Oracle® Retail Predictive Application Server

Administration Guide for the Fusion Client

Release 13.2.2

January 2011

ORACLE®

Oracle® Retail Predictive Application Server Administration Guide for the Fusion Client, Release 13.2.2

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

3 Domain Administration

4 Setting Up the RPAS Fusion Client

5 User Maintenance

6   System Administration

7 Hierarchy Management

8 Data Management

11   Internationalization

12   Commit as Soon as Possible

13   Batch Processes and RPAS Utilities

14   RPAS ODBC/JDBC Driver

15   Publishing Measure Change Events

A   Appendix: Curve Administration Guide

B   Appendix: Manually Configuring Single Sign-On

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 13.2.2 documentation set:

- *Oracle Retail Predictive Application Server Release Notes*

- *Oracle Retail Predictive Application Server Installation Guide*

- *Oracle Retail Predictive Application Server Administration Guide for the Classic Client*

- *Oracle Retail Predictive Application Server User Guide for the Fusion Client*

- *Oracle Retail Predictive Application Server Online Help for the Fusion Client*

Customer Support

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

https://support.oracle.com

When contacting Customer Support, please provide the following:

- Product version and program/module name

- Functional and technical description of the problem (include business impact)

- Detailed step-by-step instructions to re-create

- Exact error message received

- Screen shots of each step you take

Review Patch Documentation

When you install the application for the first time, you install either a base release (for example, 13.1) or a later patch release (for example, 13.1.2). 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
```

Send Us Your Comments

RPAS Administration Guide for the Fusion Client, Release 13.2.2

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`.

1

Introduction

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

Administrator Overview

After the RPAS Server and Fusion Client have been installed, administrators must set up the RPAS Fusion 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
- Dimensions
- 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. Dimensions 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.

Dimensions

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

RPAS supports many alternative dimensions 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.

Refresh and Export Data

The Refresh feature allows the user to update a workbook with the data that is currently stored in the domain.

Refresh the Data in a Workbook The Refresh feature lets the user work with the most current data without having to rebuild the workbook. However, refresh rules in the configuration governs the refresh process. Please see the *RPAS Configuration Tools User Guide* for more information on setting up refresh rules.

Export the Current View to an Output File

Navigate: On the Global toolbar, in the File menu, click **Export**. The Export window appears.

For more information on exporting the current worksheet view to an output file, refer to the *RPAS User Guide for the Fusion Client*.

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.

Overview of the RPAS Fusion Client

The RPAS Fusion Client is the Web-based Rich Client for the Retail Predictive Application Server (RPAS) platform developed using the latest Oracle Application Development Framework (ADF). It includes all the features available in the RPAS Windows-based Client and delivers an enhanced user experience that meets the performance and scalability requirements set for the RPAS platform.

Planning is one of the most important and complex processes in a retail business. It typically involves a complex set of activities that need to be followed as part of a workflow. Unlike the RPAS Windows-based client, the RPAS Fusion Client includes an Activity Task Flow feature that provides a robust workflow capability to make each planning activity easier to track and maintain.

The RPAS Fusion Client uses the same RPAS server as the current RPAS Windows-based client. You can continue to configure an RPAS solution and connect to it using the Web-based interface or the RPAS Windows-based client. Apart from the enhanced user experience, the Fusion Client provides access to a larger number of users and a certain degree of platform independence.

Terminology Differences Between Clients

There are some key terminology differences between the Fusion Client and Classic Client. Understanding these differences is useful if you use both clients. The following table describes the differences.

Table 1–1   Terminology in Fusion Client and Classic Client

Fusion Client Term	Term Description	Classic Client Term
Dimension	Grouping of a particular type of information. Typical dimensions are for products, locations, time, and measures. For instance, a Product dimension could contain information about items, item groups, departments, and divisions.	Hierarchy
Level	A subdivision of a dimension. Levels group information of the same type. For instance, a level within the Product dimension could be Department. The Department level would contain all the departments (men's shoes, women's shoes, children's shoes) that exist.	Dimension

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.

2

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.

- MKS Toolkit 8.7 Developer's Toolkit 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 following example, 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 1.5 or 1.6. Make sure 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 Environment Variables section 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 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.

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
`[-fullinstall | -patchinstall | -testinstall]`	Required argument. Indicates the type of installation to be performed, where: ■ `-fullinstall` - Builds a full domain and loads the hierarchy data files. ■ `-patch install` - Patches an existing domain. Updates or unregisters/registers measures that have changed (as necessary). ■ `-testinstall` - Used for testing only. Only generates configuration files.
`-ch` *config_home*	Required argument. Path to the directory containing the configuration file.
`-cn` *config_name*	Required argument. The name of the configuration.
-in *input_home*	Required argument. Path that includes the directory containing the input files for the domain to be created. If this directory includes a users.db file, then the users and groups within that file are added upon domain creation. For more information, see Managing Users - usermgr Utility.
`-log` *log_name*	Required argument. Path that includes the name of the log file to be created or updated.
`-configdir` *config_directory*	The path to the directory containing the xml files used by RPAS. This is a required argument if the user wants to supply globaldomainconfig.xml or calendar.xml.
`-dh` *domain_home*	The path to directory in which the domain will be created. Use if and only if a globaldomainconfig.xml is not used.
`-p` *dim_name*	The partitioning dimension. Use if and only if global domain is being implemented without the use of globaldomainconfig.xml.
`-rf` *function_name*	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.
`-updatestyles`	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

Once 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 Fusion 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 Fusion Application Server, 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. Once you have created one user, you can use the Add a User step 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.

Additional Procedures

Refer to other sections of this guide for the following procedures:

- Loading data using the loadmeasure utility: Loading Measure Data - loadmeasure.

- Disabling the timestamp header: Common Information and Parameters for RPAS Utilities.

- Information on Measure Loading Properties: Arguments Used by the loadmeasure Utility.

Upgrading and Patching Domains

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

> **Note:** See the solution-specific installation and administration guides for potential solution upgrade limitations.

This similar process should also be followed when upgrading between minor (patch) releases of RPAS (for example, RPAS 13.1.1 to RPAS 13.1.2). The following overview of this process calls out upgrade steps that are unique to upgrading to a major or minor release.

The instructions below describe the process that must be followed to upgrade an RPAS domain and configuration.

1. Acquire the new version of the platform, which includes the following key components:

   - RPAS Server

   - RPAS Client

   - RPAS Configuration Tools

2. Create a backup copy of any domains that will be upgraded.

3. Commit any existing workbooks that have not yet been committed (if commit is required).

   RPAS does not guarantee the upgrade of existing saved workbooks or styles as part of the upgrade process from one major release or patch to the next. If users are unable to build workbooks after upgrading a domain, system administrators must delete all the existing style and user/selections folders to correct the problem.

4. Remove any `-commit` and `-refresh` scheduled batch jobs from the workbook batch queue (`wbbatch` utility). If upgrading an 11.0.x environment to 12.1.x (or later), all `-build` scheduled batch jobs should also be removed from the queue.

   > **Note:** Skip this step if upgrading a 12.1 or newer environment.

5. Run the `convertDomain` utility on existing domains. In global domain environments this utility is only executed from the master domain. See the convertDomain section for more information.

   > **Note:** Skip this step if upgrading a 12.1 or newer environment.

6. Run the `upgradeDomain` utility on existing domains. In global domain environments this utility is only executed from the master domain. See the upgradeDomain section for more information.

7. Upgrade existing configurations. See the Upgrade Configurations section for more information on this step.

> **Note:** Skip this step if upgrading from 13.0 release or higher.

8. Run the `rpasInstall` utility with the `-patchinstall` argument. See the rpasInstall section for more information.

> **Note:** When using the `-updatestyles` flag, deleting existing styles is not necessary.
>
> After running the `upgradeDomain` utility, the patch installation (through `rpasInstall`) must be run over the domain even if there are no configuration updates. This is done to ensure that configuration within the domain is synchronized with any changes made as the result of the RPAS and Configuration Tools upgrade.

Patching a domain does not update any hierarchy or dimension changes in the configuration with the exception of user-defined dimensions. It does not patch changes to the hierarchy purge age, the security dimension, 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.

Patching a domain allows for any changes to measure properties, rule sets, rule groups, and workbook templates. During the process, the rule sets, rule groups, and rules are completely rebuilt, but measures and templates are updated with the changes.

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

Patching Measure Properties

Certain measure properties cannot be modified without unregistering and reregistering measures, which results in the loss of measure data. Updateable measure properties are those that can be modified, and not updateable measures are those cannot be modified. If all of the measure properties that are changed are able to be modified, simply update the measure properties, and measure data will not be lost in the domain during the patch process. If even one of the measure properties that change are not able to be modified, the patch process results in un-registration and re-registration of the measure. Therefore, the data that was in the domain for that measure before the patch process will not be there after the patch process is complete. For more information, see the Register Measure - regmeasure Utility section.

convertDomain

Migrating data from 11.0, 11.1, and 12.0 versions of RPAS to RPAS 12.1 (or newer versions) requires the convertDomain utility. It contains the necessary functions to read the pre-12.1 database and array formats and writes to the new 12.1 formats. It processes both simple and global domains, regardless of subdomain locations, optionally removes the previous .gem files, and recovers from a partial conversion. Upon successful completion, it will tag the domain as converted, which is visible in the domain history and visible with the domaininfo utility.

Usage

