sends them all at once. This reduces the number of packets that are sent over the network.

Two settings are used for Nagle's algorithm, `RPAS_TCPNODELAY` for the Server and `FC_TCPNODELAY` for the Fusion Client.

Server

Setting `RPAS_TCPNODELAY` to **on** disables Nagle's algorithm on the server side when it sends the data back to the client. This setting is useful for improving the performance of the Fusion Client in some cases.

Setting `RPAS_TCPNODELAY` to **off** or not setting it enables Nagle's algorithm.

Fusion Client

Setting `FC_TCPNODELAY` to **true** disables Nagle's algorithm on the Fusion Client side when it sends the data from the client to the server. This setting is useful in improving the performance of the Fusion Client in some cases.

Setting `FC_TCPNODELAY` to **false** or not setting it enables Nagle's algorithm.

