

Oracle® Retail Predictive Application Server

Administration Guide for the Classic Client

Release 13.3

January 2012

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

Primary Author: Barrett Gaines

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

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

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

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

This software and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. 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.

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	xvii
Preface	xix
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
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	2-2
Get the Input Files Ready	2-3
Building the Domain Manually or Through 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
Updating Domain Content	2-7
Upgrade Configurations	2-8
Command-line Support for Configuration Upgrading	2-10
Upgrade Application Configurations	2-10
convertDomain	2-10
convertDomain Usage	2-11

upgradeDomain	2-12
upgradeDomain Usage	2-12

3 Domain Administration

Configuring the Classic Client: eConfigure.....	3-1
eConfigure Menu Bar	3-1
eConfigure Main View	3-2
The Connection Group	3-2
The Domains Group	3-2
Advanced Settings Dialog	3-3
Domain Daemon	3-4
DomainDaemon Usage	3-4
Starting the DomainDaemon.....	3-5
Monitoring the DomainDaemon.....	3-5
Stopping the DomainDaemon.....	3-5
Losing a Client-Server Connection.....	3-6
Graceful Termination of the Existing Session	3-6
Autosave of Workbooks.....	3-6

4 User Maintenance

Access the User Administration Tab.....	4-2
Add User	4-3
Add User Group	4-5
Delete User	4-6
Delete User Group	4-7
Edit User.....	4-8
Managing Users: usermgr Utility.....	4-10
usermgr Usage.....	4-10
XML Schema	4-12
Use Cases.....	4-14
Exporting from an Existing Domain	4-14
Importing into a Domain	4-14
Converting Between XML and Database	4-15

5 System Administration

Auto Workbook Maintenance Wizard.....	5-2
Add a Workbook to the Auto Build Queue	5-2
Delete a Workbook from the Auto Build Queue	5-4
Edit Workbook Settings	5-4
Workbook Batch Category Management.....	5-5
Workbook Batch Category Management Wizard	5-5
Adding a Category.....	5-6
Deleting a Category	5-7
Changing a Category Label	5-8
Hierarchy Maintenance Workbook.....	5-9
Hierarchy Maintenance Example	5-9

Hierarchy Maintenance Wizard.....	5-10
Hierarchy Maintenance Worksheet.....	5-10
Access Hierarchy Maintenance.....	5-12
Maintain a User-Defined Dimension within a Hierarchy.....	5-12
Password Policy Administration Workbook.....	5-13
Security Administration Workbook.....	5-15
Security Overview.....	5-15
User Logon Security.....	5-15
Measure Level Security.....	5-15
Position Level Security.....	5-16
Workbook Security.....	5-17
Security Administration Workbook.....	5-18
Workbook Template Rights Worksheet.....	5-19
Workbook Template Measure Rights Worksheet.....	5-19
Measure Rights Worksheet.....	5-20
Dimension Modification Rights Worksheet.....	5-20
Position Level Security Worksheets.....	5-21
Workbook Template Limits Worksheets.....	5-21
Max Domain Session Limit Worksheet.....	5-21
Max User Session Limit Worksheet.....	5-21
Using the Security Administration Workbook.....	5-22
Access Security Administration.....	5-22
Set or Modify User Access to Workbook Templates.....	5-22
Set Measure Availability for Workbook Templates.....	5-22
Assign or Restrict User Access to Measures.....	5-23
Change User Ability to Modify Dimensions.....	5-23
Set or Modify Access to Positions.....	5-23
Limit the Number of Workbooks that a User Can Save.....	5-24
Set or Modify the Maximum Domain Session Limit.....	5-24
Set or Modify the Maximum User Session Limit.....	5-24
Measure Analysis Workbook.....	5-25
Building a Measure Analysis Workbook.....	5-25
Measure Analysis Worksheet.....	5-26
Review and Edit Sales or Other Registered Measure Data.....	5-26

6 Hierarchy Management

Loading Hierarchies: loadHier.....	6-1
loadHier Usage.....	6-3
loadHier Notes.....	6-4
Position Label Translation.....	6-5
Integer Indexing.....	6-6
Changing the Bit Size.....	6-8
Reindexing Domains: reindexDomain.....	6-9
reindexDomain Usage.....	6-10
ReindexDomain Usage: Reindex.....	6-10
ReindexDomain Usage: Analysis.....	6-11
ReindexDomain Usage: Domain Properties.....	6-11

When to Reindex	6-12
Run Reports	6-14
How to Reindex a Dimension	6-16
Reindex Entire Domain as Needed	6-16
Reindex Entire Domain with Force Option.....	6-16
Reindex Dimensions.....	6-16
Prepending Calendar Dimensions	6-16
Reindex Domain Examples.....	6-18
Condition 1: Change of Bit Size	6-18
Condition 2: Using the -force Option.....	6-18
Condition 3: Fragmentation Checks.....	6-19
Optimizing Domains: optimizeDomain.....	6-20
optimizeDomain Usage.....	6-21
Usage Examples	6-22
Adding New Dimensions to Hierarchies	6-23
Exporting Measure Data	6-23
Exporting Hierarchy Data.....	6-24
Purging Hierarchy Data	6-24
Adding New Dimensions: hierarchyMgr.....	6-24
Reloading Formal Hierarchy Data.....	6-25
Informalizing DPM Positions.....	6-25
Reloading Measure Data	6-25
Adding New Hierarchy Dimensions Sample Script.....	6-26
Filtering Hierarchies: filterHier.....	6-28
filterHier Usage	6-28
filterHier Notes.....	6-29
Error Conditions.....	6-29
Position Repartitioning.....	6-30
Loading RDF and Curve Parameters after Repartitioning	6-30
Syntax	6-30
Reconfiguring Partitions of a Global Domain: reconfigGlobalDomainPartitions.....	6-31
reconfigGlobalDomainPartitions Usage	6-32
Using an Input File.....	6-33
Notes, Assumptions, and Limitations.....	6-34
Renaming Positions: renamePositions.....	6-34
renamePositions Usage	6-34
Setting Properties for Dimensions: dimensionMgr	6-36
dimensionMgr Usage	6-36
Changing the Label of a Dimension	6-37
Exporting Hierarchy Data: exportHier	6-38
exportHier Usage	6-38
Informal Position Manager: informalPositionMgr	6-39
InformalPositionMgr Usage	6-39
Basic Operations: Informalize, Formalize, Remove	6-39
Error Handling	6-40
Create Informal Positions in Bulk.....	6-41
Example 1: Create Informal Positions on the STYL Dimension.....	6-43

Example 2: Create Informal Positions on the SZ12 Dimension.....	6-43
Position Naming Convention.....	6-44
Data Slice Copying.....	6-45
Merge Informal Positions to Formal Positions	6-46
Managing Position Lists as PQDs: pqdMgr.....	6-47
pqdMgr Usage	6-48
pqdMgr Notes.....	6-48

7 Data Management

Loading Measure Data: loadmeasure	7-1
Load File Names and Load Behavior	7-2
Loading Multiple Measures from One File	7-4
CSV Files	7-4
Fixed Width Files	7-4
Loading Data from Below the Base Intersection of the Measures	7-4
Staging Measure Loads	7-5
Running Pre-Load or Post-Load Scripts	7-5
Purging Old Measure Data	7-5
Behavior in a Global Domain Environment.....	7-6
loadmeasure Usage.....	7-6
Loading Image Paths for Positions.....	7-9
Exporting Measure Data: exportMeasure	7-9
exportMeasure Usage	7-10
Exporting Measure Data: exportData	7-12
exportData Usage	7-13
-useLoadFormat Parameter	7-15
Mapping Data Between Domains: mapData	7-16
mapdata Usage	7-16
Transferring Data: transferData.....	7-17
Masks	7-18
No Mask with -clearDest	7-18
No Mask without -clearDest	7-18
Mask with -clearDest.....	7-18
Mask without -clearDest.....	7-18
transferData Usage.....	7-19
Scenarios	7-20
Simple Domain to Simple Domain.....	7-20
Simple Domain to Global Domain	7-20
Global Domain to Simple Domain	7-20
Global Domain to Global Domain.....	7-20
Moving Data between Arrays: updateArray	7-21
updateArray Usage.....	7-21
Scan Domain Data: scanDomain.....	7-22
scanDomain Usage.....	7-23

8 Operational Utilities

Find Alerts: alertmgr Utility	8-1
alertmgr Usage	8-2
Copying Domains: copyDomain Utility	8-3
copyDomain Usage	8-3
copyDomain: Format of the XML Configuration File	8-5
Move a Domain: moveDomain Utility	8-6
moveDomain Usage	8-6
moveDomain: Format of the XML Configuration File	8-7
Assumptions and Requirements	8-7
Minimum Space Requirement	8-8
Setting Miscellaneous Domain Properties: domainprop Utility	8-8
domainprop Usage	8-8
Available Properties	8-8
Calculation Engine: mace Utility	8-10
mace Usage	8-11
Managing the Workbook Batch Queue: wbbatch Utility	8-13
wbbatch Usage	8-14
Workbook Manager: wbmgr Utility	8-16
wbmgr Usage	8-16
Register Measure: regmeasure Utility	8-17
regmeasure Usage	8-17
Register Token Measure: regTokenMeasure Utility	8-23
regTokenMeasure Usage	8-24

9 Informational Utilities

Retrieving Domain Information: domaininfo Utility	9-1
domaininfo Usage	9-1
Checking the Validity of a Domain: checkDomain Utility	9-3
checkDomain Usage	9-3
Determining RPAS Server Version: rpasversion Utility	9-4
rpasversion Usage	9-4
List Contents of a Database: listDb Utility	9-4
listDb Usage	9-4
Printing Data from Arrays: printArray Utility	9-4
printArray Usage	9-4
Printing Data from Measures: printmeasure Utility	9-6
printmeasure Usage	9-6

10 Internationalization

Translation	10-1
Translation Administration	10-2
Translation Administration Workbook	10-4
Hierarchy Labels Worksheet	10-4
Dimension Labels Worksheet	10-5
Workbook Template Group Labels Worksheet	10-5

Workbook Template Labels Worksheet	10-5
Measure Labels Worksheet.....	10-5
Measure Descriptions Worksheet.....	10-5
User Group Labels Worksheet.....	10-5
Message Labels Worksheet.....	10-5
RGRP Labels Worksheet	10-5
11 Commit as Soon as Possible	
Using Commit ASAP	11-1
Managing the Workbook Queue: showWorkbookQueues	11-2
showWorkbookQueue Usage.....	11-2
Commit ASAP Settings: configCommitAsap	11-3
configCommitAsap Usage.....	11-3
Logging and Technical Information.....	11-4
12 Batch Processes and RPAS Utilities	
CSV File Format.....	12-1
RPAS Utilities Logging Options	12-2
Log Levels	12-2
Utilities with Standard Logging.....	12-3
Scripts	12-3
Utilities with Multi-Process Logging	12-4
Utilities with Special Logging	12-5
Using Shell Scripts to Run Batch Processes.....	12-7
A Sample Shell Script	12-7
Common Information and Parameters for RPAS Utilities.....	12-7
Configuration Tools Log Files	12-8
RPAS Intraday Enabler.....	12-8
ride Usage.....	12-10
Scenarios	12-11
Scenario 1	12-11
Scenario 2	12-12
Scenario 3	12-13
Scenario 4	12-13
Scenario 5	12-14
Scenario 6	12-15
Domain Lock Status Utility: domainStatus	12-15
domainStatus Usage	12-15
13 RPAS ODBC/JDBC Driver	
ODBC Configuration.....	13-2
Defining the ODBC Server Configuration Settings.....	13-2
Add RPAS ODBC Manager in the Management Console	13-3
Configure Using the ODBC Manager	13-9
Import the Configuration Changes	13-14
Defining the ODBC Client Configuration for Windows	13-15

Starting the RPAS ODBC Server Process	13-19
Testing the Connection Using Interactive SQL.....	13-19
ODBC Client Configuration for UNIX	13-20
Client Configuration.....	13-20
Testing the Connection.....	13-21
Installing and Using the RPAS JDBC Driver.....	13-22
Updating Environment Variables for the JDBC Driver on Windows	13-22
Updating Environment Variables for JDBC Driver on UNIX/Linux.....	13-23
Using the RPAS JDBC Driver	13-23
Enabling Spy for RPAS JDBC Driver	13-24
Using the jdbcisql Utility Provided with RPAS JDBC Driver	13-25
Using Oracle SQL Developer	13-25
Using Oracle JDeveloper.....	13-26
Using a Java Program.....	13-26
Running the Program.....	13-28
Data Query.....	13-28
Limitations	13-28
Contention.....	13-28
Workbook Queries.....	13-28
Metadata.....	13-29
Fact and Dimension Tables.....	13-30
Measure Security in the ODBC Driver	13-31
Use Cases.....	13-32
Using Metadata Tables to Explore the Structure of a Domain or a Workbook	13-32
Querying Fact Data.....	13-33
Connecting to a Workbook.....	13-34
Requesting Additional Aggregate Tables.....	13-34
Clients.....	13-34
Oracle Business Intelligence Enterprise Edition (OBIEE)	13-35
Configuring the ODBC Client for OBIEE	13-35
Connecting OBIEE to an RPAS Domain.....	13-36
Microsoft Access.....	13-36
JDeveloper	13-37
XML Publisher.....	13-38
Interactive SQL (ISQL) Utility	13-39
Supported and Unsupported SQL Functions	13-39
Supported SQL Functions.....	13-39
Numeric Functions	13-40
String Functions	13-42
Time / Date Functions	13-46
System Functions	13-50
Aggregate Functions.....	13-50
Other Functions.....	13-51
Unsupported SQL Functions.....	13-57
Handling of NULLS.....	13-57
Schema Information.....	13-57
Data Definition Language (DDL)	13-57

Insert	13-57
SELECT Syntax.....	13-57
Value Expressions.....	13-57
Value Functions.....	13-58
Date/Time Functions	13-58
Advanced Value Expressions.....	13-58
Row Value Constructor.....	13-59
Predicates	13-59
Join Operators.....	13-59

14 Publishing Measure Change Events

Configuring Subjects and Measures for Monitoring.....	14-2
Configuring the RPAS JMS Publisher	14-3
Command Line.....	14-4
Configuration Settings.....	14-5
General Settings	14-5
Vendor-Specific Settings	14-6
File Based JNDI Object Store	14-6
LDAP Based JNDI Object Store	14-6
Configuring the RPAS JMS Subscriber	14-7
Command Line.....	14-7
Configuration Settings.....	14-8