```
convertDomain -d domain [-loglevel level] [-purge] [-forcetag]
```

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

Table 2–2 Arguments Used by the convertdomain Utility

Argument	Description
-d *path*	Path to the domain being converted.
-purge	Inclusion of this argument indicates that existing (Acumate) .gem database files should be removed after conversion to reduce the space required for the process. The .gem file will be removed immediately after it has been converted.
-forcetag	This argument indicates that the domain should be tagged as converted even if some databases/arrays could not be converted.

The conversion process is done directly to the simple or global domain path specified with the -d argument. A copyDomain is **not** automatically performed beforehand by the utility.

The process works by first opening the R_SUBDOMAINPATH%1 array in the data/configmeasdata database. If this fails then it is assumed that the domain is a simple domain and not a global one. If it succeeds then the subdomain paths are read from the array. Next the process searches for all .gem files in the subdomains and the master, recursively and converts each one. The order of the conversion should not be considered predictable; there is no assurance that subdomains will be converted in order, or that the master will be converted first or last. In order to ensure recoverability the data/configmeasdata database is not removed until the entire domain has been converted. Finally, if no errors occurred during conversion (or the -forcetag option is supplied) then the new domain is tagged with the CONVERT tag and time stamped.

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

In the case of a partial domain conversion it is possible to restart convertDomain. Currently every database will be reconverted unless the -purge option is specified. In that case the conversion will restart with the last database that was being converted. There is no current way to skip converted databases if -purge was not used, nor can individual databases be converted.

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.

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` *path*	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`	Required 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
> ```

Upgrade Configurations

Once the RPAS Configuration Tools have been installed and the domains have been upgraded, an administrator should upgrade existing configurations to version 13 or the latest version 13 patch.

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 2–2   Converting the Configuration

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

Figure 2–3   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 2–4   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 RDF, Grade, or Curve, see Upgrade Application Configurations for additional steps to complete the application upgrade.

Command-line Support for Configuration Upgrading

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

3

Domain Administration

This chapter describes domain administration tasks such as using domainDaemon and setting environment variables.

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. Once 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. The system administrators can start, stop, and monitor the DomainDaemon processes by using scripts that are provided in this directory.

Usage

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

Table 3–1   DomainDaemon Utility Arguments

Argument	Description
-version	Prints the RPAS version, revision, and build information of the utility.
-start	Starts a DomainDaemon on the specified port.
-port *portNum*	The *portNum* must be between 1025 and 65535 (inclusive). If *portNum* 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 *milliseconds*	Specifies the number of milliseconds to wait for the server to start. A value of -1 means no timeout.
-server *serverProgramName*	Specifies the name of the RPAS database server program. Defaults to RpasDbServer.

Table 3–1   (Cont.) DomainDaemon Utility Arguments

Argument	Description
-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 *pathToDomain*	Adds the specified domain to the list of domains managed by a DomainDaemon.
-activate *pathToDomain*	Reactivates a previously deactivated domain. Specify the port number and the complete path to the domain.
-deactivate *pathToDomain*	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.
-remove *pathToDomain*	Removes the specified domain from the list of domains managed by a DomainDaemon.
-showActiveServers	Shows all active server processes. Specifying a port number is required. For each active server, the DomainDaemon shows the process ID, domain, and user ID.
-stopActiveServers	Stops all active servers. Specify a port number and a process ID.
-stopServer *processId*	Stops the server using the specified processed.
-stopUser *userId*	Stops the server using the specified userId.

Starting the DomainDaemon

In order to start the DomainDaemon, execute the script called DomainDaemon in the installation directory. 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.

> **Note:**   In the following examples, [RPASDIR] stands for the full path to the directory where the RPAS Server is installed.

Example 3–1   Command to start a Domain Daemon on port 55278

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

```
([RPASDIR]/bin)$ DomainDaemon -port 55278 -start &
DomainDaemon on port 55277 is alive.
```

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.

```
([RPASDIR]/bin)$ DomainDaemon -port 55277 -ping
```

Stopping the DomainDaemon

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

Example:

```
([RPASDIR]/bin)$ DomainDaemon -port 55277 -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. Once 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.

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, or numeric data type. 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 &
> ```

Lock Timeout Variable

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 two lines below indicate 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
```

The RPAS_LOCK_TIMEOUT environment variable is also used to handle issues with firewalls and the RpasDbServer. The RpasDbServer now checks RPAS_LOCK_ 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.

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.

Log File Backups

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.

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 `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).

Database Settings

These variables are used for RPAS B-tree storage performance.

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 unavailable, RPAS uses a page split percentage of 90.

RPAS_PAGE_SIZE

This variable sets the page size for the RPAS B-Tree. Valid values are 4K, 8K, 16K, 32K, and 64K. RPAS automatically uses a page size of 2K for arrays where the logical size of the dimension space is less than 1000 cells, and a page size of 4K for arrays where the logical size of the dimension space is between 1000 and 200,000 cells. For all other dimension space sizes, RPAS uses a page size of 8K or the value of this environment variable.

Since RPAS caches page data during all its processes, it is recommended that large page sizes are used when dimension spaces are large and when it is known that most processing happens in an unraveling order of positions. This is always the case when processing is driven by the innermost dimension. In most RPAS domains, dimensions of the Calendar hierarchy are the innermost dimension for most arrays.

RPAS_PROCESSES

This variable is used when doing batch measure registration. The default value is 1. 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 will spawn off that number of parallel processes to register or unregister measures in each local domain.

4

Setting Up the RPAS Fusion Client

Once you install the RPAS Fusion Client, before letting users access the application, you must configure it based on your business need. This chapter describes the configuration options available with the RPAS Fusion Client. It also includes information on how you can configure additional domains that are created after installing the RPAS Fusion Client as well as information about logging into the Fusion Client. This chapter includes the following sections:

- Maintaining Layout and Formatting
- Setting Up the Configuration Properties File
- Setting Up Pivot Table Styles
- Configuring Additional Domains
- Logging in to the Fusion Client

Maintaining Layout and Formatting

The RPAS Fusion Client layout and formatting capabilities are used along with Save Format / All (and sometimes Group) for each template. Mass, consistent formatting is important in this case.

Periodically, the layout and formatting of an installed application will need to be revised, and these capabilities will be utilized once again.

With RPAS, when the domain's configuration is patched in a way that includes structural changes, or sometimes when the RPAS software is upgraded, the format database in the domain can be invalidated. In these cases, all saved formatting (Template All, Group, and User, plus the workbooks themselves) must be cleared out and recreated. Mass, consistent formatting is especially important in these cases.

Setting Up the Configuration Properties File

The Configuration Properties (config.properties) file enables you to set up the session timeout properties.

To set up the config.properties file:

1. Navigate to the following location where the RPAS Fusion Client is installed:

   ```
   <RPAS_Fusion_Client_Installation>/config/
   ```

2. In a text editor, open the *config.properties* file, and set the following parameters:

   - session.remindin - set a time (in minutes) when the users are reminded before the session expires. The default value is 35 minutes.

- session.timeout - set a time limit (in minutes) after which the current user session will expire. The default value is 5 minutes.

- *isSSO* - set the value to true to indicate that the application is running in an Single Sign-On environment.

> **Note:** For information about Single Sign-On, see the *Oracle Retail Predictive Application Server Installation Guide.*

- *profile.request.param* - if SSO is enabled, specify the name of the parameter that will be used to pass the profile name to the RPAS Fusion Client. This enables the installer to change the parameter name when there is a conflict with an existing parameter that is already used.

3. Save the file and for the changes to take effect, restart the WebLogic server.

Setting Up Pivot Table Styles

The RPAS Fusion Client installs with the following default formatting styles for the cells that appear on the pivot table in the worksheets:

- Read-only cells - Cell with a light grey background.

- Protected cells - Cell with a light grey background.

- Edited cells - Text within the cell appear in italics with no background color.

You can change the default formatting styles (set a different background color or font style) using the PivotTableStyles.properties file included in the **<RPAS_Fusion_ Client_Installation>/config/rpas** directory. To change the default pivot table cell formatting styles:

1. Navigate to the following location in the RPAS Fusion Client installation directory:

   ```
   <RPAS_Fusion_Client_Installation>/config/rpas/
   ```

2. In a text editor, edit the **PivotTableStyles.properties** file, and set a relevant value for the following parameters:

   - pivotTable.protected.style - style for the protected cells in the pivot table.

   - pivotTable.unprotected.style - style for the un-protected cells in the pivot table.

   - pivotTable.readonly.style - style for the read-only cells in the pivot table.

   - pivotTable.edited.style - style for the edited cells in the pivot table.

   - pivotTable.edited.text.style - style for the edited text in a cell.

   - pivotTable.unedited.style - style for the unedited text in a cell.

   - pivotTable.undefined.style - style for the undefined cells.

   - pivotTable.*ambiguous.char* - character to display when an ambiguous aggregation occurs.

   To set a background color, specify background-color:`<hexadecimal code for the color>`;. For example, background-color:#DDDDDD;.To set a font style, specify `font-style:<text-style>`;. For example, `font-style:italic;`.

3. Save and close the **PivotTableStyles.properties** file.

4. For the changes to take effect, restart the WebLogic server.

Configuring Additional Domains

During the RPAS Fusion Client installation, the installer automatically sets up the RPAS domain connection configurations in the ProfileList.xml file. In case you choose to set up the domain connection after the installation or set up an additional domain you must manually set up the connection.

To set up a domain connection manually, perform the following steps:

1. Once you have the additional domain set up, navigate to the following location in the RPAS Fusion Client installation directory:

   ```
   <RPAS_Fusion_Client_Installation>/config/rpas/
   ```

2. In a text editor, edit the ProfileList.xml file, and append a connection specification for your domain in the following syntax:

   ```
   <connection-spec name="{domainConnectionName}">
   <server-name>{RPASServerName}</server-name>
           <daemon-port>{Port Number}</daemon-port>
            <domain>
             <domain-name>{domainConnectionName}</domain-name>
             <domain-path>{Location where the domain is installed}</domain-path>
            </domain>
       </connection-spec>
   ```

   > **Note:** In the example above, placeholders have been used as values for the parameters. Ensure that you replace the placeholders with the relevant values.

   Where,

   - connection-spec name - Name for the connection specification. This will match the name for the taskflow configuration and resource files. Each must have a unique name.

   - server-name - The host where the RPAS Server is installed and the DomainDaemon process is running.

   - daemon-port - The port specified when starting the DomainDaemon.

   - domain-name - The name that will be displayed to the user on the login screen

   - domain-path - Location where the domain is installed.

     > **Notes:**
     >
     > - Multiple domains on the same server can use the same configuration.
     >
     > - When generating taskflow files with a text editor, you must run the native2ascii command from the JAVA_HOME/bin folder. This ensures that the characters in the taskflow are displayed properly in the Fusion Client.

3. Save and close the ProfileList.xml file.

4. In the **<RPAS_Fusion_Client_Installation>/config** directory, create a new folder, and specify the domain connection name or connection specification name as the folder name.

5. Navigate to the **fusionClient** directory located within the directory where the domain is installed. This folder contains the following files that define the Activity Taskflow configuration:

■ **taskflowBundle.properties**

■ **taskflow.xml**

These files exist in the domain's fusionClient directory after the domain is created with the rpasInstall script. These files can also be created directly from the report generation tool in the RPAS Configuration Tools. For more information, see the *RPAS Configuration Tool User Guide*.

6. Copy the two files to the domain name folder you created in step 4.

7. Rename the **Taskflow.xml** file to include the domain name. The new file name must now read as **Taskflow_<domainConnectionName>.xml**. Where **<domainConnectionName>** is the domain name.

> **Note:** Due to Java resource handling restrictions, the **domainConnectionName** cannot contain periods.

8. Create a **resources** subfolder within this location and move the **taskflowBundle.properties** file to the resources subfolder.

9. In the resources subfolder, rename the **taskflowBundle.properties** to include the domain connection name or connection specification name. The new file name must now read as **<domainConnectionName>Bundle.properties**. Where **<domainConnectionName>** is the domain name.

> **Note:** Due to Java resource handling restrictions, the **domainConnectionName** cannot contain periods.

10. For the changes to take effect, restart the RPAS and WebLogic servers.

> **Note:** You can choose to set up a global domain and a local domain using the same activity taskflow in the same connection specification.

Example

The following example illustrates a sample configuration where a local domain and a global domain have been specified in the same connection specification:

```
<?xml version="1.0" encoding="UTF-8" ?>
  <connection-spec-list>
   <connection-spec name="MyDomain">
    <server-name>RPAS_host</server-name>
    <daemon-port>12345</daemon-port>
    <domain>
     <domain-name>MyDomainGlobal</domain-name>
        <domain-path>/folder/GloblaDomain</domain-path>
    </domain>
    <domain>
     <domain-name>MyDomainLocal1</domain-name>
     <domain-path>/folder/GlobalDomain/LocalDomain1</domain-path>
    </domain>
   </connection-spec>
  </connection-spec-list>
```

Logging in to the Fusion Client

Before you log on to the Fusion Client, ensure that your system meets the recommended configuration. For more information, see the *RPAS Installation Guide.*

Once you check the configuration, obtain the following information:

- Uniform Resource Locator, URL - you must enter the URL or the Web address of the application in the Web browser to access the application. For example, http://<fullyqualifieddomainname>:<port>/<applicationcontextroot>

- User name, Password, and Profile Name - based on the tasks you want to perform, obtain a user account (that includes user name and password) to log on to the application.

To log in to the Fusion Client:

1. Start Internet Explorer.

2. In the Address bar, enter the Fusion Client URL, and press **Enter**. The Fusion Client Login page appears.

3. On the Login page, enter the username and password, and select the relevant profile in the Profile drop-down list.

4. Click **Login**.

5

User Maintenance

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

The following procedures are discussed in this chapter:

- Add a User
- Add a User Group
- Delete a User
- Delete a User Group
- Edit a User

These procedures can be performed through the RPAS Fusion Client by accessing the User Administration task in the Administration activity in the taskflow (Figure 5–1).

In addition, the usermgr utility is also described in this chapter. This utility allows you to manage users and users groups through a command line interface. For more information, see the Managing Users - usermgr Utility section.

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

Access User Administration Activity

To access the User Administration workbooks:

1. In the taskflow, click the **Expand** icon next to **Administration**. The Administration panel appears.

2. In the Administration panel, click the **Expand** icon next to **User Administration**. The steps within the User Administration task appear.

Figure 5–1   User Administration
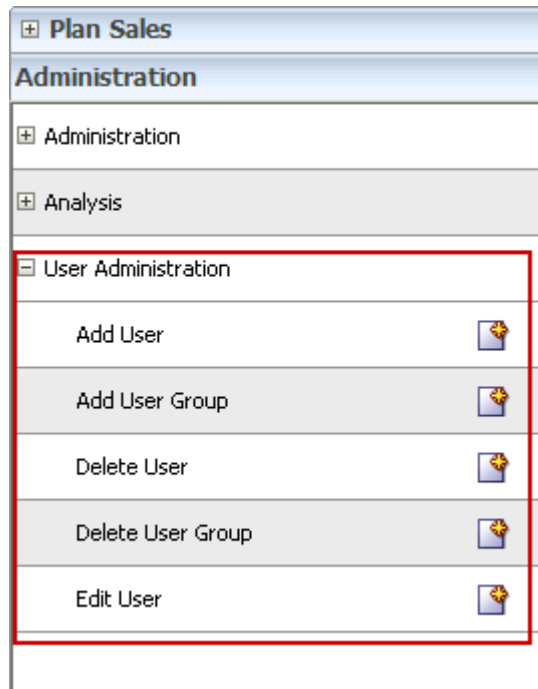
Add a User

To add a user:

1. Under User Administration, click the **Create New Workbook** icon next to **Add User**. A **Workbook Wizard** window appears.

2. In the **Workbook Wizard** window, 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 no spaces (the underscore character is acceptable).

   ▪ **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 File - Open dialog box to identify the owner of a given workbook, and on the Forecast Approval worksheet to specify which user approved a given forecast.

   ▪ **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.

3. Select the **Administrator** check box to assign administrator privileges to the user.

> **Note:** Administrator privilege enables users to perform the Save Format option on the View menu, which creates new system-wide default styles for workbook templates.
>
> Administrative users have special privileges and the read-only status may not apply to them.
>
> If you are not sure whether a user should be granted this ability, you can modify the Admin status later in the Users worksheet (in 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.

4. Select the **Force password change** check box to force the users to change the password when they log on for the first time.

5. Select the **Lock user account** check box to temporarily disable the user's account.

6. 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 Security Administration workbook. For more information, see Access Security Administration.

Add a 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.

To add a user group:

1. Under User Administration, click the **Create New Workbook** icon next to **Add User Group**. A **Workbook Wizard** window appears.

2. In the **Workbook Wizard** window, enter relevant information in the following fields:

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

   - In the **Group Label** field, type a descriptive label for the group. This label is displayed when referring to the group throughout RPAS.

3. Click **Finish** to add the user group to the database.

Delete a User

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

To delete a user:

1. Under User Administration, click the **Create New Workbook** icon next to **Delete User**. A **Workbook Wizard** window appears.

2. In the **Workbook Wizard** window, select the user you want to delete.

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

Delete a 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.

> **Caution:**   Deleting a user group will delete every user in that group.

To delete a user group:

1. Under User Administration, click the **Create New Workbook** icon next to **Delete User Group**. A **Workbook Wizard** window appears.

2. In the **Workbook Wizard** window, select the user group you want to delete.

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

Edit a User

To edit a user's profile:

1. Under User Administration, click the **Create New Workbook** icon next to **Edit User**. A **Workbook Wizard** window appears.

2. In the **Workbook Wizard** window, select the user you want to edit, and click **Next**.

3. Make the necessary changes to the user's profile. You can change anything except the User name. For more information on the fields, see Add a User.

4. 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. Once 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. Once 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.

Once 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.

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 -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 5–1   Arguments Used by the usermgr Utility

Argument	Description
-d *domainPath*	Specifies the path to a domain to add, remove, or get information about a user.
-add *userName*	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.
-label *label*	Use this argument to specify the label of the user or group to add to the domain.
	Use this argument with -add and -addGroup
-group *grp*	Use this argument to specify the user group of the user to add to the domain.
-admin	Use this argument to specify that the user to add to the domain has administrative rights.
-addGroup *groupName*	Use this argument to add a group with a specified name.
	Use -label to specify the label for the group.
-remove *userName*	Use this argument to remove the user with the specified name from the domain.
-removeLabel *label*	Use this argument to remove all users with this label.
-removeGroup groupName	Use this argument to remove a group with this groupName.
-list	Use this argument to list all the users registered to the specified domain.
-print	Use this argument to print the specified user or group information.
-user *username*	Use this argument to specify the user name in the specified domain to print.
	This argument is only applicable to -print option.
-group *groupname*	Use this argument to specify the group in the specified domain name to print.
	This argument is only applicable to -print option.
-importDb	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 -replace is used.
-replace	Use this argument to update existing users when using -importDb.
	The user label, groups that user belongs to, admin status, and account lock status is updated. Password information is not affected by the update.
-exportXml *path*	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.

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

Argument	Description
`-exportDb` *path*	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.
`-convertDbToXml`	Use this argument to convert a user database to a user XML file. Passwords are not included in the conversion.
`-convertXmlToDb`	Use this argument to convert a user XML file to a user database. This prompts you for temporary passwords for admin and user accounts.
`-src` *path*	Use this argument to specify the source file used in `-convertDbToXml` and `-convertXmlToDb`.
`-dest` *path*	Use this argument to specify the destination file used in `-convertDbToXml` and `-convertXmlToDb`.

XML Schema

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

Table 5–2   XML Schema

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

> **Note:**   All XML tags must be in all caps.

Below is a sample users.xml file.

Figure 5–2   Sample users.xml File

```xml
<?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. Once 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.

6

System Administration

This chapter describes the following system administration workbooks and wizards:

- Auto Workbook Maintenance Wizard
- Password Policy Administration Workbook
- Security Administration Workbook
- Hierarchy Maintenance Workbook
- Workbook Batch Category Management

These workbooks and wizards are found in the Administration task in the taskflow.

Figure 6–1   Administration Workbooks and Wizards

> **Note:**   The Alert Manager Wizard is described in the *RPAS User Guide for the Fusion Client*.
>
> The Translation Administration workbook is described in the Internationalization chapter.

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 Workbook Auto Build Maintenance Wizard.

The Workbook Auto Build Maintenance wizard steps the user through the processes of adding and/or deleting workbooks from the auto build queue.

Accessing the Auto Workbook Maintenance Wizard

1. On the Activity Taskflow pane, click the **Expand** icon next to **Administration**. The **Administration** panel appears.

2. In the **Administration** panel, click the **Expand** icon next to **Administration**.

3. Under **Administration**, select the **Create New Workbook** icon next to **Auto Workbook Maintenance**. The **Workbook Wizard** window appears.

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. On the Activity Taskflow pane, click the **Expand** icon next to **Administration**. The **Administration** panel appears.

2. In the **Administration** panel, click the **Expand** icon next to **Administration**.

3. Under **Administration**, select the **Create New Workbook** icon next to **Auto Workbook Maintenance**. The **Workbook Wizard** window appears.

4. In the **Workbook Wizard** window, select **Add a workbook**, and click **Next**.

5. Select a workbook template type, and then click **Next**.

6. Select an owner for the workbook, and then click Next.

7. Enter relevant information in the following fields:

   - **Build Label** - Name of the auto build workbook.

   - **Build Frequency** - Frequency (in days) with which the workbook should be built.

   - **Next Build Date** - Date (in mm/dd/yyyy format) for the next build.

8. In the **Please select the saved access for this workbook** area, select **User**, **Group**, or **World** to specify the saved access for the workbook.

9. In the **Please select the group that owns this workbook** area, select the group that owns the workbook.

10. Click **Next** to initialize the wizard for the workbook template selected in step 5 above. The choices made are saved under the name specified in the **Build Label** field.

Delete a Workbook from the Auto Build Queue

1. On the Activity Taskflow pane, click the **Expand** icon next to **Administration**. The **Administration** panel appears.

2. In the **Administration** panel, click the **Expand** icon next to **Administration**.

3. Under **Administration**, select the **Create New Workbook** icon next to **Auto Workbook Maintenance**. The **Workbook Wizard** window appears.

4. In the **Workbook Wizard** window, 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.

Edit a Workbook in the Auto Build Queue

1. On the Activity Taskflow pane, click the **Expand** icon next to **Administration**. The **Administration** panel appears.

2. In the **Administration** panel, click the **Expand** icon next to **Administration**.

3. Under **Administration**, select the **Create New Workbook** icon next to **Auto Workbook Maintenance**. The **Workbook Wizard** window appears.

4. In the **Workbook Wizard** window, select **Edit workbook settings**, and click **Next**.

5. Ensure that the **OK** radio button is selected, and then click **Finish**. The **Edit Auto Workbook Settings** window appears.

6. In the Edit Auto Workbook Settings window, make the relevant changes in the following columns:

   ■ **Frequency (in days)** - frequency (in days) with which the workbook should be built.

   ■ **Save Group** - from the drop-down list, select the group that owns the workbook.

   ■ **Workbook Name** - name of the workbook.

   ■ **Next Build Date** - date (in mm/dd/yyyy format) for the next build.

   ■ **World Access** - select this check box to make the workbook accessible to all users.

   > **Note:**   The **Owner** and **Template** columns are read-only.

7. Once you make your selection, for the changes to take effect, click the **Commit** icon on the global toolbar. You can also choose to click **Commit** in the **File** menu.

   You can choose to save this workbook by clicking the Save icon on the global toolbar or clicking the **Save** option from the **File** menu.

8. To close the workbook, click **Home** on the top-right corner of the screen.

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 in order 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.

Keep in mind that any hierarchy in RPAS can have user-defined dimensions within it as long as they are set up by your company at the time of installation. The examples in this section refer to the Product hierarchy, but other hierarchies could 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.
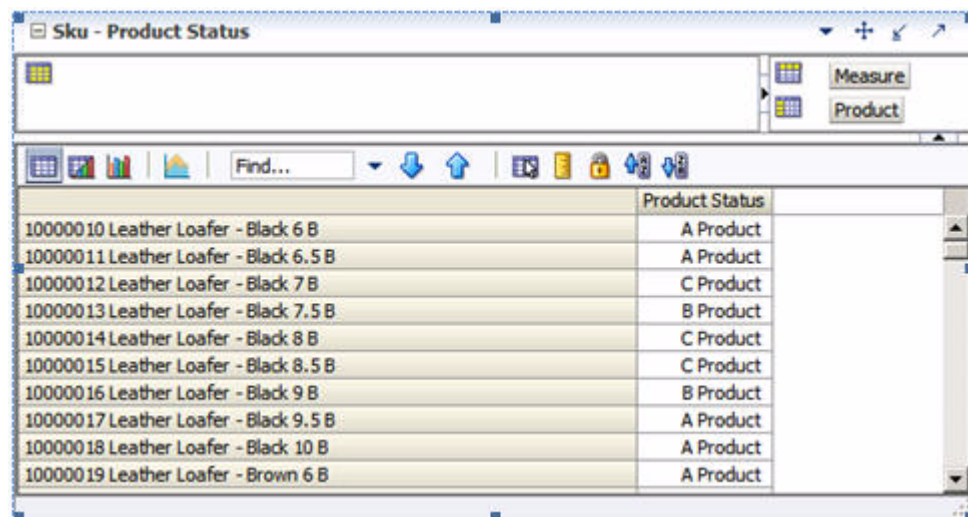
Hierarchy Maintenance Worksheet

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 the following example, products that were previously 'Unassigned' are now designated as A, B, or C Products.

Figure 6–2   Hierarchy Maintenance Worksheet

> **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.

The following figure displays the results when, in a wizard, access a quick menu 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 6–3   Product Status Dimension Results
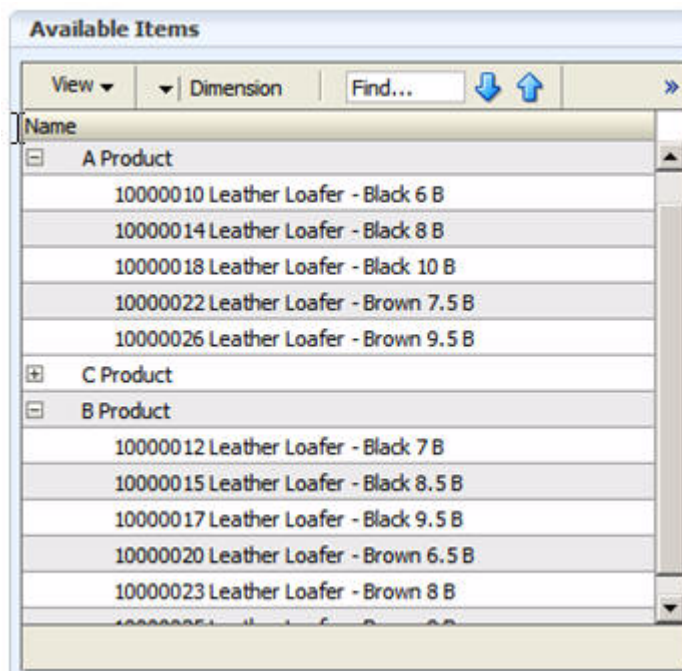
Access Hierarchy Maintenance

> **Note:** If a solution is built in a global domain environment, it is only required and possible to perform all the administrative activities in this section in the master domain.

1. On the Activity Taskflow pane, click the **Expand** icon next to **Administration**. The Administration panel appears.

2. In the **Administration** panel, click the **Expand** icon next to **Administration**.

3. Under **Administration**, select the **Create New Workbook** icon next to **Hierarchy Maintenance**.

Maintain a User-Defined Dimension within a Hierarchy

1. On the Activity Taskflow pane, click the **Expand** icon next to **Administration**. The **Administration** panel appears

2. In the **Administration** panel, click the **Expand** icon next to **Administration**.

3. Under **Administration**, select the **Create New Workbook** icon next to **Hierarchy Maintenance**. A **Workbook Wizard** window appears.

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 appear here.

5. Click **Next**.

6. Select the user-defined dimension to be updated. The number and names of available custom dimensions are set at installation.

7. Click **Next**.

8. From the **Available Items** in the selection wizard, select the items to be mapped to positions within the custom dimension.

9. Use the Add, Add All, Remove, and Remove All buttons or the drag and drop action to move your selection to the **Selected Items** section.

10. Once all items to appear in your workbook have been selected, click **Finish**. The **Hierarchy Maintenance** workbook appears.

11. 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.

12. Once you make your selection, for the changes to take effect, click the **Commit** icon on the global toolbar. You can also choose to click **Commit** in the **File** menu.

    You can choose to save this workbook by clicking the **Save** icon on the global toolbar or clicking the **Save** option from the **File** menu.

13. To close the workbook, click **Home** on the top-right corner of the screen.

14. 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.

15. Select **Commit Now** from the File menu to commit the changes to the master database. If desired, the user may also save the workbook by selecting Save from the File menu.

16. 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.

To access the Password Policy Administration workbook:

1. Under Administration, click the **Create New Workbook** icon next to **Password Policy Administration**. A **Password Policy Measures Settings** workbook appears.

2. Make the necessary changes to the password policies. For more information on the fields, see the table below.

3. Click **Commit** to commit and then click **Save** to save the changes.

Figure 6–4 Password Policy Measures Settings

Table 6–1   Password Policy Measures

Measure	Description	Allowable Values	Default Value
Password Attempts	The number of consecutive failed login attempts (due to an invalid password) before the account is locked.	Choose from the set values in the pick list: 0, 1, 2, …, 10. Choose 0 if you do not want to lock the account.	5
Password Expiration Time	The number of days that a password is valid for an account. Once 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
Password Failure Lockout Time for Administrators	The amount of time that an administrator account is locked after consecutive failed login attempts. Once 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. Once 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.

> **Note:** If a solution is built in a global domain environment, it is only required and possible to perform the administrative activities included in this section in the master domain.

Security Overview

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

User Logon Security

User accounts may 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 may be set or cleared by the domain administrator by using the usermgr utility. For more information, see Managing Users - usermgr Utility.

The account may be marked as **must change password**.

This is useful for brand new accounts. The user will be allowed to logon with the current password and then forced to select a new password.

Must change password may be set or cleared by the domain administrator using the usermgr utility. For more information, see Managing Users - usermgr 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.

Account Lockout can be enabled through the domainprop utility by using the `-lockAccount` flag.

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 expiration can be enabled through the domainprop utility using the `-expirePassword` flag.

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.

Password history may be enabled through the domainprop utility using the `-passwordHistory` flag.

> **Note:** If a solution is built in a global domain environment, it is only required and possible to perform the administrative activities described in this section in the master domain.

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:** Refer to the *RPAS User Guide for the Fusion 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. Once 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 that were 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 6–2   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 through the Insert Measure command.

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

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

Measure Rights Worksheet

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

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

The worksheet contains a drop-down list for each available user and registered measure combination. Three security options are available: Denied, Read-only, Read/Write. Denied prevents the user from viewing data. Read-only allows the user to view the data. Read/Write allows the user to edit data values. However, a commit rule must be configured for a measure for data to be committed to the RPAS datastore.

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

> **Note:** 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 only available to system administrators.

Access Security Administration

1. On the Activity Taskflow pane, click the **Expand** icon next to **Administration**. The **Administration** panel appears.

2. In the **Administration** panel, click the **Expand** icon next to **Administration**.

3. Under **Administration**, select the **Create New Workbook** icon next to **Security Administration**.

Set or Modify Users' Access to Workbook Templates

The Workbook Template Rights worksheet enables you to set the one of the following access privileges to each workbook template for a user:

- Denied
- Read-only
- Full Access

To set or modify a user's access to workbook templates:

1. On the Activity Taskflow pane, click the **Expand** icon next to **Administration**. The **Administration** panel appears.

2. In the **Administration** panel, click the **Expand** icon next to **Administration**.

3. Under **Administration**, select the **Create New Workbook** icon next to **Security Administration**.

4. On the **Workbook Template Rights** worksheet, for each workbook template, select one of the access privileges for the users you want.

5. Once you make your selection, for the changes to take effect, click the **Commit** icon on the global toolbar. You can also choose to click **Commit** in the **File** menu.

You can choose to save this workbook by clicking the **Save** icon on the global toolbar or clicking the **Save** option from the **File** menu.

To close the workbook, click Home on the top-right corner of the screen.

Set Measure Availability for Workbook Templates

1. On the Activity Taskflow pane, click the **Expand** icon next to **Administration**. The **Administration** panel appears.

2. In the **Administration** panel, click the **Expand** icon next to **Administration**.

3. Under **Administration**, select the **Create New Workbook** icon next to **Security Administration**.

4. On the **Workbook Template Measure Rights** worksheet, select each registered measure that must be available for inclusion in the associated workbook template. For measures that must not be included in the associated template, make sure there is no check mark.

5. Once you make your selection, for the changes to take effect, click the **Commit** icon on the global toolbar. You can also choose to click **Commit** in the **File** menu.

You can choose to save this workbook by clicking the **Save** icon on the global toolbar or clicking the **Save** option from the **File** menu.

To close the workbook, click **Home** on the top-right corner of the screen.