15 In-Context Launch

Launching RPAS.....	15-1
Issuing a Launch Request from a Web Page	15-2
Embedding an RWL Applet on a Web Page.....	15-3
Using JavaScript with the RWL Applet.....	15-4
User Authentication.....	15-8
Deploying In-Context Launch.....	15-10
Use Cases.....	15-10
RPAS Launch Context Format	15-11
Client Settings.....	15-12
Server Settings	15-12
Domain Settings	15-12
Workbook Settings.....	15-13
XML Schema	15-13
Open Workbook	15-17
Specifications	15-17
Examples	15-18
Build Workbook	15-19
Specifications	15-19
Example.....	15-22
buildWorkbookContext Utility	15-24
buildWorkbookContext Usage	15-24
Workspace and Oracle Single Sign-On.....	15-24

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-5
<SelectPositions>	B-5
<SetWizardData>	B-5
<WizardFinish>.....	B-6
Workbook Messages	B-7
<OpenWorkbook>	B-7
<Fetch>	B-7
<SetValues>	B-8
<CommitWorkbook>	B-8
<RefreshWorkbook>	B-9
<MenuEvent>	B-9
<CreatePOPMeasure>	B-9
<SaveWorkbook>.....	B-10
<CloseWorkbook>	B-10
RTD use with HP LoadRunner	B-11
Prerequisites	B-11
Header File	B-11
lr_rtd.h	B-11
LoadRunner Scripting	B-12
Script Settings	B-12
Script Sections	B-12
Script Implementation	B-12
Script Section: vuser_init.....	B-12
Script Section: Action (simple test).....	B-13
Script section: Action (real-world).....	B-13
Considerations for Scenario Use	B-14

C Appendix: Environment Variables

Required Settings	C-2
Database Settings	C-2
RPAS_CACHE	C-2
RPAS_PAGE_SIZE	C-3
RPAS_PAGE_SPLIT_PERCENTAGE	C-3
RPAS_LOCK_TIMEOUT	C-3
RPAS_REQUEST_TIMEOUT	C-4
RPAS_USER_MODE	C-4
Log Settings	C-5
RPAS_LOG_BACKUPS	C-5
RPAS_LOG_LEVEL	C-5
RPAS_DAEMON_LOG_PATH	C-5
Profiling Logging	C-6
RPAS_PROFILING_ENABLE	C-6
RPAS_PROFILING_PATH	C-6
Date and Time Setting	C-6
RPAS_TODAY	C-6
RPAS_TODAY_STATIC	C-6
Numeric Precision	C-7
RPAS_INCAGGPRC	C-7
epsilon	C-7
RPAS_PROCESSES	C-7
LDR_CNTRL=NAMEDSHLIB=RPASZONE	C-8

Send Us Your Comments

Oracle Retail Predictive Application Server Administration Guide for the Classic Client, Release 13.3

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

Our goal is to make Oracle products, services, and supporting documentation accessible to all users, including users that are disabled. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Accessibility standards will continue to evolve over time, and Oracle is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For more information, visit the Oracle Accessibility Program Web site at <http://www.oracle.com/accessibility/>.

Accessibility of Code Examples in Documentation

Screen readers may not always correctly read the code examples in this document. The conventions for writing code require that closing braces should appear on an otherwise empty line; however, some screen readers may not always read a line of text that consists solely of a bracket or brace.

Accessibility of Links to External Web Sites in Documentation

This documentation may contain links to Web sites of other companies or organizations that Oracle does not own or control. Oracle neither evaluates nor makes any representations regarding the accessibility of these Web sites.

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/support/contact.html> or visit <http://www.oracle.com/accessibility/support.html> if you are hearing impaired.

Related Documents

For more information, see the following documents in the Oracle Retail Predictive Application Server Release 13.3 documentation set:

- *Oracle Retail Predictive Application Server Release Notes*
- *Oracle Retail Predictive Application Server Installation Guide*
- *Oracle Retail Predictive Application Server Licensing Information*
- *Oracle Retail Predictive Application Server Administration Guide for the Fusion Client*
- *Oracle Retail Predictive Application Server User Guide for the Classic Client*
- *Oracle Retail Predictive Application Server User Guide for the Fusion Client*
- *Oracle Retail Predictive Application Server Configuration Tools User Guide*

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, 13.3) or a later patch release (for example, 13.3.1). If you are installing the base release, additional patch, and bundled hot fix releases, read the documentation for all releases that have occurred since the base release before you begin installation. Documentation for patch and bundled hot fix releases can contain critical information related to the base release, as well as information about code changes since the base release.

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/technology/documentation/oracle_retail.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

Navigate: This is a navigate statement. It tells you how to get to the start of the procedure and ends with a screen shot of the starting point and the statement “the Window Name window opens.”

`This is a code sample`

`It is used to display examples of code`

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](#)

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 sub-domains) 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.

Building and Upgrading Domains

This chapter describes how to build and upgrade domains.

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 installing the domain 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 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 regarding 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 such that the zip file will include 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 the example below, 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 install has the correct version of Java.

Unzip the Configuration File

1. Find the location where 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 should be installed on the server that will store the domains. During that process, the necessary environment variables for RPAS and the Configuration Tools should 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 is very important 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 need to 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 the [Loading Hierarchies: loadHier](#) section for more details.

Building the Domain Manually or Through Command-Line Interface

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

rpasInstall

This section provides the details required for the upgrade of `rpasInstall`.

rpasInstall Usage

```
rpasInstall <arguments>
```

The following table 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>	Required argument. 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.
<code>-ch config_home</code>	Required argument. Path to the directory containing the configuration file.
<code>-cn config_name</code>	Required argument. The name of the configuration.
<code>-in input_home</code>	Required argument. 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: <i>usermgr</i> Utility .
<code>-log log_name</code>	Required argument. Path that includes the name of the log file to be created or updated.
<code>-configdir config_directory</code>	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>	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>	The partitioning dimension. Use if and only if global domain is being implemented without the use of <code>globaldomainconfig.xml</code> .
<code>-rf function_name</code>	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. Failure to use this flag in a patch install will result in changes to configured styles not being imported into the domain.

Notes: When building a domain for the first time, an installs directory is created inside the domain. The installs directory is essential for the patch process and should not be removed, moved, or renamed.

When sending in an issue to Oracle Retail Customer Support, if asked to provide a domain, be sure to provide the installs directory as well. The following information should 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 output below:

```
Time: 58.451  
COMPLETE
```

After Building the Domain

After building the domain, the domain 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: 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 the [Domain Daemon](#) section 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 the [Security Administration Workbook](#) section for more information.

Upgrading and Patching Domains

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


Note: If the domain is not currently on the latest 13.2.3 hot fix, you must upgrade 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 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 13.3 because you will run that utility and argument at the end of this section.

Note: See the solution-specific installation and implementation/operations guides for potential solution upgrade limitations.

Figure 2–2 shows an overview of the upgrade and patching process. The detailed steps are shown below the figure.

Figure 2–2 Upgrade and Patching Process Overview



1. Acquire the latest 13.3 RPAS release and install the following components:

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

Note: If using `rsp_manager` to install an RPAS 13.3 patch, the `-no_domain` option should be used.

2. 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 the "Defining Dimension Properties" section in the *RPAS Configuration Tools User Guide* for instructions on setting 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 upgraded configuration. This file can be created by running the hierarchy.xml report from within the Report Generator of the RPAS Configuration Tools.
4. Run the `convertDomain` utility to convert the 13.2.3.x domain. The exported hierarchy.xml file must be supplied as the input to the `convertDomain` command. In global domain environments this utility is only executed from the master domain. See the [convertDomain](#) section for more information.
5. After executing `convertDomain`, run the `upgradeDomain` utility to complete the conversion of the domain to 13.3. See the [upgradeDomain](#) section for more information.

Note: If you are using these instructions to upgrade to a RPAS 13.3.0.x hot fix, note that the call to `upgradeDomain` updates the domain to the hot fix version you are patching to.

6. At this point, if there are any additional configuration changes that should be made 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), those changes should be made to the configuration modified in step 2.
7. If any additional changes were made to the domain in step 6, those changes should now be applied to the converted domain. To apply these changes, run the `rpasInstall` utility with the `-patchinstall` argument. See the [rpasInstall](#) section for more information.

Notes: After running the `convertDomain` 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, deleting existing styles is not necessary.

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, new hierarchies can be added. 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 the existing hierarchies cannot be changed. However, the order of the new hierarchy can be between or after existing hierarchies.

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: reindexDomain](#)) as well as add or rename dimensions (see [Adding New Dimensions: hierarchyMgr](#)).

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: loadHier](#) section for more information.

Measure Properties Certain measure properties cannot be modified without unregistering and reregistering the measures, which results in the loss of measure data. The measure properties that are able to be modified without being reregistered are listed in the [Register Measure: regmeasure Utility](#) section. If all of the measure properties that are changed are able to be modified without reregistration, the measure data will not be lost in the domain during the patch process.

If even one of the measure properties that change is not able to be modified without reregistration, the patch process results in the unregistration and reregistration 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 the [Register Measure: regmeasure Utility](#) section.

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.


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 2–3 Converting the Configuration



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


Figure 2–4 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 2–5 Confirmation of Conversion



7. The upgraded configuration will now be open 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 RDE, 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.

Updating the solution configuration for each application should be done 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`
- `hmain.db` (which is completely rebuilt separately)
- `measdata.db` (also rebuilt separately)
- `master_staging.db` (use only for loadhier)
- `alitemp*.db`

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 will only contain the `<username>.db` and `<username>.db.lck` files.

The styles directory is copied conditional. 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 (based on the information above), those files and directories must be manually copied to the converted domain.

convertDomain Usage

```
convertDomain [-src srcDomain -dest destDomain | -xmlConfigFile XMLConfigFile]
              -xmlHierFile PathToHierarchy.xml -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 new `BitSize` attribute for each dimension.
3. Use the `-maxProcesses` argument to specify the number of processes can be utilized.

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

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

Argument	Description
<code>-src srcDomain</code>	Use this argument to specify the path to the source domain being converted.
<code>-dest destDomain</code>	Use this argument to specify the path to the destination domain.
<code>-xmlConfigFile XMLConfigFile</code>	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.
<code>-xmlHierFile PathToHierarchy.xml</code>	Use this argument to specify the path to the <code>hierarchy.xml</code> file that is exported from the upgraded configuration. This file should include the new <code>BitSize</code> attribute for each dimension.
<code>-maxProcesses count</code>	If specified, some parts of <code>convertDomain</code> will run in parallel, meaning that it will use a maximum of the defined processes, which are specified by <i>count</i> .

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

upgradeDomain

The `upgradeDomain` utility is used to upgrade just 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]
```

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

Table 2–3 Arguments Used by the upgradeDomain Utility

Argument	Description
-d <i>path</i>	Path to the domain being upgraded.
-verbose	Optional parameter to show the detail about each change that is applied to the domain.
-n	Optional parameter to report which changes would be applied without applying the changes.
-purgeWorkbooks	Optional parameter that purges all existing workbooks and clears the workbook batch queue.
-ignoreSharedNames	Allow upgrade even if dimensions and hierarchies share names.
-apptag	Indicate 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
```

Domain Administration

This chapter describes domain administration tasks such as configuring the Classic Client and using domainDaemon.

Configuring the Classic Client: 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. EConfigure lets you specify communication parameters and produces a file that is used as input to the client. These files must be in 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.

EConfigure 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.

eConfigure Menu Bar

The files produced by EConfigure may contain multiple connections. Each connection will be 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 eConfigure Menu Bar




eConfigure Main View

The main view has the basic connection parameters. On this view, three groups of controls are available:

- [The Connection Group](#)
- [The Domains Group](#)
- [Advanced Settings Dialog](#)

Figure 3–2 eConfigure Utility



The Connection Group

- **Database Server:** The hostname or the IP address of the server, for example, atldev03 or 10.2.1.23. This value should 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).

The Domains Group

- **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 3–2](#) 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.
- **Password:** Like the default database user, default database password is used if a domain specific password 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 rpaDbServer 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.

Domain Daemon

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. After 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] stands for the full path to the directory where the RPAS Server is installed.

DomainDaemon Usage

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

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>	The <i>portNum</i> must be between 1025 and 65535 (inclusive). If <i>portNum</i> is set to auto, it will find any free port.
-loglevel	This option enables additional logging. Note: The -debug option has been deprecated. Instead of using -debug, please 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.
-add <i>pathToDomain</i>	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/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.

Table 3–1 (Cont.) DomainDaemon Arguments

Argument	Description
<code>-remove pathToDomain</code>	Removes the specified domain from the list of domains managed by a DomainDaemon.
<code>-showActiveServers</code>	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.
<code>-stopActiveServers</code>	Stops all active servers. Specify a port number and a process ID.
<code>-stopServer processId</code>	Stops the server using the specified process ID.
<code>-stopUser userId</code>	Stops the server using the specified user ID.

Starting the DomainDaemon

To start the DomainDaemon, execute the DomainDaemon utility. The port number where the DomainDaemon will be 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 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 in to a new session. Graceful termination includes the completion of any pending processes and custom menu or rule group processing in the existing session, followed with an auto-save of the workbook and subsequent termination of the server process.

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

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 4-1](#)).

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 the [Managing Users: usermgr Utility](#) section.


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:

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

Figure 4–1 *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 4–2 Add User Wizard

Enter relevant information in the following fields:

- **User name** - Type the user name that the user will use for logging on.

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 will belong.

- **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** - Type the same password.
5. If the user should have 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. If unsure whether a user should be granted this ability, note that a user's Admin status can later be modified 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 so, access the User and Template Administration workbook. For more information, see the [Security Administration Workbook](#) section.

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 4–3 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 contains the text "RetailPlanners". The second is labeled "Group label:" and also contains 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 4–4 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, all users must be removed 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 4–5 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 4–6 *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 4–7 Edit User Wizard, Page 2

Edit User Wizard

User name:

User label:

Default group:

Other groups:

Change password:

Password verification:

Administrator: ☒

Force password change: ☐

Lock user account: ☐

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

Managing Users: usermgr Utility

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 is in order to protect 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, it 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` and `-convertXmlToDb` do not require a domain.

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

Table 4–1 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>	Use this argument to add a user with a specified name. Use the other arguments specified in the usage to add those attributes for that user. This prompts you to enter the temporary password for the user.
<code>-label label</code>	Use this argument to specify 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>	Use this argument to specify the user group of the user to add to the domain.
<code>-admin</code>	Use this argument to specify that the user to add to the domain has administrative rights.
<code>-addGroup groupName</code>	Use this argument to add a group with a specified name. Use <code>-label</code> to specify the label for the group.
<code>-remove userName</code>	Use this argument to remove the user with the specified name from the domain.
<code>-removeLabel label</code>	Use this argument to remove all users with this label.
<code>-removeGroup groupName</code>	Use this argument to remove a group with this groupName.
<code>-list</code>	Use this argument to list all the users registered to the specified domain.
<code>-print</code>	Use this argument to print the specified user or group information.
<code>-user username</code>	Use this argument to specify the user name in the specified domain to print. This argument is only applicable to <code>-print</code> option.
<code>-group groupname</code>	Use this argument to specify the group in the specified domain name to print. This argument is only applicable to <code>-print</code> option.
<code>-importDb</code>	Use the argument to import the database. The database must be located in the domain's input directory. The database is time stamped and moved to the processed directory upon successful completion. Existing user are skipped unless <code>-replace</code> is used.
<code>-replace</code>	Use this argument to update existing users when using <code>-importDb</code> . The user label, groups that user belongs to, admin status, and account lock status is updated. Password information is not affected by the update.
<code>-lock userName</code>	Use this argument to lock the specified user. This prevents the user from logging in the domain.
<code>-unlock userName</code>	Use this argument to unlock the specified user. This allows the user to log in the domain.

Table 4–1 (Cont.) Arguments Used by the usermgr Utility

Argument	Description
<code>-exportXml path</code>	Use this argument to create an XML file that contains all users and groups in the selected domain. Passwords and password histories are not exported.
<code>-exportDb path</code>	Use this argument to create a database that contains all users and groups in the selected domain. This prompts you for new temporary passwords for admin and user accounts.
<code>-convertDbToXml</code>	Use this argument to convert a user database to a user XML file. Passwords are not included in the conversion.
<code>-convertXmlToDb</code>	Use this argument to convert a user XML file to a user database. This prompts you for temporary passwords for admin and user accounts.
<code>-src path</code>	Use this argument to specify the source file used in <code>-convertDbToXml</code> and <code>-convertXmlToDb</code> .
<code>-dest path</code>	Use this argument to specify the destination file used in <code>-convertDbToXml</code> and <code>-convertXmlToDb</code> .

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 4–2 XML Schema

Outer Tag	Inner Tag	Description
GROUP	LABEL	The group's label.
USER	LABEL	The user's label.
USER	DFLT_GRP	The user's default group.
USER	OTHER_GRP	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.

Below is a sample users.xml file.

Figure 4–8 Sample users.xml File