Assign or Restrict User Access to Measures

1. On the Activity Taskflow pane, click the **Expand** icon next to **Administration**. The **Administration** panel appears.

2. In the **Administration** panel, click the **Expand** icon next to **Administration**.

3. Under **Administration**, select the **Create New Workbook** icon next to **Security Administration**.

4. On the **Measure Rights** worksheet, for each measure that a user should have access to, select **Read Only** or **Read/Write** from the drop-down list in the cell. For measures to which the user should not have access, make sure **Denied** is selected.

> **Note:** The drop-down list appears when you double-click inside the cell.

5. Once you make your selection, for the changes to take effect, click the **Commit** icon on the global toolbar. You can also choose to click **Commit** in the **File** menu.

You can choose to save this workbook by clicking the **Save** icon on the global toolbar or clicking the **Save** option from the **File** menu.

To close the workbook, click **Home** on the top-right corner of the screen.

Change a User's Ability to Modify Dimensions

1. On the Activity Taskflow pane, click the **Expand** icon next to **Administration**. The **Administration** panel appears.

2. In the **Administration** panel, click the **Expand** icon next to **Administration**.

3. Under **Administration**, select the **Create New Workbook** icon next to **Security Administration**.

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. Once you make your selection, for the changes to take effect, click the **Commit** icon on the global toolbar. You can also choose to click **Commit** in the **File** menu.

You can choose to save this workbook by clicking the **Save** icon on the global toolbar or clicking the **Save** option from the **File** menu.

To close the workbook, click **Home** on the top-right corner of the screen.

Set or Modify Access to Positions (if position level security has been enabled)

1. On the Activity Taskflow pane, click the **Expand** icon next to **Administration**. The **Administration** panel appears.

2. In the **Administration** panel, click the **Expand** icon next to **Administration**.

3. Under **Administration**, select the **Create New Workbook** icon next to **Security Administration**.

4. On the **Position Level Security** worksheet, 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 appears. 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, select the check box in the relevant cell.

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

7. Once you make your selection, for the changes to take effect, click the **Commit** icon on the global toolbar. You can also choose to click **Commit** in the **File** menu.

You can choose to save this workbook by clicking the **Save** icon on the global toolbar or clicking the **Save** option from the **File** menu.

To close the workbook, click **Home** on the top-right corner of the screen.

Limit the Number of Workbooks that a User Can Save

1. On the Activity Taskflow pane, click the **Expand** icon next to **Administration**. The **Administration** panel appears.

2. In the **Administration** panel, click the **Expand** icon next to **Administration**.

3. Under **Administration**, select the **Create New Workbook** icon next to **Security Administration**.

4. On the **Workbook Limit Per User/Template** worksheet, for a workbook template, set an appropriate value as the limit for a user.

5. Once you make your selection, for the changes to take effect, click the **Commit** icon on the global toolbar. You can also choose to click **Commit** in the **File** menu.

You can choose to save this workbook by clicking the **Save** icon on the global toolbar or clicking the **Save** option from the **File** menu.

To close the workbook, click **Home** on the top-right corner of the screen.

Limit the Number of Workbooks Per Template

1. On the Activity Taskflow pane, click the **Expand** icon next to **Administration**. The **Administration** panel appears.

2. In the **Administration** panel, click the **Expand** icon next to **Administration**.

3. Under **Administration**, select the **Create New Workbook** icon next to **Security Administration**.

4. On the **Workbook Limit Per Template** worksheet, for a workbook template, set an appropriate value as the limit.

5. Once you make your selection, for the changes to take effect, click the **Commit** icon on the global toolbar. You can also choose to click **Commit** in the **File** menu.

You can choose to save this workbook by clicking the **Save** icon on the global toolbar or clicking the **Save** option from the **File** menu.

To close the workbook, click **Home** on the top-right corner of the screen.

Limit the Number of Workbooks that a Group Can Save

1. On the Activity Taskflow pane, click the Expand icon next to Administration. The Administration panel appears.

2. In the **Administration** panel, click the **Expand** icon next to **Administration**.

3. Under Administration, select the Create New Workbook icon next to Security Administration.

4. On the **Workbook Limit Per Group/Template** worksheet, for a workbook template, set an appropriate value as the limit for the user group.

5. Once you make your selection, for the changes to take effect, click the **Commit** icon on the global toolbar. You can also choose to click **Commit** in the **File** menu.

You can choose to save this workbook by clicking the **Save** icon on the global toolbar or clicking the **Save** option from the **File** menu.

To close the workbook, click **Home** on the top-right corner of the screen.

Measure Analysis Workbook

The Measure Analysis workbook template allows the user to view data associated with any registered measure in the Oracle Retail Predictive Solutions 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 Oracle Retail Predictive Solutions 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.

Measure Analysis Worksheet

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

Figure 6–5   Example of Measure Analysis Worksheet

Figure 6–5 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.

Access Measure Analysis

1. On the Activity Taskflow pane, click the **Expand** icon next to **Administration**. The **Administration** panel appears.

2. In the Administration panel, click the **Expand** icon next to **Analysis**.

3. Under Analysis, select the **Create New Workbook** icon next to **Measure Analysis**.

Review and Edit Sales or Other Registered Measure Data

1. On the Activity Taskflow pane, click the **Expand** icon next to Administration. The Administration panel appears.

2. In the **Administration** panel, click the **Expand** icon next to **Analysis**.

3. Under **Analysis**, select the **Create New Workbook** icon next to **Measure Analysis**.

   The **Workbook Wizard** window appears.

   > **Note:**   To open an existing Measure Analysis workbook, click the **Show List of Workbooks** icon next to Measure Analysis, and then go to step 8.

4. In the **Workbook Wizard** window, select the measures you want to include in the new workbook. Use Ctrl-Click and/or Shift-Click to select multiple measures.

5. Click **Next**. Based on each hierarchy specified in the base intersection of the selected measures, a hierarchy wizard appears.

6. In the hierarchy wizard, select the relevant positions, and then click **Next**. Repeat this step for each hierarchy wizard.

7. Click **Finish** to open the Measure Analysis workbook.

8. On the Measure Analysis Worksheet(s), view the stored data associated with the measures and hierarchy positions selected in the wizard. Make any changes as required. As an administrator you can commit the changes.

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 is available only in the master domain. It allows users to 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 Workbook** icon in the Workbook Batch Category Management step.

Figure 6–6   Workbook Batch Category Management Wizard
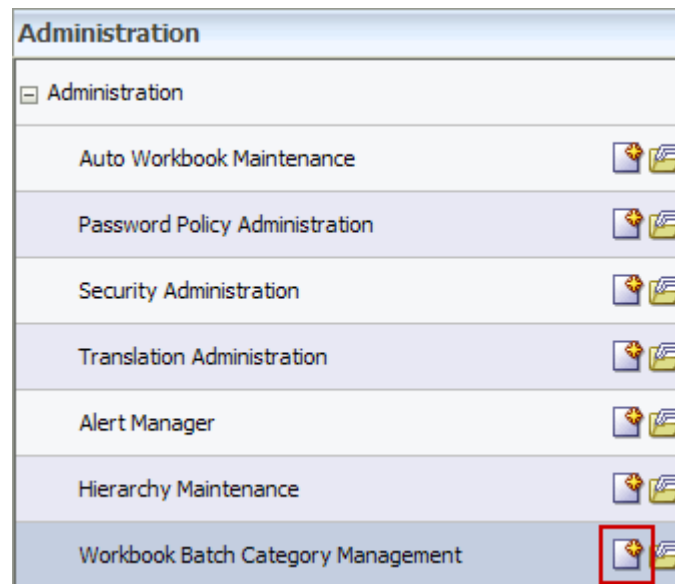
The Workbook Batch Category Management Wizard opens.
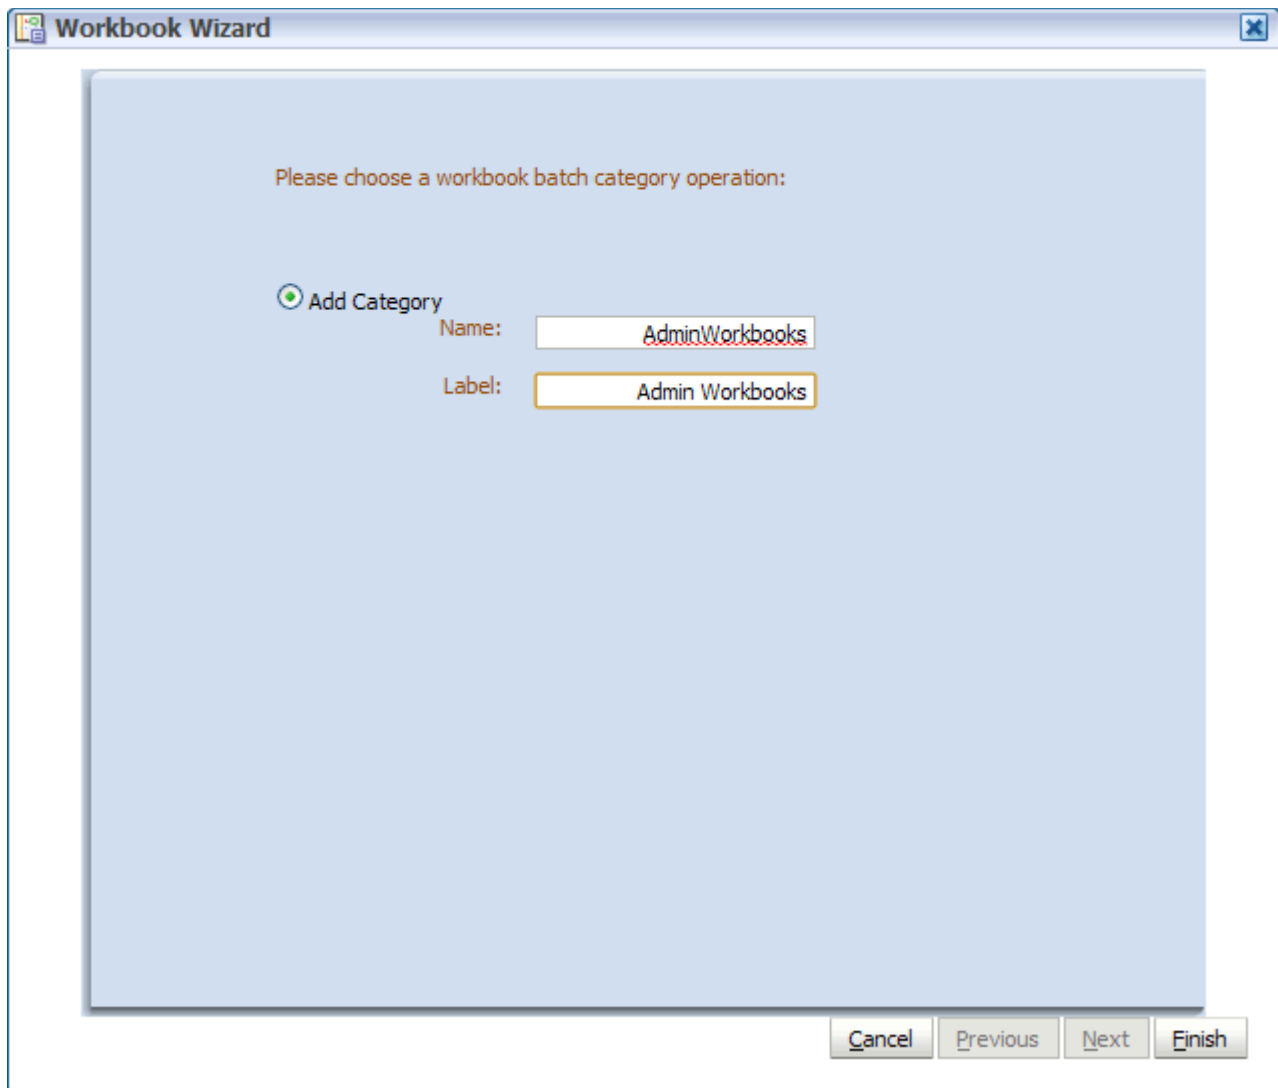
Adding a Category

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

- **Name**: The category name is restricted to standard alphanumeric characters. It cannot contain spaces. This name is used when specifying the category in the wbbatch utility. After a category has been created, this name cannot be changed.