```
<?xml version="1.0" ?>
<VERSION>1.0</VERSION>
<GROUPS>
  <GROUP NAME="adms">
    <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_GRPS>grp1,grp3</OTHER_GRPS>
  </USER>
  <USER NAME="usr3">
    <LABEL>user_3</LABEL>
    <DFLT_GRP>grp1</DFLT_GRP>
    <OTHER_GRPS>grp1</OTHER_GRPS>
  </USER>
</USERS>
```

Use Cases

Below is a list 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 used 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 the `-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.
- If 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 `rpasInstall` 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 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, the following validation rules apply:

- 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 5–1 *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.

When working with 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 5–2 Auto Workbook Maintenance Wizard

Please enter the following information for this workbook.

Build label:

Build frequency (in days):

Next build date:

Category:

Please select the saved access for this workbook.

☐ User

☐ Group

☒ World

Please select the group that owns this workbook.

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 the [Workbook Batch Category Management](#) section.
 - **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 5–3 Edit Auto Workbook Settings Worksheet



The screenshot shows a window titled "Edit Auto Workbook Settings" with a table of workbook settings. The table has columns for Category, Frequency (in days), Next Build date, Owner, Save Group, Template, Workbook Name, and World Access. The data row shows "PlannerWorkbooks" as the category, with a frequency of 1, next build date of 1/30/2010, owner "adm", save group "Admin", template "Sales Planning", and workbook name "PlannerWorkbooks". The World Access column shows a red minus sign. A "Measure" button is in the top right corner. At the bottom, there is a tab labeled "ADMQ" and a scroll bar.

	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 the [Auto Workbook Maintenance Wizard](#) and [Managing the Workbook Batch Queue: wbbatch Utility](#) sections.


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 5–4 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 5–5 Adding a Category

The screenshot shows the 'Workbook Batch Category Management Wizard' dialog box. The title bar is blue with a red 'X' button. The main area is light beige. At the top, it says 'Please choose a workbook batch category operation:'. There are three radio buttons: 'Add Category' (selected), 'Delete Category', and 'Change Category Label'. Under 'Add Category', there are two text boxes: 'Name:' with 'AdminWorkbooks' and 'Label:' with 'Admin Workbooks'. Under 'Delete Category', there is a 'Select' dropdown menu with 'Test'. Under 'Change Category Label', there is a 'Select' dropdown menu with 'Test' and a 'New Label:' text box. At the bottom, 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 5–6 *Deleting 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", "Delete Category" (which is selected with a green dot), and "Change Category Label".

For the "Add Category" option, there are text boxes for "Name:" (containing "Produce") and "Label:" (containing "Produce Workbooks").

For the "Delete Category" option, there is a "Select" dropdown menu currently showing "Admin Workbooks". A list is open below it, showing three items: "Admin Workbooks", "Planner Workbooks", and "Produce Workbooks" (which is highlighted in blue). Below the dropdown is a "New Label:" text box.

For the "Change Category Label" option, there is a "Select" dropdown menu and a "New Label:" text box.


At the bottom of the dialog are five buttons: "Cancel", "<Back", "Next>", "Finish", and "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 5–7 *Changing a Category Label*



The screenshot shows a dialog box titled "Workbook Batch Category Management Wizard" with a close button (X) in the top right corner. The main area contains the instruction "Please choose a workbook batch category operation:". There are three radio button options: "Add Category", "Delete Category", and "Change Category Label". The "Change Category Label" option is selected. To the right of the radio buttons are input fields: "Name:" with a text box containing "Produce", "Label:" with a text box containing "Produce Workbooks", "Select" with a dropdown menu showing "Admin Workbooks", and "New Label:" with a text box containing "Executive Workbooks". At the bottom of the dialog are five buttons: "Cancel", "<Back", "Next>", "Finish", and "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 will be 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 5–8](#), products that were previously unassigned are now designated as A, B, or C products.

Figure 5–8 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 5-9 displays the results when you access a quick menu in a wizard and change the rollup 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 5-9 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 the table below.

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 5–10 Password Policy Measures Settings




Table 5–1 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

Table 5–1 (Cont.) Password Policy Measures

Measure	Description	Allowable Values	Default Value
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 pick list: 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
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 out** by the domain administrator. This prevents the user from logging in to the RPAS Client. The account remains locked out until the administrator re-enables the account. Account lockouts can be set or cleared by the domain administrator by using the User Management utility.

The account may be marked as **must change password**. This is useful for brand new accounts. The user is allowed to logon with the current password and then forced to select a new password.

Must change password can be set or cleared by the domain administrator using the User Management utility.

Account Lockout may be enabled for a domain. The domain administrator selects a number of failed logon attempts after which the User account will be marked as locked out. The account will remain locked out until the administrator re-enables it.

Password expiration may be enabled for a domain. The domain administrator selects a number of days after which passwords expire. When the user logs in, the system requires a new password to be entered if the configured number of days has passed since this user entered a new password.

Password history may be enabled for a domain. The domain administrator selects a number of passwords to save. When the user attempt to change passwords, the system will not permit any password already stored in the password history to be used again.

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* from 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. The following table demonstrates how access is granted or denied based on all combinations of settings:

Table 5–2 Grant/Denial of Access by Combination of Settings

Security set by Position Denied = False Granted = True			Based on settings on left, user is Granted or Denied access
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/or 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.

Notes:

- 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 checkbox for each available user and dimension combination. A checkmark 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 Limit 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 checkbox 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: Due to its dynamic nature, formatting settings cannot be saved in the Measure Analysis workbook.

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 5–11 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 5–12 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 5–12 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 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.
- On the Analysis tab, select **Measure Analysis** and click **OK**.
- The Measure Analysis Wizard opens and prompts the user to select the measures to be displayed in the new workbook. Use Ctrl-Click and/or Shift-Click to select multiple measures. Click **Next**.
- For each hierarchy specified in the base intersection of the measures selected, the user will see a hierarchy wizard from which to select positions to view. Repeat this step for each hierarchy wizard and click **Next**.
- Click **Finish** to open the Measure Analysis workbook.
- 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, are 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.
- 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: 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 has also been enhanced to allow a simple compression method that can skip duplicated values line by line.

Note: The following notes concern 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 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](#).

loadHier 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/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.

The loadHier utility stops with an error if the loadHierBk directory exists in the data directory of the domain, which indicates there may be a non-recoverable error 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, that 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 very important to specify the measures wisely, especially with the `-includeMeasList` option. This is because there will no data moving for those measures not included in the list, that data will be lost when loadHier completes.

loadHier Usage

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

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

Table 6–1 *loadHier Utility Arguments*

Argument	Description
-d domainPath	The domain in which to load the hierarchy data.
-load hierName	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	This is a switch that prevents the removal of input files and temporary data files that are generated during the hierarchy load process. Input files will remain in the input directory of the domain after the process is completed. This option is often used for debugging or troubleshooting purposes.
-logLoadedPositions	This argument will enable 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	This argument enforces new hierarchy roll-up changes. New roll-up changes will override or dominate existing hierarchy roll-ups in the event 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	Use this argument to force NA consistency when the current NA value is different from the originally defined NA value for the measure.

Table 6–1 (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#,	Use to specify comma separated default domain paths that will be 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>	Use one of these arguments to optimize 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>.

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 positions is added.

Example:

Let's say we have a global domain that consists of two local domains, ldom0 and ldom1.

Let's assume we use the following loadHier command:

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

Let's say that in this call, there are three new partition positions (part1, part2, and part3) 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 we only add two new partition positions (part1 and part2), 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 we add 5 partition positions (part1, part2, part3, part4 and part5), 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 will be 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 were patched to make a dimension translatable but the translation measures were not loaded, RPAS users would 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 were to be 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 will be named r_skulabel.csv.ovr and will have 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 10-1, "Supported Languages with Language Identifiers"](#).

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 due to errors when preparing translation files, 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 the 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.

Before integer indexing is described in more detail, review the following terms:

- **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. Before 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.

For example, let's say you load some items into the Product hierarchy. These items are externally identified by the position names and labels (as shown in the second and third columns in the table below). The table below 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.

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 was loaded into the domain, its position ID, SKU3456, was mapped to an available integer index, which was 2. After these items were 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.

Now let's say that 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.

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 will need to reindex the dimension first.

For instance, let's say you want to add 5 more colors of sweatpants to the dimension. This dimension has 10 integer indexes, so it can hold 10 total position IDs. As the dimension is now, it has only 3 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, you would 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: `reindexDomain`](#).

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			

Now let's say that you want to 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 will 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: `dimensionMgr`](#)).

Changing the Bit Size

To change the bit size after a domain has been built or upgraded to 13.3, perform 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 *RPAS Configuration Tools User Guide*.
2. Run the `reindexDomain` utility. See [Reindexing Domains: `reindexDomain`](#) for instructions.

Or,

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

Reindexing Domains: 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 of the physical IDs makes the domain load and run faster.

Compression is achieved by removing the gaps, which develop due to hierarchy operations like purge and delete. Gaps essentially block certain physical IDs within an address space and prevents 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.


Notes: While working with a global domain, the reindexDomain utility runs from 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, workbook refresh and build operations may run slower. 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 6–1 Reindexing Domain Process Flow

reindexDomain Usage

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

There are three unique usages of the reindexDomain utility: reindex, analyze, and domain properties. These three usages are described in the sections below.

ReindexDomain Usage: 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 6–2 reindexDomain: Reindexing Usages

Usage	Use
<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.
<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 Usage: Analysis

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

Table 6–3 reindexDomain: Analyzing Usages

Usage	Use
<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 Usage: Domain Properties

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

Table 6–4 reindexDomain: Domain Properties

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

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

Table 6–5 *reindexDomain Usage*

Argument	Description
-d	Use this argument to specify 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	Use this argument to force 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	Use this argument only when you are prepending the calendar hierarchy. This argument 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.
-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, bitsize, number of available and used position IDs, the threshold ratio, and whether reindexing is necessary.
-prop	Use this argument to see 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 the [Run Reports](#) section.) 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 positions 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 Bit Size 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 as any initial value. The default value is 10% 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**

In this case, the calculation would produce 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 defrag 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**

Then the calculation would produce 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 defragging 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

There are three reports you can run with the reindexDomain that 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 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 was 15, but then reindexing caused that same SKU to be at index 2, then the dimregistry_version would be 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 they are named after their version number. The path up to the version is the same across all versions to enable data to be processed from one version to another.

Each dim registry version stores dimension index information in a database in the domain. The name of this database is dimRegistryX.db, where X is the dim registry version. The first dim registry version is numbered 0; therefore, the database is named dimRegistry0.db. The reindexing process does not remove old dim registry version databases. Over time, the number of these databases can become large, and you may want to remove some of them. However, do not remove the current nor penultimate database. For instance, if your current dim registry version is 15, you can remove versions 0-13. You should keep dim registry version 14 in case the previous reindex process failed or was cancelled.

- **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 reindexing process starts, the domain should 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 usage 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. You can 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 seperated 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 `-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) which 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.

Below 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 which 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 example on the previous page.

Reindex Domain Examples

This is not a comprehensive list of scenarios that can occur. This is just to highlight some examples of the reindexDomain functionality.

Condition 1: Change of Bit Size

Table 6–6 *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 6–7 *-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 6–8 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 $C-M < TR$.
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.

Table 6–8 (Cont.) Fragmentation Check Scenarios

Scenario	Utility Option	Result
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: 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 running this utility 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, 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 is due to the fact that 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 after a full hierarchy purge that you run the optimizeDomain utility with the -cleanOnly argument to remove all stale data from the domain.

optimizeDomain Usage

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

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

Table 6–9 *optimizeDomain Usage*

Argument	Description
-d domainpath	Use this to define the path to the domain that you want to optimize. If no options are specified, by default optimizeDomain defrags and cleans the entire measure set that have arrays which 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	This option selectively defrags 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.
-cleanOnly	This option 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	This option forces the utility to defrag the entire measure arrays set. This should be used in combination with the -defragOnly option. This option makes the defrag process run longer. This option overrides the Btree density threshold (80%) and defrag all measure arrays regardless of their BTree density.

Table 6–9 (Cont.) optimizeDomain Usage

Argument	Description
-excludeMeasList	Use this option to define the path to exclude list XML file. Optimize performance by excluding this list of measures while defragging or cleaning. This option cannot be used with -includeMeasList.
-includeMeasList:	Use this option to define the path to include list XML file. Optimize performance by including only this list of measures while defragging or cleaning. This option cannot be used with -excludeMeasList.
-returnFailure	Use this option to return a non-zero value if any of the databases or arrays fail to defrag or clean. The optimizeDomain process will still run 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

Say we have the following measures list and corresponding BTree density ratios. SKU dimension had some positions purged from the domain and marked for cleaning.

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

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

	Meas1	Meas2	Meas3	Meas4
No specified options	X	X	X	
-includeMeasList Meas1,Meas3, Meas4	X		X	
-excludeMeasList Meas1,Meas3		X		
-defragOnly	X		X	
-cleanOnly	X	X		
-defragOnly -force OR -force	X	X	X	X
-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 below. Each step is described in greater detail in the following sections.

Caution: Due to its invasive impact, adding new dimensions is different from 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.

Notes: 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 *RPAS 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: hierarchyMgr](#).
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 use 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 an HBI measure 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 sub-domain. The file names contain an internal sub-domain index, for example `sales.0.csv.rpl`, `sales.1.csv.rpl`, and so on.

For more information about `exportMeasure`, see [Exporting Measure Data: `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: `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,cld`). For more information about `loadHier`, see [Loading Hierarchies: `loadHier`](#).

Adding New Dimensions: `hierarchyMgr`

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.

Notes:

- 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 will cause an error. All loaded positions will be 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: 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: informalPositionMgr](#).

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 CSV format with header line is supported. The exported files will have one measure per file.

For more information, see [Loading Measure Data: loadmeasure](#).

Adding New Hierarchy Dimensions Sample Script

Below 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

# -----
# 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

```

Filtering Hierarchies: filterHier

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 don't 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 would have their relevant hierarchy load data in one master file, but using loadhier, the retailer would not be able to load just what is relevant to the domain from the master. System integrators would 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 ridding them of the need to write custom scripts. The utility analyzes the target domain and trims down the master file to only have 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}
```

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

Table 6–10 *filterHier Utility Arguments*

Argument	Description
-d domainPath	The domain in which to load the hierarchy data.
-input inputpath	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 has introduced 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 will combine 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 should 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 will be arranged to match the hierarchy format of the target domain. Any dimensions from the input files that are not present in the output domain will be filtered out.

The output file will be a properly formatted .csv.dat file, and will be 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 proceed 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 need to 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 that was 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 need to 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: 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/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:

- Add a new position along the partition dimension and allocate it to an existing or new local domain.
- Remove an existing position from the partition dimension.
- Remove local domains (this is automatic if all partition-level positions in a local domain are removed or moved).
- Move an existing partition position from one local domain to an existing or new local domain.

Runs loadHier to apply the position addition/removal to 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 should run the utility reconfigGlobalDomainPartitions to by specifying the sub-domain 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 adding positions (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.

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
```

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

Table 6–11 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 will be 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, utilizing up to the given number of processes.

Table 6–11 (Cont.) reconfigGlobalDomainPartitions Utility Arguments

Argument	Description
-forceInputRollups	This argument will prevent this utility from failing in instances where there is a roll-up conflict in the input file provided to the utility. This argument enforces new hierarchy roll-up changes such that they dominate existing hierarchy roll-ups in case they conflict with the roll-ups 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.

Below 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.

The following table provides descriptions of the arguments used by the input file.

Table 6–12 Arguments Used by the input File

Argument	Description
command_name	Valid values are add or move.
path_to_local_domain	Path to the local domain to which positions are added or moved.
sample_pos	One or more positions that are being added or moved to the designated local domains.

Notes, Assumptions, and Limitations

- Position names are separated by commas and must be valid external position names without the prefix of a dimension (for instance, 0102,0144,0152,0160).
- The utility backs up the required data and will automatically restore 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 adding positions 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: 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.

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. One line 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 an already used name. 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 max 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), you must perform two separate operations: (a->tempa,b->tempb) and (tempa->b, tempb->a).

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

Table 6–13 Arguments Used by the renamePositions Utility

Argument	Description
-d <i>domainPath</i>	Specifies the full path to the domain.
-hier <i>hierName</i>	Hierarchy in which positions are being renamed.
-input <i>inputDirectory</i>	Input directory where input file with positions to rename is located. This is optional. Utility looks for hierarchy data files with "rn" between hierarchy name and .dat extension (for instance, prod.rn.dat).
-log <i>logFileName</i>	Optional parameter to generate log file to file name other than default. The default file name is hierRename.log. 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.

Setting Properties for Dimensions: 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}
```

The table below 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 6–14 *dimensionMgr Utility Arguments*

Argument	Description
<code>-d <i>pathToDomain</i></code>	Specifies the path to the domain.
<code>-dim <i>dimensionName</i></code>	Specifies the name of the dimension to which the settings will apply.
<code>-label <i>newLabel</i></code>	Use this argument to specify a new dimension label.
<code>-enableDPM</code>	This argument enables Dynamic Position Maintenance (DPM) for the specified dimension (<code>-dim dimensionName</code>). 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, LNGS).
<code>-enableTranslation <i>width</i></code>	This argument 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 6–14 (Cont.) dimensionMgr Utility Arguments

Argument	Description
-specs	This argument 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. Below 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/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>	This argument 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	This argument changes 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	Use this argument to set 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: 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]
```

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

Table 6–15 Arguments Used by the exportHier Utility

Argument	Description
-d <i>pathToDomain</i>	Specifies the path to the domain.
-hier <i>hier_name</i>	The name of a hierarchy in the domain from which the .DAT file will be generated.
-datFile <i>dat_File</i>	The path/location where the .DAT file will be created. This .DAT file can then be used with loadHier to load the hierarchy into a domain.
-fixedWidth	This argument indicates that the output .DAT file is a fixed-width file instead of a comma-separated value (CSV) file format.
-onlyFormal	This argument exports only formal positions to the .DAT file. If this is option is specified, informal positions will be skipped.
-onlyInformal	This argument exports only informal positions to the .DAT file. If this is option is specified, formal positions will be skipped.
-genheader	This argument generates a header line for CSV export. The line contains fields to identify the dimension and its label column.
-udd	This argument exports the user-defined definitions. It 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.
-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	This argument 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: informalPositionMgr

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

There are three main areas of use for informalPositionMgr, which are described in the following sections:

- [Basic Operations: Informalize, Formalize, Remove](#)
- [Create Informal Positions in Bulk](#)
- [Data Slice Copying](#)

Basic Operations: Informalize, Formalize, Remove

Informalizing, formalizing, and removing informal positions all have similar command line options.

```
informalPositionMgr -d domainPath -hier hierName -operation [remove | formalize |
informalize] [-dir inputDir | -file inputFile] [-retain]
```

The following table provides descriptions of the arguments used in basic operations of the informalPositionMgr utility.

Table 6–16 Arguments Used by the informalPositionMgr Utility: Basic Operations

Argument	Description
-d <i>domainPath</i>	The domain to run the utility.
-hier <i>hierName</i>	Use this option to specify the hierarchy to operate on.
-operation <i>informalize formalize remove</i>	Specify one of the three basic operations: informalize, formalize, or remove. You can use only one of these operations at a time. ■ <i>remove</i>: Removes informal positions of the specified hierarchy. ■ <i>formalize</i>: Formalizes positions of the specified hierarchy. ■ <i>informalize</i>: Informalizes positions of the specified hierarchy.

Table 6–16 (Cont.) Arguments Used by the informalPositionMgr Utility: Basic

Argument	Description
-file <i>inputFile</i> or -dir <i>inputDir</i>	Use one and only one of these arguments to specify the input files to use during the operation specified. ▪ -dir <i>inputDirectory</i>: Processes all applicable files under <i>inputDirectory</i>. ▪ -file <i>inputFile</i>: Processes the <i>inputFile</i> only. All files that match the naming patterns below are processed: ▪ Remove files: {hierName}.remove[.extension] ▪ Formalize files: {hierName}.formalize[.extension] ▪ Informalize files: {hierName}.informalize[.extension] For example: prod.informalize.20100220 The format of the returned content 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>
-retain	Use this argument 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 6–17 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 and a warning message is logged and the operation continues.

Create Informal Positions in Bulk

This feature allows 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 exception is thrown.

If the specified dimension 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 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 to 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.

Note: To manage the performance impact of bulk creation, a limit is placed on how many positions can be added at one time. The default limit is 5000 per dimension, which can be overridden using the following domain property:

```
MAX_NEW_BULK_DPM_POS
```

If the limit is exceeded, the program stops and no positions are created in the domain.

The following table provides descriptions of the arguments used by the create bulk positions functionality of the informalPositionMgr utility.


Table 6–18 Arguments Used by the informalPositionMgr Utility: Create Bulk Positions

Argument	Description
<code>-d <i>domainPath</i></code>	The domain to run the utility.
<code>-create</code>	Use this option to batch create new positions with auto-generated names.
<code>-dim <i>dimName</i></code>	Use this option to specify the dimension to operate on.
<code>-rollups <i>dim1:pos1,dim2:pos2,...</i></code>	Use this option to specify 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.
<code>-spreads <i>dimA:ratioA,dimB:ratioB,...</i></code>	Use this option to specify the spread counts for the auto-generated positions. This is optional. It defaults to 1 if not specified.

Table 6–18 (Cont.) Arguments Used by the informalPositionMgr Utility: Create Bulk

Argument	Description
<code>-n posCount</code>	Use this argument to enter the number of positions to create on the specified dimension.

The following sections show some examples of creating informal positions in bulk on the product hierarchy. [Figure 6–2](#) displays a product hierarchy with alternative branches and four DPM-enabled dimensions (STYL, STCO, SKU, and SZ12).

Figure 6–2 Sample Product Hierarchy with Alternative Branches

Example 1: Create 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 two options for SZ12:

- Option 1: Specify the rollup for SZ12
- 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.

Option 1

```
informalPositionMgr -d domainPath -create -dim STYL -n 100 -rollups
SCLS:scsls1001,CLR:clrwhite,SPLR:splr100,SZ12:sz12null -spreads STCO:2,SKU:5
```

This command creates the following:

1. 100 new informal positions on the STYL dimension. These positions all roll up to SCLS:scsls1001.
2. Since the spread ratio for STCO dimension is 2, 200 new informal positions are created at the STCO level, and every two of them roll up to one STYL 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

```
informalPositionMgr -d domainPath -create -dim STYL -n 100 -rollups
SCLS:scsls1001,CLR:clrwhite,SPLR:splr100,SIZ1:siz1null,SIZ2:siz2lt02 -spreads
STCO:2,SKU:5
```

This command creates the following:

1. 100 new informal positions on the STYL dimension. These positions all roll up to SCLS:scsls1001.
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, 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: Create 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 two options:

- Option 1: Specify a rollup position.
- 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).

Option 1

```
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. 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, 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

```
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. 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, 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 4 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 of one position to those of another. It is especially 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.

```
informalPositionMgr -d domainPath -copy -file inputXMLFile [-retain]
```

The following table provides descriptions of the arguments used by the data slice copying functionality of the informalPositionMgr utility.

Table 6–19 Arguments Used by the informalPositionMgr Utility: Data Slice Copying

Argument	Description
-d <i>domainPath</i>	The domain to run the utility.
-copy	Use this option to copy data slices within measures.
-retain	Use this argument if you do not want to move the input file to the processed subdirectory after a successful run.
-file <i>inputXMLFile</i>	Use this argument to specify the XML file that contains the configuration settings for the copying operations.

[Table 6–20](#) shows the available settings for the XML input file. This file defines the configuration settings for the copying operation.

Table 6–20 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

[Figure 6–3](#) below shows a sample XML input file.

Figure 6–3 XML Input File Schema in XSD Format

```

<?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="measures" 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>

```

Figure 6–4 XML Input File Example

```

<?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>

```

Merge Informal Positions to Formal Positions

Below 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: pqdMgr

`pqdMgr` 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]
```

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

Table 6–21 Arguments Used by the pqdMgr Utility

Argument	Description
-d domainPath	Specifies the path to the domain.
-load xmlFile	Use this argument to load position lists from an input XML file. Position lists with an operation attribute of delete are ignored.
-delete xmlFile	Use this argument to delete 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	Use this argument to delete all PQD lists from the domain.
-validate xmlFile	Use this argument to validate the XML file and report any errors. No impact on the existing PQD files in the domain.
-export outFile [-user userName -world]	Use this argument to export existing PQDs in the domain for a user or world access level. The file specified by outFile is overwritten. Requires one of the following options: ■ -user userName exports PQDs for the provided userName 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

- 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.

Loading Measure Data: 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` 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 7-1 Loading Measure Data



- The base intersection is SKU.
- The clearint is dept.
- Data is present for SKUs 3, 5, and 11; which fall under dept 1 and dept3.
- Data will be 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) will be 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 overwrite, 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 will be switched to NA. In other words, all data at the base intersection for the measure will be removed before cells are populated with the data from the incoming file.
- .inc - (Increment) Increment mode should only be used with numeric measures where 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 would be 5 (2 incremented with 3).
- .clr - (Clear) Clear mode is a variation of replace mode. It is meant to be used when measure data is loaded in parts or staggered in time, such 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 will be 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 there is data 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, but 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 would be 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.

`loadmeasure` 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.

Notes:

- 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.
-
-

`loadmeasure` 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 would 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

loadmeasure allows multiple measures to be loaded from a single file. You can loaded 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 multiple measures, loadmeasure uses the start positions of measures as defined in the Data Interface Tool for the fixed file format 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 start position (for fixed file format) for the measure values have been configured to be 40, 110, and 70 respectively, then when using the CSV file multiple.csv.ovr, loadmeasure will assume that after the dimension columns, the first column is measure A, then C, and then B because 40 (A) is less than 70 (C) is less than 110 (B).

It is not necessary to load all measures in a multiple measure file. Administrators can choose to load only a subset of measures from the multi-measure file. However, if Administrators wish to reuse the same file for loading different measures at different times in a batch, they must use the -noclean option to ensure that loadmeasure does not move the file to the processed folder after processing the first load request.

Note: All the measures contained in the CSV file must be loaded in one loadmeasure command. A single measure cannot be loaded from a CSV file that contains data for multiple measures.

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 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 fault in 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 substitutes 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. loadmeasure 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 loading multiple measures in a single call, only the preprocessing script for the first listed measure will have 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 N/A 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 5, then at data load time, all data that is older than 5 days before `RPAS_TODAY` will be purged.

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 will load 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 will then split the input files and load 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 will only occur 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}
```

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

Table 7–1 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, <code>measurename1</code> , <code>measurename2</code> , <code>measurename3</code>). If more than one measure is specified, all the measures must be in the same input file.
-applyloads	Use this argument to apply any staged loads for the named measure. If the measure is registered to be a stage-only measure, <code>loadmeasure</code> will put the load in a staging area but will not update the measure until <code>loadmeasure</code> is called again with this argument. Upon the use of this argument, <code>loadmeasure</code> will apply all loads that have been queued up in the staging area. It will clear out the staged loads unless the measure's <code>loadsToKeep</code> property has been set to a non-zero number. In that case, it will not clear out the latest <code>loadsToKeep</code> loads. Note that only <code>.ovr</code> , <code>.inc</code> , and <code>.clr</code> loads can be staged. <code>.rpl</code> 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 should not be used for measures that are not stage-only.
-processes <i>max</i>	This argument 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 5 as the maximum number of processes, then up to 5 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 7–1 (Cont.) Arguments Used by the loadmeasure Utility