- **Label**: The category label is displayed on the workbook template wizard pages. It can be in any language and can contain spaces. Category labels are case sensitive and must be unique.

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

Figure 6–7   Adding a Category
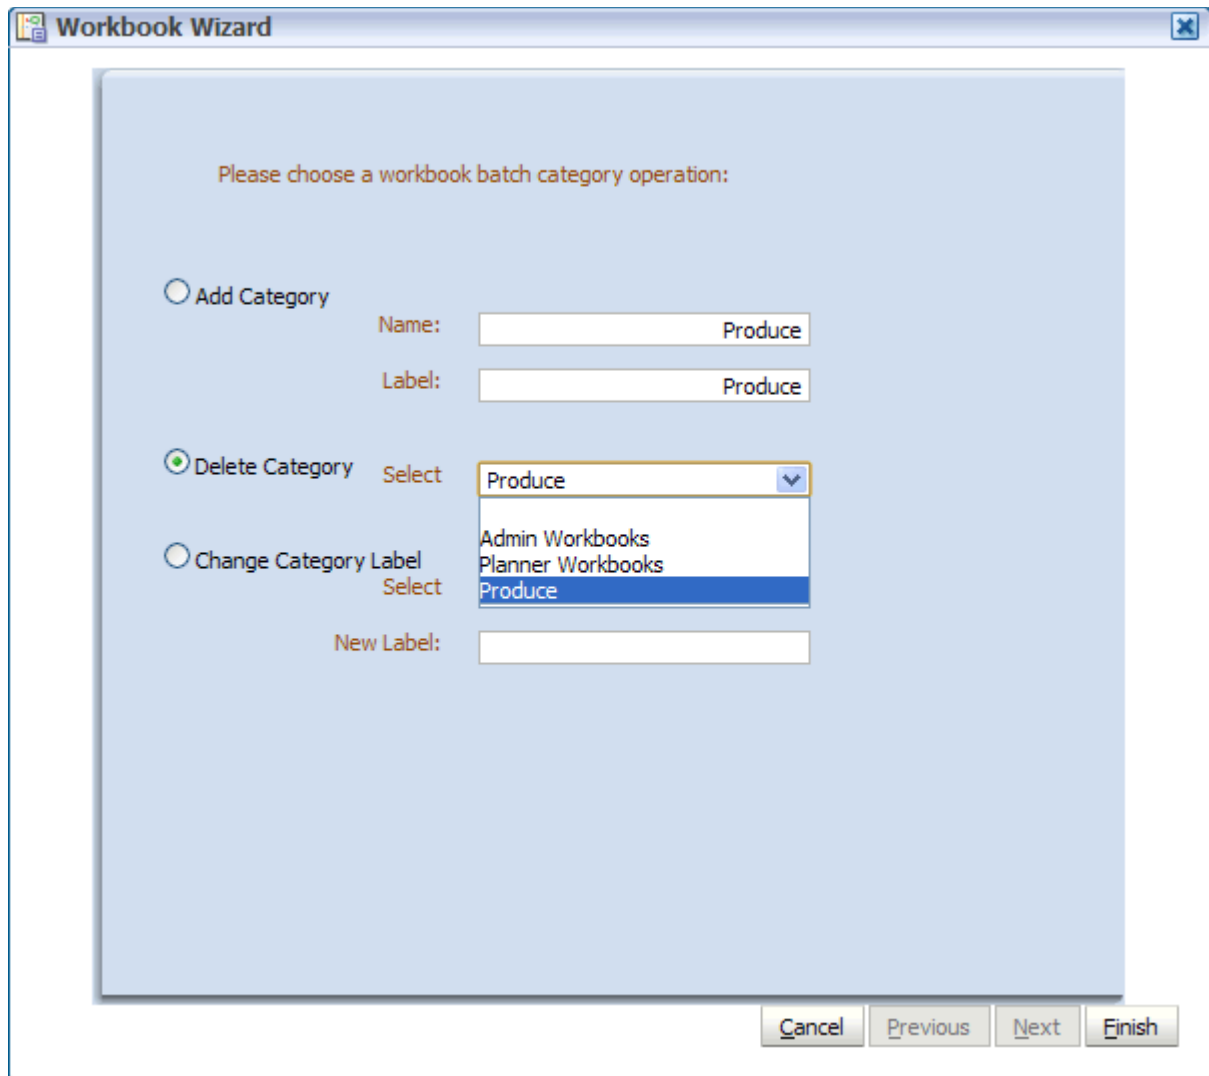
Deleting a Category

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

Figure 6–8   Deleting a Category

Changing a Category Label

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

> **Note:** Category names cannot be changed.

Figure 6–9   Changing a Category Label

7

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.

- The process of updating the data structures in the domain is commonly referred to as reshaping. This process is handled automatically by the system.

- 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. Use of dummy positions defers the time consuming process of reshaping until a scheduled time, thus saving valuable time in the batch window for time-constrained batch processes.

- 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 requires the use of dummy positions and 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.

Placeholder Positions in the Domain

RPAS provides the ability for a domain to contain dummy hierarchy positions, which are placeholder positions that allow new positions to be added to a hierarchy without having to update the data structures in the domain (process referred to as reshaping). These dummy positions are not visible to applications or users, so they cannot be seen in workbooks. They should not be confused with any form of 'dummy' or 'placeholder' positions that are part of a business process and visible to users, such as in the process of new store or item introductions.

RPAS provides the ability for any dimension, other than those in the calendar hierarchy, to contain dummy positions. The number of dummy positions is a percentage of total positions for a given dimension in the domain. For example, if the SKU dimension of the PROD hierarchy contains 1 million positions; a dummy position buffer of 1% will allow for 10,000 dummy SKU positions.

When dummy positions are enabled for a dimension, positions have an internal name (known only to RPAS and not visible to users or applications) and an external name. Dummy positions are positions with an internal name, but no external name. Reshaping is required whenever the collection of internal names changes.

New positions are added to a domain through the hierarchy load process (`loadHier` utility). If a position is new, RPAS will map an existing unused dummy position to the newly added position so that it can be used in the domain without having to reshape the domain.

Similarly, as old positions are deleted, the external name for the internal position is removed from the mapping table (and data for the positions is removed from the arrays), and the position effectively become a dummy position. Therefore, deletes can happen without the need to reshape the domain.

As dummy positions are consumed, the number of available dummy positions will be reduced. Dummy positions are held in the buffer, and the process of updating this is rebuffering. The buffer can be updated automatically, but unpredictably, when required (based on the lower and upper bounds that are defined for a PNI dimension). It is recommended that the rebuffering process is scheduled to ensure that batch process windows are predictable. The manual rebuffering process is executed with the `positionBufferMgr` utility.

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 subdomain(s). 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/
```

Enabling Dummy Positions

The use of dummy positions is enabled in a domain per dimension. Dimensions are enabled for dummy positions using the Configuration Tools both before and after a domain has been built. There are two properties in the Hierarchy Tool, and these properties can be set for each dimension except the calendar hierarchy. These properties are the PNI Buffer Percent for both the lower and upper bounds of dummy positions.

New positions are added to a domain by including new positions in the hierarchy data input files, and then running `loadHier`. Newly added positions will be immediately available for use unless all dummy positions have been consumed, which launches an automatic rebuffering process. Positions that have been assigned a new parent that require movement between domains will be automatically processed as part of the position repartitioning process.

Configuring and Scheduling the RebufferingProcess

Administrators need to determine the process by which they wish to rebuffer the dummy positions. Rebuffering can be completed automatically as part of the hierarchy load process when a domain runs out of dummy positions, or it can be scheduled using an RPAS utility.

If it is desired to have a predictable batch window, it is recommended that administrators schedule the rebuffering process rather than use the automated process. Scheduling rebuffering will minimize the possibility that the automated rebuffering process occurs during a critical batch process. The automated rebuffering would then only be used as a backup and run as an exception.

If customers do wish to reschedule the rebuffering, they will need to determine an approach that fits their business and batch processes. One approach might be to schedule all dimensions (with dummy positions enabled) in all domains to be rebuffered together on a weekly or monthly basis when there is a large amount of system down-time. Another approach could be to rebuffer a few domains on a cyclical basis, such as a few a day or week.

There are high and low settings for the dummy position buffer. Within the Configuration Tools, these settings are the **Buffer % Low** and **Buffer % High**. When executing a scheduled rebuffering process using the `positionBufferMgr` utility, the buffer for a dimension in a given local domain is updated when the number of dummy positions falls outside the high or low target buffer percentages. The number of dummy positions is calculated by taking the average of the high and low percentages multiplied by the total number of positions of the dimension in the local domain.

Deciding the buffer percentages will depend on a number of criteria. The goal is to have sufficient dummy positions to allow for growth in the local domains without having to execute an automated rebuffering process. Determining the targeted number of dummy positions will be a product of the anticipated growth in a given time period (for instance, 100 SKU's per week) and the frequency of the scheduled rebuffering process (for instance, rebuffering once a week or month).

These buffer settings and rebuffering processes are managed by the `positionBufferMgr` utility.

Loading Hierarchies - loadHier

The `loadHier` utility is used to load and refresh a hierarchy. All hierarchy data files are saved in fixed width (or space delimited) files with a .DAT file extension. The width of positions (number of characters) is specified in a configuration file before a domain is built. The width of positions 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.

`loadHier` supports comma separated value (CSV) or fixed width flat files. The utility has also been enhanced to allow a simple compression method that can skip duplicated values line by line. To use a CSV file, instead of a fixed-width file, you must include .csv before the file extension.

Example of loadHier using a CSV File:

```
loadHier -load calendar.csv.dat
```

`loadHier` allows the use of a simple compression technique to reduce the size of the input file and to reduce the file I/O time. CSV input files can be prepared to substitute a column's value with a '?' character if it matches the value from the same column in the previous row. The '?' character can be used for multiple consecutive rows to reuse the value from the row preceding the first use of '?'. Note that `loadHier` assumes that the '?' character will only be used to take advantage of RPAS' proprietary CSV compression. Therefore if a single '?' is meant to be the value of a column, it should be enclosed between opening and closing double quotes; for example, "?".

Due to the compression algorithms used in the CSV format, the `-nosort` argument is automatically used when CSV is detected. This argument is used to prevent the reordering of rows so that RPAS will not lose track of which column value is to be cached for a subsequent row. If the order of records is important, CSV files should be sorted before passing them to `loadHier`.

> **Note:** The `exportHier` and `loadHier` utilities will skip any user-defined dimensions. This is true for both fixed width and CSV formats. At this time, there is no way to export or import user-defined dimensions.

The hierarchy load utility also handles the process of reshaping data in the domain after adding or removing positions. This reshaping process is required to update the underlying data structures to reflect the hierarchical changes; however, note that the use of dummy positions can reduce the frequency of reshaping (see Enabling Dummy Positions for more information).

`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.

By default, loadHier stops with an error if the loadHierBk directory exists in the data directory of any domain. If the domain is in a good state and the backup directory is not valid, it is the responsibility of the user to remove the loadHierBk directory. The switch, -useExistBackup, is used to indicate that the data in the backup directory is valid; loadHier will not overwrite the data in the loadHierBk directory with the current hierarchy data in 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. Because, there will no data moving for those measures not included in the list, that data will be lost when loadHier completes.

Position Label Translation

To enable translation of position labels for the desired dimension(s) 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 11–1, " Supported Languages with Language Identifiers" in the Internationalization chapter.

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. Once 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.

Usage

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

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

Table 7–1   Arguments Used by the loadHier Utility

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.
-checkParents	Verifies the one-to-one mapping of a child position being loaded to its parent position.
-purgeAge purgeage	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 purgeage is reached.
-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.
-noSort	Use this argument for instances when the RPAS process of sorting the hierarchy file prior to loading the values is not needed due to preprocessing (external to RPAS). This flag will prevent the input <hier.dat> file from being sorted while loading the hierarchy.
	This argument is often needed when loading calendar hierarchy data. This option can be eliminated if an alphabetic sort arranges calendar data in chronological order.
	Due to the compression algorithms used in the CSV file format, the -noSort option will automatically be used when CSV is detected. It is recommended that you presort the data before exporting it into a CSV format.
-logLoadedPositions	This argument will enable the logging of successfully loaded input file lines into a loaded[HIERNAME].dat file under the processed directory.
-maxProcesses count	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 count.
-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 7–1   (Cont.) Arguments Used by the loadHier Utility

Argument	Description
`-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:
	`loadHier -defaultDomain ldom1, ldom2, ldom3`
`-returnReshapeFailure`	Returns a non-zero value if one or more arrays fail to reshape. The default behavior is to log a warning with the array name and continue. The `loadHier` process still runs to completion even if a non-zero value is returned.
`-useExistBackup`	Use this argument if the data in the `loadHierBk` directory is valid. `loadHier` will not overwrite the data in the backup directory with the current hierarchy data in the domain.
`-excludeMeasList` *listName* or `-includeMeasList` *listName*	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 *listName*. ■ Use `-includeMeasList` to optimize performance by including only the list of measures in *listName*.

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
```

FilterHier Utility

Sometimes, a retailer has a master file of hierarchy data that needs to be loaded into multiple domains. Some of these domains may be missing one or more levels from the master hierarchy, mostly because the planning levels in these domains are higher than the lowest level in the master and the domains 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 file will be a .csv.dat file that can be subsequently used by the loadHier utility.

Usage

filterHier -d domainPath -input inputPath [COMMAND] {OPTIONS}
The following table provides descriptions of the arguments used by the `filterHier` utility.

Table 7–2 Arguments Used by the `filterHier` *Utility*

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.

Table 7–2   (Cont.) Arguments Used by the `filterHier` ***Utility***

Argument	Description
`-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.

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.

Reconfiguring the 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.

Usage

```
reconfigGlobalDomainPartitions -d pathToMasterDomain -add posName1,posName2, ... -
sub pathToSubDomain
reconfigGlobalDomainPartitions -d pathToMasterDomain -remove posName1, posName2,
...
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 7–3   Arguments Used by the `reconfigGlobalDomainPartitions` *Utility*

Argument	Description
`-d` *pathToMasterDomain*	Specifies the path to the master domain in a global domain environment.
`-add` *posName1, posName2, …*	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` *posName1, posName2, ...*	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` *posName1, posName2, ...*	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` *pathToSubDomain*	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` *pathToInputDir*	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` *count*	If specified, some parts of `reconfig` utility will run in parallel, utilizing up to the given number of processes.
`-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 7–4   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 being 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 (e.g., 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.

- A log file (loadHier.log) will be created in the root directory if loadHier fails.

Updating or Reporting the Dummy Position Buffer - positionBufferMgr

The position buffer manager is used to manually rebuffer dummy positions in the domain or to report on the status of dummy positions. This utility can only be used if dummy positions have been enabled in the domain environment

Usage

```
positionBufferMgr -d pathtodomain [-rebuffer|-report] {-hier hierName}* {-
partitionPositions "pos1,pos2..."}
```

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

Table 7–5   Arguments Used by the `positionBufferMgr` *Utility*

Argument	Description
–d *pathToDomain*	Specifies the path to the domain.
–hier *hiername*	The hierarchy that will be rebuffered. If no hierarchy is specified, the utility will rebuffer or report on all hierarchies configured to have dummy positions other than calendar.
-partitionPositions *"pos1, pos2, …"*	By specifying positions at the partition level, the utility will only rebuffer local domains to which the specified positions belong.
	Positions specified that are not at the partition level will be ignored.
-rebuffer	Adjust all dimensions of the specified hierarchy to have the number of dummy positions based on the configuration of the dimension for the environment.
	If a hierarchy is not specified with this argument, the utility will run on all hierarchies other than calendar configured to have dummy positions.
-report	Report the number of buffer positions for all dimensions of the provided hierarchy. The minimum buffer size is also displayed if applicable.
	If a hierarchy is not specified with this argument the utility will run on all hierarchies configured to have dummy positions other than calendar.
-maxProcesses *max*	If specified, some parts of positionBufferMgr will run in parallel, utilizing up to the given number of processes.

Notes

The minimum buffer size appears as a new column when `positionBufferMgr` is called with the `-report` argument. Also, the minimum buffer size is factored into the output indicating whether rebuffering is needed by `positionBufferMgr`.

The following example of output from `positionBufferMgr -report` where the STYL dimension needs to be rebuffered since it has 3 available positions but a minimum size of 4.

Table 7–6   positionBufferMgr -report Example

Percent Low	Percent High	Min Size	Available	Allocated	Dimension
0	0	0	0	5	CLR
0	0	0	0	10	CLSS
0	0	0	0	7	DEPT
0	0	0	0	1	DVSN
0	0	0	0	2	PGRP
0	0	0	0	12	SCLS
0	0	0	0	1	SDCD
0	0	0	0	4	SIZ1
0	0	0	0	5	SIZ2
25	35	0	6	24	SKU
0	0	0	0	12	SPLR
10	20	0	2	24	STCO
0	0	4	3	18	STYL should add 1 positions
0	0	0	0	5	SZ12

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 that is located in a designated input directory (specified when running utility) and that follows the naming convention **hierName.rn.dat** (for instance, prod.rn.dat).

After the hierarchy data file(s) has been updated and placed in the designated location, an administrator must run the `renamePositions` utility.

Usage

```
renamePositions -d domainPath -i inputDirectory -hier hierName {-log logFileName}
{-n}
```

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

Table 7–7   Arguments Used by the renamePositions Utility

Argument	Description
-d *domainPath*	Specifies the full path to the domain.
-i *inputDirectory*	Input directory where input file with positions to rename is located.
	Utility looks for hierarchy data files with "rn" between hierarchy name and .dat extension (for instance, prod.rn.dat).
-hier *hierName*	Hierarchy for which positions are being renamed.
-log *logFileName*	Optional parameter to generate log file to file name other than default. The default file name is hierRename.log.
-n	Does not apply changes, but generates a report that identifies changes that would be applied.

> **Note:**   The -n is a dry run, which means that it does everything as a true run (for instance, writing to a log file) except that it does not actually commit the changes to the domain.
>
> -log is an optional argument that can be used to name the log file other than the default, which will be created as hierRename.log in the current directory.

Input File

There will be three columns of data in each input line. These columns correspond to the dimension name, the old position name, and the new position name. The three fields will be tab-delimited. Any line not formatted this way will be ignored, and empty lines are also ignored.

Old position names must be an existing position name.

New position names cannot be an already used external name or an existing internal name. Lines having invalid position names will be ignored and added to the log file.

Old Position Name and New Position Name should not be prefixed with the name of the dimension.

The 'width' attribute in the domain must be greater than or equal to the max length of the new external names in the input file. If the width of the new name is greater than the width attribute of the corresponding dimension, RPAS will print an error in the log file and ignore the record.

Dimensions specified in the input file should belong to the hierarchy that is specified in arguments. Otherwise, the record will be ignored, and RPAS will print an error in the log file.

Setting Properties for Dimensions -dimensionMgr

The dimension manager utility is used for setting a number of parameters for dimensions and positions. These parameters are set when using the functionality of Position Name Indirection (PNI). This feature provides the ability to have position names that are longer than the default 24 characters and for a dimension to have dummy positions.

Usage