Argument	Description
-noClean	Use this option to prevent the input files from being moved to the processed directory. This option is meant for use 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	Force 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 N/A 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 5, at data load time all data that is more than 5 days before RPAS_TODAY will be purged. This option does not require you to load any new data.
-loglevel <i>level</i>	Use this argument to set the logger verbosity level. Possible values: all, profile, audit, information, warning, error, or none.
-noSplit	Use this argument to load the pre-split input files (created by -splitOnly) into the local domains. This option should only be used in global domain environments.
-splitOnly	This argument 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% of the load time. System integrators may be able to improve batch performance by breaking away file-splitting from actual measure loading. This is specifically useful if a multi-measure file is being used such 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	This argument can be used to defragment 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 7–1 (Cont.) Arguments Used by the loadmeasure Utility

Argument	Description
<code>-recordLogLevel <i>level</i></code>	This argument is used to set 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 with the use of the <code>-recordLogLevel <i>level</i></code> 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 <code>-loglevel</code>). 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.
<code>-inDir <i>InputDirectory</i></code>	Only .rpl files can be used with this option, and only CSV format with header line is supported. The header line is used to correspond 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 measure that may contains 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 <code>exportMeasure -hier</code> only. Manual name-indexing of the files is not recommended. The filename property of the measure is not considered with the <code>-inDir</code> 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 should be stored in a measure called `r_images_<dimension name>` where `<dimension name>` should 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
```

`<domain path>` is the path to the domain.

Exporting Measure Data: exportMeasure

`exportMeasure` is a command-line utility that may be used to export domain or workbook measure data from RPAS in either CSV or 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, or a range specified on Calendar dimension, or a combination of the two.
- Supports multiple processes for better performance in a global domain environment.

Note: `exportMeasure` 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.

exportMeasure Usage

```
exportMeasure -d pathToDomain -out outFile [COMMAND] [OPTIONS]
```

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

Table 7–2 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.
-intx <i>intxString</i>	Specifies the intersection at which to export measures. If measure's base intersection is higher than the export intersection, replication is used to spread the measure down to the export intersection. If measure's base intersection is lower than 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 ("_") should 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 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 current output file. If not specified, 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 will be 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 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.

Table 7–2 (Cont.) Arguments Used by the exportMeasure Utility

Argument	Description
<code>-outDir outputDirectory</code>	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.
<code>-upperCase</code>	This argument 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.
<code>-meas "measSpec, measSpec ..."</code>	Must specify one. measSpec is measName.modifier. The -meas argument may be repeated to export multiple measure arrays to the same output file. modifier include the following: .precision<double>, specify the precision for numeric measure .format<formatString>, specify the user defined export format The examples below provide valid measure specifications given MeasNameA 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 MeasNameB 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: 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 should only be called to export data from the master domain. The utility will extract the data from either the local domains or the master domain depending on where the data resides, which depends on the level at which the global domain environment is partitioned.

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 the arrays appear in the -array parameter is the order that they will 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* would export 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}
```

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

Table 7–3 Arguments Used by the exportData Utility

Argument	Description
-d <i>domainPath</i>	Specifies the domain that contains the data that to export.
-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>	Use this argument 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 7–3 (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	This parameter 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 7–3 (Cont.) Arguments Used by the exportData Utility

Argument	Description
-precision <i>precisionValue</i>	This parameter 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	This argument enables the use the NA value of the array instead of the NA value specified in measSpec or arraySpec.
-upperCase	This argument 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	This argument 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	This argument enables the use of the format as specified by the measure property. The level at which the data is stored in the domain will be used. The -dim parameter is not allowed with the -useLoadFormat.

-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 will be exported into the same column in the output file. If the full measure export specs are not provided including the cellFormat, naValue and naFormat, the default format will be used. The default export format for each type of measure is listed below:

- Integer: %<width>.0f
- Real: %<width>f
- String: %<width>s
- Date: %Y%m%d
- Boolean: TRUE or FALSE as string

All value will be exported right aligned as in the measure loading file.

If users give a full measure specs, user-specified cellFormat, naValue, and naFormat will be used rather than the default format.

Users can either use the default format by specify the measure name only, or give the full specs. It is not allowed to give a partial measure spec.

If users specify multiple measures to be exported into the same file, these measures will each 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 will be exported by default. The external position name will be 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 will be left aligned in the export file.

Mapping Data Between Domains: 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 running this utility, 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}
```

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

Table 7–4 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>	Apply 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>	Apply mapdata only on the given array. The database in which the array resides must be specified with the -db argument.

Transferring Data: transferData

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 is 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 desired.

Note: You cannot use the transferData utility for forced non-HBI (higher based intersection) measures. Forced non-HBI measures are measures that should be an HBI measure 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 described in the [Adding New Dimensions to Hierarchies](#) section.

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 the [transferData Usage](#) section.

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.

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 below.

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 7–5](#) illustrates the results of the mask and flag combinations.

Table 7–5 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
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}
```

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

Table 7–6 *transferData Utility Arguments*

Argument	Description
-d pathToDestDomain	Use this argument to specify the path to the domain containing the destination measure. This is required.
-src pathToSrcDomain	Use this argument to specify the path to the domain containing the source measure. This is required.
-measMap	Use this argument to enter a 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	Use this argument to define the maximum number of processes to run in parallel in a global domain setup.
-clearDest	When used with a mask option, this argument ensures that the destination measures array is not cleared and its NA value is not changed. The utility copies the populated/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	Use this argument to ensure only formal positions that have a matching external name in the destination measure are copied over.

Scenarios

The following sections describe the four scenarios of 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 if the destination measure is a non-HBI measure, 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 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 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.

Moving Data between Arrays: 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}
```

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

Table 7-7 Arguments Used by the updateArray Utility

Argument	Description
-destArray <i>dbPath.arrayName</i>	Required argument to specify the destination array where the data will be copied. <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.

Table 7–7 (Cont.) Arguments Used by the updateArray Utility

Argument	Description
-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.
-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 <pre>-srcRange WEEK200811011:WEEK200811022</pre> <pre>-destRange WEEK200811011:WEEK200811022</pre>
-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: 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 attempting to perform a repair of the databases, the utility has a 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 [-?|-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 could run two consecutive commands for detecting corruption and data loss, although this would be less efficient.

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

Table 7–8 Arguments Used by the scanDomain Utility

Argument	Description
-version	Prints the version of the utility.
-d <i>domainPath</i>	Required argument to specify the path to a global, non-partitioned, or local domain.
-detectDataLoss	Checks for data loss in the specified domain.
-detectCorruption	Checks for database corruption in the specified domain.
-repairCorruption	Repairs the database corruption in the specified domain. Note: This argument cannot be used with -detectDataLoss or -detectCorruption.
-backup	Optional argument to back up database files before attempting to repair them. Note: This argument can only be used with -repairCorruption.
-processes <i>maximumNumberOfProcesses</i>	Optional argument to specify the maximum number of processes to be started to repair Btree database corruptions. Note: This argument can only be used with -repairCorruption.
-loglevel <i>level</i>	Optional argument to set the logger verbosity level. Possible values: all, profile, debug, audit, information, warning, error, or none.
-noheader	Optional argument to disable the timestamp header.
-? -help -usage	Optional argument to get the usage text.

Scan Domain Data: scanDomain

Operational Utilities

This chapter describes the following operational utilities of RPAS:

- Find Alerts: [alertmgr Utility](#)
- Copying Domains: [copyDomain Utility](#)
- Move a Domain: [moveDomain Utility](#)
- Setting Miscellaneous Domain Properties: [domainprop Utility](#)
- Calculation Engine: [mace Utility](#)
- Managing the Workbook Batch Queue: [wbbatch Utility](#)
- Workbook Manager: [wbmgr Utility](#)
- Register Measure: [regmeasure Utility](#)
- Register Token Measure: [regTokenMeasure Utility](#)

Find Alerts: [alertmgr Utility](#)

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 when 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 {-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 8–1 provides descriptions of the arguments used by the alertmgr utility.

Table 8–1 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 -alerts or -categories list is not specified, findAlerts is run on all alerts. findAlerts can be run from either Master or Local Domains.
-alerts <i>a1 a2...</i>	Evaluate 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>	Evaluate 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	-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.

Note: alertmgr 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: copyDomain Utility

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 upgrading the domains 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 it copies each local domain into 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 8–2 provides descriptions of the arguments used by the copyDomain utility.

Table 8–2 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>	This argument allows copyDomain to copy each subdomain into user-instructed specific locations. This argument should not be used with -d OR -dest. See the copyDomain: Format of the XML Configuration File section 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>	Use this argument to copy a subset of a domain environment. Copies only positions specified in a format as dim1,pos1,...,posn:dim2,pos1,...,posn where the sequence dim1,pos1,...,posn specifies the selected positions along dim1. 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 config directory located in the source domains.

Table 8–2 (Cont.) Arguments Used by the copyDomain Utility

Argument	Description
-skipEmptyDir	Do not copy the empty directory located in the source domain.
-maxProcesses count	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 pathToDestDomain	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	Export 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	Compress 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: (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" 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: (PGRP100 INFO) : "ldom3" (PGRP200 INFO) : "ldom3" (PGRP21 INFO) : "ldom0" (PGRP22 INFO) : "ldom1" (PGRP300 INFO) : "ldom3" (PGRP31 INFO) : "ldom2" (PGRP32 INFO) : "ldom2"

copyDomain Notes

- By not providing the `-dest` or `pathToDestDomain`, the utility no longer makes the paths to the subdomains relative paths. Instead the `-relativizePaths` argument should be used.
- 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.
- `-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 sub-domains. 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 sub-domain. 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 will validate the configuration xml file first before any files are copied. If any of the subdomain source paths do not match a sub-domain path of the global domain being copied, a "can't find source sub-domain 'sub-domain' error will be report. If the global domain being copied contains any sub-domain that doesn't have a matching `srcPath` tag, a "sub-domain 'sub-domain' doesn't have a subdomain xml tag" error will be reported. If the global domain `srcPath` tag doesn't 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 will be 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 sub-domains to be copied. These options will still limit the number of sub-domains that are copied; however the configuration file must still contain entries for all domains.

Move a Domain: moveDomain Utility

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 will be simple and designed to fit the required task. It will contain fields for the locations of the source master domain and destination master domain as well as source and destination fields for each of the sub-domains that need to be moved.

moveDomain Usage

```
moveDomain -version
moveDomain -xmlConfigFile filename
moveDomain -d master -srcSubDomain src -dstSubDomain dst
```

[Table 8–3](#) provides descriptions of the arguments used by the `moveDomain` utility.

Table 8–3 Arguments Used by the moveDomain Utility

Argument	Description
<code>-xmlConfigFile pathToXmlConfigFile.xml</code>	This argument allows <code>moveDomain</code> to move a subdomain into user-instructed specific locations based paths specified in an xml file. This argument should not be used with the <code>-d</code> , <code>-srcSubDomain</code> , and <code>-dstSubDomain</code> parameters. See the moveDomain: Format of the XML Configuration File section for the file format.
<code>-d pathTomaster</code>	This argument will allow <code>moveDomain</code> to move each sub-domain based on the user-specified paths. Enter the path to the master domain.

Table 8–3 (Cont.) Arguments Used by the moveDomain Utility

Argument	Description
-srcSubDomain <i>src</i>	Path of the sub-domain to be moved.
-destSubDomain <i>src</i>	Path where the sub-domain is to be moved.

moveDomain: Format of the XML Configuration File

The XML configuration being used will be simple and designed to fit the required task. It will contain fields for the locations of the source master domain and destination master domain as well as source and destination fields for each of the sub-domains 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 will not be 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; i.e. 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 sub domain paths are required only for the domains that need to be moved.
- The sub domains that need to be moved can be specified or those sub domains that will remain under the master domain can be left out. If a sub domain 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 (just) master domain plus the size of the largest local domain.

Setting Miscellaneous Domain Properties: domainprop Utility

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 8–4 provides descriptions of the arguments used by the domainprop utility.

Table 8–4 Arguments Used by the domainprop Utility

Argument	Description
-d <i>pathToDomain</i>	Specifies the domain path.
-property <i>propertyname=value</i>	Used to specify the property to be changed. See the table of available properties below that can be set with this utility. To view the current property setting, use the property command with no value.
-reportSubDomains	Use this option to show property values for subdomains (should all match).

Available Properties

Table 8–5 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.

Table 8–5 (Cont.) Available Properties for the domainprop Utility

Domain Property Name	Type	Description
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.
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 need to be in "WebHelp" format. This path would normally be 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 be raised 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 the Translation Administration section.
meas_fillclr_precedence	Boolean	By default, when deciding which color to fill a particular cell with, RPAS grid uses the following order or 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 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 8–5 (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 they 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 that 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.

Calculation Engine: mace Utility

The mace utility (Multi-dimensional Array Calculation Engine) allows the administrator to evaluate rule groups or expressions in order to manipulate measures. mace 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:

1. The utility must be invoked on a master domain.
2. Parallelization is only applicable to single-expression evaluation (-run -expression argument). Parallelization does not apply to rule group evaluation.
3. The evaluated expression cannot be a SpecialExpression.
4. 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 running mace 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 / 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 8–6 provides descriptions of the arguments used by the mace utility.

Table 8–6 Arguments Used by the mace Utility

Argument	Description
<code>-d domainPath</code>	Specifies the domain in which to load the measure.
<code>-find string</code>	Use this argument to search all expressions for the specified string, printing all the rules and rule groups that have these expressions.
<code>-newRule {-ruleName ruleName}</code>	Use this argument to create a new empty rule. If desired, use the <code>-ruleName</code> argument to specify a name for the rule.
<code>-label {ruleLabel groupLabel}</code>	Use this argument to specify the label of the rule with the <code>-newRule</code> argument or label of the group with the <code>-newGroup</code> argument.
<code>-processes numProcesses</code>	Use this argument to specify the number of child processes to be run in parallel.
<code>-delRule ruleName</code>	Use this argument to remove the specified rule.
<code>-addRule groupName:ruleName</code>	Use this argument to add the specified rule to the group specified by <code>groupName</code> .
<code>-removeRule groupName:ruleName</code>	Use this argument to remove the specified <code>ruleName</code> from the group specified by <code>groupName</code> .
<code>-newGroup groupName</code>	Use this argument to create a new rule group with the specified name.
<code>-removeGroup groupName</code>	Use this argument to remove the specified group and non-shared rules in it.
<code>-addExpression ruleName</code>	Use this argument to add an expression to the specified rule. <code>-expression</code> should be used with this argument.
<code>-check</code>	Use this argument to validate the specified expression. <code>-expression</code> should be used with this argument.
<code>-run</code>	Use this argument to evaluate the specified expression or rule group. <code>-expression</code> should be used with this argument.
<code>-resolve groupName</code>	Use this argument to order (does not evaluate) expressions within rule group. Requires a comma-separated list of edited measures.
<code>-transit workbookName</code>	Use this argument to run 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.
<code>-print {ruleList groupList true}</code>	Use this argument to print all the specified rules and rule groups. The <code>ruleList</code> is a comma-separated list of rule names. The <code>groupList</code> is a comma-separated list of group names. If "true" is supplied for either <code>ruleList</code> or <code>groupList</code> , all rules or rule groups are printed.
<code>-purgeRules</code>	Use this argument to remove all rules not contained in any rule groups.
<code>-removeAllRuleData</code>	Use this argument to remove all rule groups and all rules.

Table 8–6 (Cont.) Arguments Used by the mace Utility

Argument	Description
-validate {calc general refresh}	Use this argument to validate 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	Use this argument to generate a file "mace.log" in the working directory for logging RuleEngine specific debug information.
-expression <i>exprString</i>	Use the argument to specify the expression. This argument is used in conjunction with the -addExpression, -check, and -run arguments.
-group <i>groupName</i>	Use this argument to specify the rule group to evaluate using the -run argument.
-measures <i>measureList</i>	Use this argument to specify the measures to resolve.
-group <i>groupList</i>	Use this argument to specify a list of group names, separated by commas. Use this argument in conjunction with the -transit and -print arguments.
-rule <i>ruleList</i>	Use this argument to specify 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	Use this argument to create a new rule group with the specified name

Managing the Workbook Batch Queue: wbbatch Utility

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 8–7 describes the arguments used by the wbbatch utility.

Table 8–7 Arguments Used by the wbbatch Utility

Argument	Description
<code>-d pathToDomain</code>	Specifies the domain containing the workbooks.
<code>-build queueIndex</code>	Runs workbook build for provided <i>queueIndex</i> .
<code>-refresh workbookName</code>	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.
<code>-commit workbookName</code>	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.
<code>-processes count</code>	Used with either <code>-build</code> or <code>-refresh</code> to build or refresh workbooks in the auto-workbook queue in parallel using the specified number of parallel processes.
<code>-scheduleRefresh</code>	Schedules a workbook to be refreshed later by adding it to the workbook refresh batch queue. If the <code>-category</code> option is specified, the scheduled workbook will be in that category. Otherwise, it will be in the default category.
<code>-unscheduleRefresh workbookName</code>	Removes a workbook from the workbook refresh batch queue.
<code>-scheduleCommit workbookName</code>	Schedules a workbook to be committed later by adding it to the workbook commit batch queue.

Table 8–7 (Cont.) Arguments Used by the wbbatch Utility

Argument	Description
-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.
-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>	Use this to update 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: <pre>wbbatch -d pathToDomain -queue build -updateCategories queueIndex1:newCatName1, queueIndex2:newCatName2</pre> Refresh example: <pre>wbbatch -d pathToDomain -queue refresh -updateCategories workbookName1:newCatName1, workbookName2:newCatName2</pre>

Workbook Manager: wbmgr Utility

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 8–8](#) provides descriptions of the arguments used by the wbmgr utility.

Table 8–8 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.
-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.

Register Measure: regmeasure Utility

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 will cause 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 properties can be modified, such as type, base intersection, 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} {-nvalue 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 8–9 provides descriptions of the arguments used by the regmeasure utility.

Table 8–9 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 8–9 (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 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.
-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 will not be 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 which 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 8–9 (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 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 will never be purged. clearPurgeAge is only available with the -modify option. When using the -modify option, -purgeage or -clearPurgeAge can be specified.

Table 8–9 (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 8–9 (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 na value as a regular value.

Table 8–9 (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.

Register Token Measure: regTokenMeasure Utility

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 will be 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

Will be evaluated as:

VM1 = a + b

The following expression is not valid, because TM1 is used without prefixing it with @:

TM1 = "sth"

If evaluated using mace, mace 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.

regTokenMeasure Usage

```
regTokenMeasure -version
regTokenMeasure -d pathToDomain -add tokenMeasure=valueMeasure {-fnhbi}
regTokenMeasure -d pathToDomain -list
regTokenMeasure -d pathToDomain -remove tokenMeasure=valueMeasure
```

Table 8–10 provides descriptions of the arguments used by the regTokenMeasure utility.

Table 8–10 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 will be stored in each local domain rather than the global domain, although by definition the token measure should be always be HBI measure since its 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.

Informational Utilities

There are numerous RPAS utilities that 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 to not make any changes to a domain or data in a domain.

- [Retrieving Domain Information: domaininfo Utility](#)
- [Checking the Validity of a Domain: checkDomain Utility](#)
- [Determining RPAS Server Version: rpasversion Utility](#)
- [List Contents of a Database: listDb Utility](#)
- [Printing Data from Arrays: printArray Utility](#)
- [Printing Data from Measures: printmeasure Utility](#)

Retrieving Domain Information: domaininfo Utility

The `domaininfo` utility is used to provide miscellaneous details about a domain, such as the type of domain (simple, master, or sub/local), and the upgrade/version history of the domain.

domaininfo Usage

```
domaininfo -d pathToDomain [Command]
domaininfo -expectedversion
```

Note: Domain path (-d) is required for all options except -expectedversion.

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

Table 9–1 Arguments Used by the `domaininfo` Utility

Argument	Description
-d	Path to the domain. Required for all options except -expectedversion.
-domainversion	Display the RPAS version of the specified domain.
-expectedversion	Displays the expected RPAS version of the domain that the utility expects to find.

Table 9–1 (Cont.) Arguments Used by the domaininfo Utility

Argument	Description
-type	Command to display 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/master domain of a global domain environment. A Sub domain 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 dims for subdomains.
-all	Displays the domain version, expected domain version, domain type, associated application, history, subdomains, and master domain info (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 there are 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.
-subdomain dim,pos	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.

Table 9–1 (Cont.) Arguments Used by the domaininfo Utility

Argument	Description
<code>-showrelativepaths</code>	When listing subdomains, indicates if paths are relative. Only relevant in combination with <code>-listsubdomains</code> or <code>-all</code> .
<code>-terse</code>	Removes the header information from the output, therefore returning only the requested information.

Checking the Validity of a Domain: checkDomain Utility

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}
```

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

Table 9–2 Arguments Used by the checkDomain Utility

Argument	Description
<code>-d pathToDomain</code>	Path to the domain that needs to be validated.
<code>-type expectedType</code>	Expected type of domain: simple, master, or sub.
<code>-q</code>	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 sub-domains for:
 - The sub-domain directory exists and is of type "sub"
 - If the master domain and the sub-domain have a repos directory
 - The measures, rules, rule groups, templates, and functions are the same in the global and sub-domain

If it is run on a sub-domain, 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 sub-domain.

Determining RPAS Server Version: rpassversion Utility

Use the `rpassversion` utility to determine which version of the RPAS Server is running in a particular location.

rpassversion Usage

```
rpassversion -l pathToLibrary
```

List Contents of a Database: listDb Utility

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
```

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

Table 9–3 Arguments Used by the listDb Utility

Argument	Description
<code>-db pathToDb</code>	Specifies the database to list the contents.
<code>-row</code>	List array information in a row format.
<code>-pageUsage</code>	Show btree page usage. Requires <code>-row</code> switch to be active.
<code>-standardOptions</code>	List only standard options.

Printing Data from Arrays: printArray Utility

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"} {-cellsprow num} {-nuposnames}
printArray -array db.array -allpopulatedcells {-format "formatString"}{-
cellsprow num} {-nuposnames}
```


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

Table 9–4 Arguments Used by the `printArray` Utility

Argument	Description
<code>-array db.array</code>	Specifies the array to print. Specify the full path to the database containing the array. Required for all commands except <code>-version</code> . <i>db</i> is a full or relative path to a database. Do not specify the <code>.ary</code> suffix. If no other commands are included, the array defaults to <code>-allpopulatedcells</code> with cells per row 1. The <code>-allpopulatedcells</code> command is still available, but now functions as a useful default action. The <code>-nposnames</code> , <code>-cellsperrrow</code> , and <code>-format</code> parameters may still be specified when relying on the implicit <code>-allpopulatedcells</code> behavior.
<code>-specs</code>	Prints the specifications of the array and positions along each dimension.
<code>-popcount</code>	Outputs only the <code>popcount</code> of the specified array. Useful to shell script writers to get the <code>popcount</code> value into a shell script variable. For example, export <code>POPCOUNT=`printArray -array hmaint.dim_year -popcount`</code>
<code>-cell CELLSPEC</code>	Prints a specific cell value from the array. Must not contain spaces. Must identify a single of 1-D slice. Specify using the format <code>"dim1:pos1,dim2:pos2,..."</code>
<code>-cellplain CELLSPEC</code>	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 <code>"dim1:pos1,dim2:pos2,..."</code>
<code>-slice CELLSPEC</code>	Prints a one-dimensional slice from the array. Must not contain spaces. Must identify a single of 1-D slice. Specify using the format <code>"dim1:pos1,dim2:pos2,..."</code>
<code>-allpopulatedcells</code>	Print all populated cells including the <code>navalue</code> of the array.
<code>-format "fmtstr"</code>	If <code>-format</code> 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
<code>-cellsperrrow num</code>	For multi-cell output commands (<code>-slice</code> and <code>--allpopulatedcells</code>), indicates how many cells should be printed on each line.
<code>-nposnames</code>	Suppresses the output of position names, only cell values are shown.

Printing Data from Measures: printmeasure Utility

Use the `printmeasure` utility to print measure information.

printmeasure Usage

```
printmeasure -d domainPath {-wb wbName} {-m measure} [COMMAND]
```

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

Table 9–5 Arguments Used by the `printmeasure` Utility

Argument	Description
<code>-d pathToDomain</code>	Specifies the domain that contains the measure to print. Requires the <code>-m</code> parameter.
<code>-m measure</code>	Specifies the measure to print.
<code>-wb workbookName</code>	Specifies the workbook associated with the measure to print. If <code>-wb</code> is not used, the domain measure information is printed. Requires the <code>-m</code> parameter.
<code>-list</code>	<code>printMeasure</code> will return a list of all registered measures in the domain. This argument does not require <code>-d domainPath</code> .
<code>-listHBIMeasures</code>	In a global domain, <code>printMeasure</code> will return a list of all measures registered at or above the partition dimension.
<code>-specs</code>	<code>printMeasure</code> returns the list of measure properties. Requires the <code>-m</code> parameter.
<code>-listDataIntersections</code>	<code>printMeasure</code> will return the base intersection of the measure

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 10-1](#), the language identifier is used for position labels. For more information, see the [Position Label Translation](#) section. The Windows Language ID is in the foundation.ini file. For more information, see the [Translation Administration](#) section.

Table 10–1 Supported Languages with Language Identifiers

Language	Language Identifier	Windows Language ID
Chinese (Simplified)	CHINESE_SIMPLIFIED	2052
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 the [Position Label Translation](#) section.

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 setup 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 10–1](#).

Or, you can set the following option:

```
[Options]
ResourceDLL=LANG.dll
```

LANG.dll represents a value in the Language Identifier column of [Table 10-1](#).

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).

Notes:

- The files will be 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/Traditional, Russian, you will need to 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 should 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 of 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 *RPAS User Guide*.

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.

Using Commit ASAP

After attempting to commit a workbook using Commit ASAP (File\Commit ASAP), the user will see a message in the client that the workbook has been scheduled for a commit. The user can continue with his/her work. The system will then try 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 using 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 checkbox **All Users** is not checked, the user will only be able to view his/her entries. If it is checked, you will see the entries for all users. The checkboxes in the **Status To Display** group allow you to filter the output so that the user sees 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.

Notes:

- 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 there can only ever be one pending Commit ASAP in the queue for a given workbook/user/template name combination.
- Workbooks must have been saved at least once before attempting a Commit ASAP. A workbook has not been saved if the label says "untitled."

Managing the Workbook Queue: 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 usage of this utility follows below.

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]*
```

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

Table 11–1 *showWorkbookQueues Arguments*

Argument	Description
<code>-version</code>	Prints the RPAS version, revision, and build information of the utility.
<code>-d domainPath</code>	Specifies the path to the domain.
<code>-show</code>	Lists the contents of the queue in the order in which the parameter is specified. Possible values: <code>all</code> , <code>pending</code> , <code>waiting</code> , <code>working</code> , <code>success</code> , and <code>failed</code> .
<code>all</code>	Used with the <code>-show</code> parameter. This lists all of the workbooks in all statuses.
<code>pending</code>	Used with the <code>-show</code> parameter. This lists all workbooks that are waiting to be committed.
<code>waiting</code>	For Oracle Retail development use only.
<code>success</code>	Used with the <code>-show</code> parameter. This lists all workbooks that have been successfully committed.
<code>failed</code>	Used with the <code>-show</code> parameter. This lists all workbooks that did not successfully commit.

Table 11–1 (Cont.) showWorkbookQueues Arguments

Argument	Description
<code>-purge date</code>	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. Administrator must select to purge commit processes that either succeeded or failed.

Commit ASAP Settings: 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 for 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 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]
```

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

Table 11–2 configCommitAsap Arguments

Argument	Description
<code>version</code>	Prints the RPAS version, revision, and build information of the utility.
<code>-maxProcs numProcs</code>	Sets the maximum number of concurrent commit processes where <code>numProcs</code> 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.
<code>-deadline time</code>	The time of the day when all outstanding commit ASAP operations will timeout. 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.

Table 11-2 (Cont.) configCommitAsap Arguments

Argument	Description
-display	Displays the current commit ASAP settings.
-loglevel level	Use this argument to set the logger verbosity level. Possible values: all, profile, information, warning, error, or none.
-noheader	To disable 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 the 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. If 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/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

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 need to be paired.

The following table shows examples of the formatting.

Table 12-1 CSV 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

This is a list of the standard log levels, controlled by the `-loglevel` option. Not all programs use these levels, but most do. 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 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. There should be no output 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 gets recorded in `rpas.log` under each user's working directory. Information related to domain activities, such as user sign-on and sign-off, gets 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. The table below displays the utilities that can use the `-loglevel` option.

Table 12–2 *loglevel Utilities*

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 indicated the sub-domain 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 APPNAMEYYYYMMDDHHMMbXX, where APPNAME is the utility name, 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 is two digits used to make the directory name unique. The framework will attempt 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 send 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. 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_DAEMON_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_Dyyyymmddhhmmbxx.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 and xx is some number used to make the file name unique.

The number of these log files will be limited to the number provided in the environmental variable RPAS_LOG_BACKUPS.

Otherwise, the log file name will be Daemon.log. Any existing log file is renamed to Daemon.old.

At midnight the current log file is closed and a new one opened, with naming as 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 not set then it defaults to logging messages at the warning level.

This utility creates log files in the domain/users/client directory, where domain is the current domain path and client is the current client. The actual file name used will be either rpas Dyyyymmddhhmmbxx.log or rpas.log base 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/date stamp, then the process id (pXXXX), and then a 2 digit number to avoid conflicts (bYY).

If there are any problems 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-partition hierarchies are reported as shown in the following:

```
<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 the loadHier entry 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/deregistration. Each process creates a log file in a newly created directory in the domain output directory. The newly created directory name will be 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 is two digits used to make the directory name unique. The RPAS libraries will attempt to create at most eight directories for any single application.

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 will be in the [DOM]/scripts/err/ directory. Examples of tools include Kornshell, Python, and Perl.

A Sample Shell Script

The following is a sample shell script that 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 will execute the script. In this example, it is defined to be the Korn shell. Therefore, this script will always be 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 12–3 Standard Arguments for RPAS Utilities

Argument	Description
-version	Use this argument to get the version information of the utility (for instance, RPAS 13.2.0). It does not require -d domainPath.
-d pathtodomain	Common to most utility this 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. Using this option will show the administrator what would change, but makes no actual changes to the system or data. See the usage of a specific utility to see whether this option is applicable.
-noheader	To disable the use of a timestamp in the header of the log file.
-help -? -usage	Any of these arguments will 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 when running a given utility. An administrator can set these levels for each RPAS utility. The available logger verbosity levels are as follows:

- none - There should be 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, that 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` sub-directory 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 will 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 will not be 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

Users that enter workbooks into the commit ASAP queue, while an exclusive lock is in place, will have the entries processed after the exclusive process is complete.

In these same domains, users will not be able to 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 a user tries to access one of these operations after the exclusive lock is obtained, a message is provided stating that an exclusive process is running. A default message is provided or it can be replaced by providing a message as part of the call to the ride utility.

When a user is 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 the Examples section for more details on the domain access during different types of ride processes.

Configuration functionality is provided so that a custom menu can be marked 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/or 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 would conflict with the ride process.

In order for a batch job to run over a domain without interference by 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 12–1 Domain Access Control

This figure 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 in order for the job to run. After the lock is received, no other workbook operations are able to get 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 Utility: `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}
```

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

Table 12–4 ride Utility Arguments

Argument	Description
<code>-d domain</code>	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.
<code>-process pname</code>	The name of the process to execute. This parameter cannot be used with the <code>-script</code> parameter.
<code>-script sname</code>	The name of the script to execute. This parameter cannot be used with the <code>-process</code> parameter.
<code>-args args</code>	Process arguments passed to the script or process to be executed. The <code>-args</code> parameter must be the last parameter or switch for this application. All parameters or switches after the <code>-args</code> parameter are passed on to the process or script to be started.

Table 12–4 (Cont.) ride Utility Arguments

Argument	Description
-message <i>messageString</i>	This optional argument is the override message presented to the user when 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>	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	Indicates that, when running over a global domain environment, the master domain will be 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 outlines several scenarios that are possible with the ride utility. The scenarios outline 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 is an example of running ride specifying only the master domain. This will lock all domains.

Usage example: `ride -d master -script script.ksh`

Table 12–5 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

Table 12–5 (Cont.) Running the ride Utility Specifying Only the Master Domain

Workbook Operation		Master Domain	Local Domain 1	Local Domain 2	Local Domain 3
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 is an example of running `ride` specifying only one local domain (local domain with partition position 100).

Usage example: `ride -d master -script script.ksh -partitions 100`

Table 12–6 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 is an example of running `ride` specifying only one local domain (local domain with partition position 100) and the `masterInBatch` option.

Usage example: `ride -d master -script script.ksh -partitions 100 -masterInBatch`

Table 12–7 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 is an example of running `ride` specifying two local domains (local domain 1 with partition position 100 and local domain 2 with partition position 200).

Usage example: `ride -d master -script script.ksh -partitions 100,200`

Table 12–8 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

Table 12–8 (Cont.) Running the ride Utility Specifying Two Local Domains

Workbook Operation	Master Domain	Local Domain 1	Local Domain 2	Local Domain 3
DPM	NA	Blocked	Blocked	Allowed

Scenario 5

This scenario is an example of running ride specifying two local domains (local domain 1 with partition position 100 and local domain 2 with partition position 200) and the masterInBatch option.

Usage example: `ride -d master -script script.ksh -partitions 100,200 -masterInBatch`

Table 12–9 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 is an example of running ride to lock the master domain.

Usage example: `ride -d master -script script.ksh -masterInBatch`

Table 12–10 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 Utility: 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 will 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 the cause for the ride process not running. This should provide them with enough information to either notify the user that they are 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
```

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

Argument	Description
<code>-d domain</code>	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.
<code>-autoRefresh refreshPeriod</code>	Refreshes the lock status information every number of seconds specified by the <code>refreshPeriod</code> .

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 *RPAS 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/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:


Add 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 the snap-in is not registered.

Figure 13–1 Console Manager Window Opened for First Time



2. To create a new work space, select **File** and then **New**.


Figure 13–2 Console Manager Window Showing Console Root Directory



To connect 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** to start the wizard to add a new snap-in.

Figure 13–3 Add/Remove Snap-in Dialog Box




4. To add a new snap-in, click **Add**.

Figure 13–4 Add Standalone Snap-in Dialog Box



5. Select **Oracle Retail RPAS ODBC Server Manager** and click **Add**.

Figure 13–5 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 13–6 Add Manager Dialog Box Showing Connection Information to be Added for the Agent

Add Manager

Agent
Please enter the connection information to your Agent.
Click Finish to complete the wizard.

Specify the host on which your Agent is running.
The host can be specified by either the host name or the host address.

Host:

The port is the TCP port on which your Agent is listening.
The TCP port must be in the range 1024 to 65534.

TCP Port:

Encrypted (SSL): ☐

< Back Finish 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.

Configure Using 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 13–7 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 13–8 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. The following figure shows the exact screen you should see.

Figure 13–9 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 the following figure, 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 13–10 New Attribute Dialog Box




- e. Click **OK**.
5. Expand Data Source Settings on the Console Manager window. You should find 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 are NOT to be modified.


- Under gdom, select IP Parameters.

Figure 13–11 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 13–12 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.

To connect to a domain on the local host, the following is an example of the path: "DOMAIN_PATH=C:\RPAS\11\Test\ODBC\nt_testGlobalDomain"

8. On the right panel of the window, double click the DataSourceIPCustomProperties attribute. The following window appears.

Figure 13–13 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.

Import 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 `<install_dir>\admin\oaadmin60.msc` for the file name. This will overwrite 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:

1. From the **Start** menu, select "ODBC Administrator" under the "Oracle RPAS odbc driver" menu item.

Note: Upon successful installation of RPAS ODBC Client, 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 13–14 ODBC Data Source Administrator Dialog Box



3. In the Create New Data source window, select **Oracle RPAS**.

Figure 13–15 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 13–1 lists the available custom properties (all optional):

Table 13–1 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 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 will present 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 will only contain columns for this dimension and its immediate parent dimension. If set to No, the dimension tables will contain columns for this dimension and all parent dimensions within the hierarchy.

Figure 13–16 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 13–17 Logon to Data Source Dialog Box



If the connection is successful, the following dialog box is displayed.

Figure 13–18 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 13–19 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 32 and 64-bit ODBC Client.

Note: For the remainder of this chapter, the 32 or 64-bit ODBC Client are referred to as ODBC Client, and `odbcclient32` or `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/odbc-trac.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.
 - ServerDataSource should be set 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 13–20 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, run the executable odbcisql. 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 13–21 Environment Variables Dialog Box



3. Select the **CLASSPATH** from the **System variables** list and click **Edit**. The Edit System Variable dialog box appears.

Figure 13–22 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**.

Where driver_home is the location where the jdbcclient was installed. If your jdbcclient was installed in C:/jdbcclient, you would have 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/Linux

On UNIX/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 RPAS JDBC Driver

Spy is a logging facility for JDBC driver. To enable spy for the RPAS JDBC connection:

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 13–2 lists the available attributes:

Table 13–2 Attributes Available

Key and Value	Description
log=System.out	Redirects logging to the Java output standard, System.out.
log=(file)filename	Redirects logging to the file specified by <i>filename</i> . For example, C:\temp\spy.log
linelimit=numberofchars	The maximum number of characters, specified by <i>numberofchars</i> , that Spy will log 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 / 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 RPAS JDBC Driver

Oracle Retail suggests that you use the jdbcisql.bat (for Windows) or jdbcisql.sh (for UNIX/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 -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.OpenAcce
ssDriver;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 by 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. Please note the double quotes must be included due to the semicolon.

Data Query

This section provides the details of data query, including the limitations, metadata, and dimension tables.

Limitations

It is important to 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) on 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 which 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 what is 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 13–23 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.

The fact and dimension tables have foreign key relationships between them to represent the intersection and maintain data integrity between the dimensions and the facts. Dimension tables have foreign key relationships with other dimension tables to represent the hierarchical relationships between them.

Figure 13–24 Example of Star - Denormalized Schema to Represent Facts and Dimensions in RPAS



At connection time, all intersections at which any measure is stored at its base level are available as fact tables within the database. Additional aggregate level intersections may be made available in the database by specifying them in a custom connection property. These fact tables are a part of the set of database entities that will be visible to reporting tools at connection time.

However, the RPAS ODBC/JDBC driver supports dynamic aggregate level fact tables that can be queried even though they are not available at connection time. These tables include all intersections that are logically above the base intersection fact tables and have at least one measure in them when manifested. If the measure existence condition is not met, the driver returns an error that the fact table could not be found.

These dynamic fact tables are queried in the same fashion as the tables that are available at connection time. The name of the fact table can be constructed by piecing together dimension names (not labels) that make up the intersection in the order in which they would exist within the domain. For example, if someone wants to query facts at the store/class/day level but the fact table is not available at connection time, they can construct the fact table name as: FACT_STR_CLSSDAY_. Note that dimension names have been concatenated in the same order as the intersection and have been prefixed with 'FACT_'. Also, note that a dimension name is assumed to be four 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

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. To understand the structure of a particular hierarchy (for example: CLND) the following table will list 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

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

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

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', the value of this custom property would be '`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), the property key value pairs must be separated by a semicolon. Using examples above, the content of the custom properties input box would appear 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 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 (re)starting the OBIEE server.

Configuring the ODBC Client for OBIEE

The following example provides a sample of configuring 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 `odbc 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 enter 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. Select **ODBC 3.5** for Connection Type, choose the appropriate DNS for the RPAS domain, enter the RPAS user name and password, and then click **OK**.

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 OBIEE documentation.

Microsoft Access

To connect using Microsoft Access:

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/tables that you can import or link.

- 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 Microsoft Access displays a list of the fields in the linked table. Select a field, or a combination of fields, that will 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:


- Start JDeveloper.
- On the JDeveloper left panel, select the **Connections** tab.
- Right-click on **Databases**, and select **New Database Connection**.
- On the first screen of the Create New Database Connection wizard, enter a connection name and select **Third Party JDBC driver** for Connection Type.
- On the second screen, enter the user name and password, and then click **Next**.
- 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 which is configured on the ODBC server.

Figure 13–25 Register JDBC Driver Dialog Box



- Follow the instructions to finish creating the connection.

After the connection is established, a user can expand the Connection (and the nodes under the Connection) to browse the objects in the RPAS datastore. A user can also open a SQL worksheet (by selecting **SQL Worksheet** from the Tools menu) to write/execute SQL statements.

XML Publisher

This section describes how to make the connection from XML Publisher to RPAS using 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 and 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 which 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 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 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/doe@rpasDomain
```

Where `john/doe` 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*doe@rpasDomain
```

After connected, users 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 functionalities that are not supported.

Supported SQL Functions

Use extreme 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 13–3 Numeric Functions

Function	Description
<code>ABS(<i>numeric_exp</i>)</code>	Returns the absolute value of <i>numeric_exp</i> . For example: <pre>SELECT ABS(-1.0), ABS(0.0), ABS(1.0) FROM emp WHERE empno = 1;</pre> This returns 3 result columns with values 1, 0, and 1.
<code>ACOS(<i>float_exp</i>)</code>	Returns the arccosine of <i>float_exp</i> as an angle, expressed in radians. For example: <pre>SELECT ACOS(-1) FROM emp WHERE empno = 1;</pre> This returns 3.14159.
<code>ASIN(<i>float_exp</i>)</code>	Returns the arcsine of <i>float_exp</i> as an angle, expressed in radians. For example: <pre>SELECT ASIN(-1.0) FROM emp WHERE empno = 1;</pre> This returns -1.57079.
<code>ATAN(<i>float_exp</i>)</code>	Returns the arctangent of <i>float_exp</i> as an angle, expressed in radians. For example: <pre>SELECT ATAN(45.0) FROM emp WHERE empno = 1;</pre> This returns 1.54857.
<code>ATAN2(<i>float_exp1</i>, <i>float_exp2</i>)</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: <pre>SELECT ATAN2(35.175, 129.44) FROM emp WHERE empno = 1;</pre> This returns 0.2653399.
<code>CEILING(<i>numeric_exp</i>)</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: <pre>SELECT CEILING(123.45), CEILING(-123.45), CEILING(0.0) FROM emp WHERE empno = 1;</pre> This returns 124, -123 and 0.
<code>COS(<i>float_exp</i>)</code>	Returns the cosine of <i>float_exp</i> , where <i>float_exp</i> is an angle expressed in radians. For example: <pre>SELECT COS(14.78) FROM emp WHERE empno = 1;</pre> This returns -0.59946542.
<code>COT(<i>float_exp</i>)</code>	Returns the cotangent of <i>float_exp</i> , where <i>float_exp</i> is an angle expressed in radians. For example: <pre>SELECT COT(124.78) FROM emp WHERE empno = 1;</pre> This returns -0.82045588.
<code>DEGREES(<i>numeric_exp</i>)</code>	Returns the number of degrees converted from <i>numeric_exp</i> radians. For example: <pre>SELECT DEGREES(3.143) FROM emp WHERE empno = 1;</pre> This returns 180.0806.
<code>EXP(<i>float_exp</i>)</code>	Returns the exponential value of <i>float_exp</i> . For example: <pre>SELECT EXP(378.615) FROM emp WHERE empno = 1;</pre> This returns 2.69404760606322E+164

Table 13–3 (Cont.) Numeric Functions

Function	Description
FLOOR(<i>numeric_exp</i>)	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(<i>float_exp</i>)	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(<i>float_exp</i>)	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(<i>integer_exp1</i> , <i>integer_exp2</i>)	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(<i>numeric_exp</i> , <i>integer_exp</i>)	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(<i>numeric_exp</i>)	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([<i>integer_exp</i>])	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 13–3 (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 <i>code</i> as a SQL_WCHAR value. The value of <i>code</i> should be between 0 and 65535. Example: "SELECT NCHAR(945)" returns the character α.

String Functions

Table 13–4 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 13–4 (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 13–4 (Cont.) String Functions

Function	Description
<code>LEFT(string_exp, count)</code>	Returns the leftmost count characters of string_exp. If string_exp is wide character string, the return value is a wide character string. Offset (count) 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 ename is a column of wide character data type, this returns wide character string N'jo'.
<code>LENGTH(string_exp)</code>	Returns the number of characters in string_exp, excluding trailing blanks. If string_exp is wide character string, the return value is a number of wide characters in string_exp. 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.
<code>LOCATE(string_exp1, string_exp2[, start])</code>	Returns the starting position of the first occurrence of string_exp1 within string_exp2. The search for the first occurrence of string_exp1 begins with the first character position in string_exp2 unless the optional argument, start, is specified. If start is specified, the search begins with the character position indicated by the value of start. The first character position in string_exp2 is indicated by the value 1. If string_exp1 is not found within string_exp2, the value 0 is returned. If string_exp2 is a wide character string, returns the starting position of the first occurrence of string_exp1 within the wide character string string_exp2. Offset (start) must be specified in number of characters. If string_exp2 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.
<code>LTRIM(string_exp)</code>	Returns the characters of string_exp, with leading blanks removed. For example: <pre>SELECT LTRIM(' ABC') FROM emp;</pre> This returns 'ABC'.
<code>OCTET_LENGTH(string_exp)</code>	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 13–4 (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 13–4 (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('αβγ')" returns an integer value of 945.

Time / Date Functions**Table 13–5 Time / 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 13–5 (Cont.) Time / 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 13–5 (Cont.) Time / 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 13–5 (Cont.) Time / 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.
WEEK(<i>date_exp</i>)	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.
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 13–6 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 13–7 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 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

The following table defines the precision, length, and scale keywords of the CAST function.

Table 13–8 *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
CHAR, BINARY	255

Table 13–8 (Cont.) CAST Function

Keyword	Value
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

The following tables define the length, precision, and scale keywords of the CONVERT function.

Table 13–9 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 13–10 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 13–11 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 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/Time Functions

The OpenAccess SDK SQL engine does not support the following Date/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:

1. MACE: all left hand side measures in the expression or expressions of the rule group specified in the command line.
2. Workbook Commit: all left hand side measures in the commit rule group.
3. loadMeasure: all successfully loaded measures.
4. loadHier: measures whose base intersection contains at least one affected dimension. "Affected" means adding and deleting one or more logical positions.
5. 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.

The following diagram demonstrates a JMS system in the context of RPAS. The JMS system consists of four basic components:

- JMS Provider: The central daemon which 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 14–1 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/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, consists 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 their own 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 interprocess 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 the following table. The macros are replaced by actual values at run time.

Table 14–1 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.

Below 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 [...]]]
```

The following table provides descriptions of the arguments used in the command line.

Table 14–2 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 domainPath	Path to the monitored domain. It must be a simple or master domain.
-c configFileName	Configuration file in Java properties file format. This argument is optional.
name1=value1 name2=value2 ...	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

The following table lists all vendor neutral configuration settings for the JMS Publisher:

Table 14–3 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 needs to 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 14–4 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.RefFSContextFactory
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 14–5 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 14–5 (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.

The following 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 trouble-shooting. 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 [...]]]
```

- {classpath} should include rpaspmc.jar and jms.jar under \$RPAS_HOME/lib and any vendor-specific JMS implementation jar files.
- {configFileName} is a configuration file in Java properties file format. It is optional.
- name/value pairs - If the configuration file is not specified, the name/value pairs are used. 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 the [Configuring the RPAS JMS Publisher](#) section for details.

The following table lists all the general settings for the JMS Subscriber:

Table 14–6 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.

The following 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 a planning centered, integrated UI solution for responding to events by adjusting the affected plans within the context of the changed circumstances and interacting plans. A planner's workspace can be developed where 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, that 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 Oracle Retail Workspace or other Web portal.

Launching RPAS

A system integrator is responsible for creating a Workspace or Web-based application that can launch RPAS. Creating a Workspace application involves the following steps:

1. Configure the navigational pane in Workspace to appropriately 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, 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, which 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. Find 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 preinstalled client by setting a parameter for the applet.
2. Download and install 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. Start the RPAS client.
4. Route the ICL request to the RPAS server through the RPAS client and wait 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 15–1 Attributes Required in the `<Applet>` Tag

Attribute	Description	Value
code	Applet class name	com.retek.mdap.client.launch.LaunchApplet.class
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 15–2 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
launchPreinstalledOnly	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 is 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.

The following figure 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 15–1 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]);
```

- `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 the [User Authentication](#) section.

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)
{
  ...
}
```

- `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, 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 Active Content, RWL Applet, and 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. The figure shows the case where two RPAS client processes are launched.

Figure 15-2 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 will continue to function, but its status will not be 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 which is generally at the bottom of the browser window.

[illegible]

```

CQkJCQk8ZGltIG5hbWU9IiNUUiI+DQoJCQkJCQkJPHNlbGVjdGlvbj4wMTAyPC9z\
ZWxlY3Rpb24+DQoJCQkJCQkJPHNlbGVjdGlvbj4wMTUyPC9zZWxlY3Rpb24+CQkJ\
CQkJDQoJCQkJCQkJPHNlbGVjdGlvbj4wMTUyPC9zZWxlY3Rpb24+CQkJCQkJDQoJ\
CQkJCQkJPHNlbGVjdGlvbj4wNTU3PC9zZWxlY3Rpb24+CQkJCQkJDQoJCQkJCQkJ\
PHNlbGVjdGlvbj4wNTk0PC9zZWxlY3Rpb24+CQkJCQkJDQoJCQkJCQkJPHNlbGVj\
dGlvbj4wOTU5PC9zZWxlY3Rpb24+CQkJCQkJDQoJCQkJCQkJ8L2Rpb24NCgkJCQkJ\
PC92YWxlZT4NCgkJCQkJ8L1N1cGVyVHJlZT4NCgkJCTwvV2l6YXJkUGFnZT4NCgkJ\
CTxXaXphcmRQYWdlIGlkPSJTS1UiIGZhbGxiYWNRPSJub251Ij4NCgkJCQkJ8U3Vw\
ZXJUCmVlIG5hbWU9InRyZWUxIj4NCgkJCQkJPHZhbHVlPg0KCQkJCQkJPGRpbSBu\
YWllPSJDTFNTIj4NCgkJCQkJCQkJ8c2VsZWN0aW9uPjAyMTwvc2VsZWN0aW9uPg0K\
CQkJCQkJPC9kaW0+DQoJCQkJCTwvdmFsdWU+DQoJCQkJPC9TdXB1clRyZWU+DQoJ\
CQkJ8L1dpemFyZFBhZ2U+DQoJCTwvV2l6YXJkUGFnZXM+DQoJCTxzYXZlPg0KCQkJ\
PGxhYmVsPlRlc3R3YjwvbGFjZWw+DQoJCQkJ8YWNjZXRzPkdib3VwPC9hY2Nlc3M+\
DQoJCQkJ8Z3JvdXA+QURNSU48L2dyb3VwPg0KCQkJ8L3NhdmU+DQoJPC9idWlsZD4N\
CjwvV29ya2Jvb2s+DQo8L1JQQVNMXXVvY2hDb250ZXh0Pg0K";
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>

```

The following figure shows this sample code rendered by a Web browser. Note that the status bar shows the current status of the RPAS client.

Figure 15–3 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 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, a user name and password is 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.

The following figure shows the flow of OSSO support for ICL.

Figure 15–4 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 lifecycle of the launch contextual data.

- Do not pass the user name and password to the applet. The RPAS client will prompt 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

The following diagram 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 Workspace or 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 15–5 ICL Deployment Diagram



Use Cases

This section lists 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 Workspace 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 datetime of the workbook. These parameters can be specified in the XML context specification to have RPAS open the workbook. For information, see the "Open Workbook" section.

To enable opening the latest workbook from a nightly auto-workbook-build batch, system integrators can provide the label and template name. ICL will automatically default 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 Workspace or other Web application. The following examples illustrate this concept.

- If a link for "Merchandise Financial Planning" is clicked 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 had drilled 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 needed 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 15–3 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

Server Settings

Table 15–4 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 15–5 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

Table 15–5 (Cont.) Domain Settings

Setting	Description	Type	Required
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 the [Open Workbook](#) section. For information on building workbook, see the [Build Workbook](#) section.

XML Schema

The following 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:complexType>

</xs:schema>
```

```

        </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 réf="WizardPages" />
      <xs:element réf="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: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: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 reacting to an event or alert related to these products and locations, it may be necessary to always open a particular workbook or auto-build workbook. System integrators configuring a Workspace 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 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

There are two options that can be specified when a workbook is opened.

Option 1

Table 15–6 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 15–7 Option 1- 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 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 15–8 Basic Settings

Setting	Description	Type	Required
templateName	The template name for the workbook to be built.	String	Yes
allowManualOv erride	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

Table 15–8 (Cont.) Basic Settings

Setting	Description	Type	Required
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 15–9 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 15–10 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 15–11 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

Table 15–11 (Cont.) Wizard Widget Specifications (2)

Edit, DropDownList, SingleHierSelect, and DatePicker

Setting	Description	Type	Required
default	Default value for the widget. Must be a String type.	XML section	No

Table 15–12 Wizard Widget Specifications (3)

Listbox and Tree

Setting	Description	Type	Required
name (attribute)	Widget name.	String	Yes
fallback (attribute)	Widget level fall back 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 15–13 Wizard Widget Specifications (4)

SuperTree

Setting	Description	Type	Required
name (attribute)	Widget name.	String	Yes
fallback (attribute)	Widget level fall back 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
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

Table 15–13 (Cont.) Wizard Widget Specifications (4)

SuperTree			
Setting	Description	Type	Required
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 15–14 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 />
    <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>
<save>
    <label>Testwb</label>
    <access>Group</access>
    <group>ADMIN</group>
</save>
</build>
</Workbook>
</RPASLaunchContext>

```

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 listbox 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]
```

- `[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.

Workspace and Oracle Single Sign-On

For information about Workspace and Oracle Single Sign-On (SSO), see the *Oracle Retail Predictive Application Server Installation Guide*.

Appendix: Curve Administration Guide

This chapter provides the details of the utilities used in Curve administration.

curvevalidate

`Curvevalidate` automatically executes during the domain install, and it can also be run at any time against a Master 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 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]
```

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

Table A–1 Arguments Used by the `curvevalidate` Utility

Argument	Description
<code>-d domainpath</code>	Path to the domain.
<code>-s</code>	Set defaults.
<code>-debug</code>	This argument causes temporary measures to be retained for debugging purposes.
<code>-h</code>	Standard information argument to display information about <code>curvebatch</code> in the terminal screen.
<code>-version</code>	Standard argument that displays version information.
<code>-loglevel</code>	Used to set the logger verbosity level. Available verbosity levels are as follows: ■ <code>all</code> ■ <code>profile</code> ■ <code>information</code> ■ <code>warning</code> ■ <code>error</code> ■ <code>none</code>
<code>-noheader</code>	Standard argument to disable 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).
 - There must be at least one common hierarchy 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 the details of the curvebatch utility.

curvebatch Usage

```
curvebatch -d domainpath [-level # ] [-debug] | -h | -version
```

The following table 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>	Path to the domain.
-level #	The # signifies the profile ID. When using level, a valid profile ID must be provided.
-debug	This argument causes temporary measures to be retained for debugging purposes.
-h	Standard information argument to display information about curvebatch in the terminal screen.
-version	Standard argument that displays version information.
-loglevel	Used to set 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, which 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 the [Translation Administration](#) section.) 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)
```

The following table 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 will be 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: Freeing the char buffer is necessary to prevent memory leaks.

The functions `rtd_login`, `rtd_login_ex`, and `rtd_send` indicate success by returning 1 and failure by returning 0.

The following figure 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 will be 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 checkboxes.

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 will be added to.

measure - The name of the "parent" measure for the new percent of parent measure.

hier [optional] - Specifies the hierarchy name for a relative percent of parent measure.

hiers [optional] - Specifies all hierarchies for an absolute percent of parent measure. Hierarchy names are delimited by a space.

dims [optional] - Specifies 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 will be relative. Specifying both 'hiers' and 'dims' indicates that the percent of parent measure will be 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/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 needs to reference a header which 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 advised 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 which 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 should 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 should be prefaced by the full path.

Script Section: Action (simple test)

The Actions section contains user actions; these will 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 should 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 should be identical to the prior example.

The Actions section should 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 should 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: 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/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_CACHE](#)
 - [RPAS_PAGE_SIZE](#)
 - [RPAS_PAGE_SPLIT_PERCENTAGE](#)
 - [RPAS_LOCK_TIMEOUT](#)
 - [RPAS_USER_MODE](#)
- [Log Settings](#)
 - [RPAS_LOG_BACKUPS](#)
 - [RPAS_LOG_LEVEL](#)
 - [RPAS_DAEMON_LOG_PATH](#)
- [Profiling Logging](#)
 - [RPAS_PROFILING_ENABLE](#)
 - [RPAS_PROFILING_PATH](#)

- [Date and Time Setting](#)
 - [RPAS_TODAY](#)
 - [RPAS_TODAY_STATIC](#)
- [Numeric Precision](#)
 - [RPAS_INCAGGPRC](#)
 - [epsilon](#)
- [RPAS_PROCESSES](#)
- [LDR_CNTRL=NAMEDSHLIB=RPASZONE](#)

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 *RPAS 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_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. Although, 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 composed 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 page 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 environmental 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/save the same data back to the domain. 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 environmental 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 will appear as a warning dialog box only after some new action is initiated. With the client, there will be additional warning dialog boxes 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 environmental variable that RpasDbServer uses, this environmental variable must be set prior to starting the DomainDaemon.

The environmental 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, which means that there is no need to rapidly close databases.

```
export RPAS_USER_MODE="single"
```

Value: single

Log Settings

The following environment variables can be used for log settings.

RPAS_LOG_BACKUPS

The RPAS_LOG_BACKUPS environmental 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 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_LEVEL

The RPAS_LOG_LEVEL environmental 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_DAEMON_LOG_PATH

The RPAS_DAEMON_LOG_PATH environmental 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_DAEMON_LOG_PATH=" C:/RPAS/Domains/Temp"
```

Value: The absolute path to the directory where RPAS daemon applications creates log files

Profiling Logging

The following two environment variables may be setup 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 general RPAS log file, which is controlled solely by loglevel.

RPAS_PROFILING_PATH

The `RPAS_PROFILING_PATH` environment variable defaults to `rpasProfile.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
```

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.

- 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` will cause 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

There are two variables that 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.

Below is an example of setting this variable to a value smaller than the default.

```
export RPAS_INCAGGPRC=0.000000001
```

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_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 results in all processes being 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

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 would 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 would be difficult to fill up since it is used only by RPAS processes.

```
export LDR_CNTRL=NAMEDSHLIB=RPASZONE
```