```
dimensionMgr -d pathToDomain -dim dimensionName [COMMAND]
```
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 7–8 dimensionMgr Utility Arguments

Argument	Description
`–d` *pathToDomain*	Specifies the path to the domain.
`–dim` *dimensionName*	Specifies the name of the dimension to which the settings will apply.
`-specs`	This argument displays the properties of the specified dimension. The dimension properties indicate whether DPM and translation are enabled for the dimension.
`–width` *widthVal*	This argument sets the width of position names for the specified dimension. The default width for positions of a given dimension is 24 characters. Widths can only be extended; they cannot be decreased.
`-minBufferSize` *minSize*	This argument sets the minimum number of unused positions specified by *minSize*. Using this argument, you can assign an absolute minimum size to a dimension buffer. This feature is not available in the Configuration Tools.
`–bufPctMin` *minVal*	This argument sets the minimum percent of unused positions.
`–bufPctMax` *maxVal*	This argument sets the maximum percent of unused positions.
`-enableDPM`	This argument enables Dynamic Position Maintenance (DPM) for the specified dimension (`-dim dimensionName`).
	In order to enable DPM for a dimension, the specified dimension (and all dimensions that roll up to it) must be PNI enabled, meaning that buffer percent minimum and buffer percent maximum are assigned non-zero values.
	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).

Table 7–8   (Cont.) dimensionMgr Utility Arguments

Argument	Description
`-enableTranslation` *width*	This argument enables the specified dimension to use translated labels.
	When enabling translated labels, the numeric parameter passed in the argument (`width`) defines the field width for the translated values in the data file, which is loaded using the `loadMeasure` utility.
`-maxProcesses` *max*	If specified, some parts of `dimensionMgr` will run in parallel, using up to the given number of processes defined by `max`.

Notes

To force the `dimensionMgr` utility to add positions to a dimension, both the `-bufPctMin` and `-bufPctMax` must be set. As long as minimum remains at zero, setting a higher maximum has no effect.

Multiple command arguments are allowed.

Buffer Percent Minimum Size and Buffer Percent Maximum Size are specified as a percentage of the total size of the dimension. For example, a dimension with 200 real positions and a Buffer Minimum of 5 and Buffer Maximum of 20 could have between 10 and 40 extra buffer positions at any given time.

If buffers are defined, then the Buffer Maximum must be greater than the Buffer Minimum and less than 10000. To turn off buffering for a dimension, set both Minimum Buffer and Maximum Buffer to zero.

If both the minimum buffer size and buffer minimum percentage are defined, then dimension is rebuffered to generate the appropriate buffer positions based on the higher value. For example, on a dimension with 100 currently in-use positions, if the minimum and maximum percentages were 10% and 20%, then about 15 available buffer positions would be created. However, if minimum buffer size of 25 is defined, then the dimension would get rebuffered to have 25 positions since that is the larger value. However, if the dimension then grew to 200 current in-use positions, the percentage target would grow to 30 buffer positions, and the 25 position minimum would no longer be relevant.

The primary advantage to the new minimum buffer size is that it works well for dimensions which have very few current positions (or even none).

The new minimum buffer size can be set in one of following ways:

- A new optional field is supported in the hierarchy.xml file at domain creation time as shown below:

  `<buffer_size_min>20</buffer_size_min>`
- Using the `-minBufferSize` `minSize` parameter, shown in the argument table above.

Deleting a DPM Position: When a DPM position is deleted, the underlying position is not automatically returned to the buffer positions. The prepForBatch batch job has to be run to redeem those buffer positions.

Setting Informal Positions to Formal -updateDpmPositionStatus

RPAS supports the ad hoc addition of non-calendar positions to the hierarchy by end users in the RPAS Client. This functionality is referred to as Dynamic Position Maintenance (DPM). Positions added outside of the `loadHier` process are assigned an informal status. `loadHier` ignores informal positions when they are present in input file. `updateDpmPositionStatus` changes the status of a position from informal to formal and enable loading/purging that position through `loadHier` utility and will disable further DPM activities on the position.

This utility cannot be run in subdomains in a global domain environment.

This utility can be run in parallel by specifying the processes option.

This utility expects an input file named as dpmposupdate.xml in the input directory of the domain. Input file must be the XML file format as the following:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes" ?>
<rpas>
    <dimension name="dim1">
            <positions>pos1,pos2,pos3…</positions>
      </dimension>
      <dimension name="dim2">
        <positions>pos1,pos2,pos3…</positions>
      </dimension>
      <dimension name="dim3">
            <positions>pos1,pos2,pos3…</positions>
      </dimension>
</rpas>
```

Multiple positions, which belong to multiple dimensions, can be specified using this input file in order to change their statuses from informal to formal.

In a global domain, the user will place an input file only to the input directory of the master domain. An input file for each subdomain is not required. The utility will split the input file and place it to each of the subdomains' input directories. Only the positions that reside in that subdomain will exist in that input file. The input file will be moved to processed directory under the input directory of the domain. The input file name will be appended with a timestamp.

Usage

```
updateDpmPositionStatus -d domainPath [{PARAMETERS}] {OPTIONS}
```

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

Table 7–9   Arguments Used by the `updateDpmPositionStatus` *Utility*

Argument	Description
−d *pathToDomain*	Specifies the path to the domain.
-processes *max*	When run on a global domain set defines the maximum number of processes to run in parallel.

Exporting Hierarchy Data -exportHier

exportHier is a command-line utility used to export all the positions in a hierarchy, including their rollup relations, from RPAS. By default, the utility assumes that the file will have a CSV flat file format with fixed width format as an optional argument. The utility will export all hierarchy positions, but the file may be specified to include only formal or informal positions. The resulting file could then be used as a .DAT file with loadHier.

Usage

```
exportHier -d domainPath -hier hier_name -datFile dat_file [-fixedWidth][-
onlyFormal | onlyInformal]
```

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

Table 7–10   Arguments Used by the exportHier Utility

Argument	Description
-d pathToDomain	Specifies the path to the domain.
-hier hier_name	The name of a hierarchy in the domain from which the .DAT file will be generated.
-datFile dat_File	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.

Note: The exportHier and loadHier utilities will skip any user-defined dimensions. This is true for both fixed width and CSV formats. At this time, there is no way to export or import user-defined dimensions.

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>
```

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>*

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 7–11   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.

8

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 8–1   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 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.

- .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 we load the data region by region or for a subset of regions at a time. When loading data, we want to 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.

`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 .ovr, .clr and .inc files. If fixed width and CSV files are present, then fixed width files of a particular load type are loaded before CSV files of that type; all files of one load type are loaded before moving on to the next load type. 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.

> **Note:** Please refer to the Data Interface Tool section of the *RPAS Configuration Tools User Guide* for more information.

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.

Loading Data from Below the Base Intersection of the Measures

loadmeasure 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. Please refer to 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 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, 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.

Usage

```
loadmeasure -d pathToDomain -measure measureName{,measureName,…} {-applyloads}{-
processes max} {-noClean} {-forcePurge}{-splitOnly | -noSplit} {-defrag}{-loglevel
level}{-recordLogLevel level}
```
The following table provides descriptions of the arguments used by the loadmeasure utility.

Table 8–1   Arguments Used by the loadmeasure Utility

Argument	Description
`-d` *pathToDomain*	Specifies the domain in which to load the measure.
`-measure` *measureNames*	Specifies the name of the measure(s) to load.
	Measure names must be lowercase (for example, `measurename1`, `measurename2`, `measurename3`).
	If more than one measure is specified, all the measures must be in the same input file.
`-applyloads`	Use this argument to apply any staged loads for the named measure.
	If the measure is registered to be a stage-only measure, `loadmeasure` will put the load in a staging area but will not update the measure until `loadmeasure` is called again with this argument. Upon the use of this argument, `loadmeasure` will apply all loads that have been queued up in the staging area. It will clear out the staged loads unless the measure's `loadsToKeep` property has been set to a non-zero number. In that case, it will not clear out the latest `loadsToKeep` loads.
	Note that only .ovr, .inc, and .clr loads can be staged. .rpl loads cannot be staged. Additionally, staging is only allowed for base intersection loads. RPAS cannot stage loads where load intersection is below the base intersection of the measure.
	This argument should not be used for measures that are not stage-only.

Table 8–1   (Cont.) Arguments Used by the loadmeasure Utility

Argument	Description
-processes *max*	Specifies the maximum number of child `loadmeasure` processes to run concurrently across the local domains in a global domain environment.
	If this argument is omitted or if less than two processes are specified, the application will do all processing in a single process and no child processes will be created.
	This only specifies the number of child processes and the controlling process is not included (*max* + 1 is the actual number of processes).
-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.
-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.
	File-splitting is a single process procedure, so the `-splitOnly` option is mutually exclusive with the `-processes` argument.
	This option should only be used in global domain environments.
-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.
-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 8–1   (Cont.) Arguments Used by the loadmeasure Utility

Argument	Description
`-loglevel` *level*	Use this argument to set the logger verbosity level.
	Possible values: `all`, `profile`, `audit`, `information`, `warning`, `error`, or `none`.
`-recordLogLevel` *level*	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 `-recordLogLevel` *level* argument.
	The standard log levels, error, warning, information, and profile, can be used as parameters to this argument.
	When logging, loadmeasure compares this logging level to the utility's logging level (set using `-loglevel`). If the utility's logging level is less verbose than the record logging level, then record issues will not be logged. If utility's logging level is at same or higher verbosity as the record logging level, the record issues will be logged with the log indicator as set using this argument.

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 Fusion 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.

Usage

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

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

Table 8–2   Arguments Used by the exportMeasure Utility

Argument	Description
-d *pathToDomain*	Specifies the path to the domain.
-out *outFile*	Specifies the output file name. It is required and must be a valid file name including the path.
-wb *wbname*	If specified, `exportMeasure` exports data from the specified workbook (`wbname`). A valid workbook name must be used.
-intx *intxString*	Specifies the export intersection. Measures will be exported at this intersection. If measure's baseint is higher than export intx, measure's default method will be used for aggregation. If measure's baseintx is lower than export intersection, Replication will be used to spread measure down to export intersection.
-mask *measureName*	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 *start:end*	Specifies a range of positions along the innermost dimension. Only values in the range are considered for export.
-processes *max*	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.

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

Argument	Description
–meas "*measSpec, measSpec …*"	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:
| | ```
-meas MeasNameA
-meas MeasNameA.precision(0.0001)
-meas MeasNameA.format("%13.2f").precision(0.01)
-meas MeasNameA.precision(0.01).format("%13.2f")
``` |
	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:
| | ```
-meas MeasNameB
-meas MeasNameB.format("%Y%m%d")
-meas MeasNameB.format("%d%B%Y%H%M%S")
``` |

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" (without quotes). 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.

Usage

```
exportData -d domainPath -out outputFile -params db array
exportData -d domainPath -out outputFile -meas \"measSpec\"
{-wb wbName} {options}
exportData -d domainPath -out outputFile -array \"arraySpec\" {options}
```

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

Table 8–3   Arguments Used by the exportData Utility

Argument	Description
–d *domainPath*	Specifies the domain that contains the data that to export.
-out *outputFile*	Specifies the file that will contain the exported data.
	The *outputFile* is relative to the domain unless the full path is specified.
-meas \"*measSpec*\"	Specifies the measure(s) to export.
	measSpec must be quoted, and the format is \"measName cellFormat naValue naFormat\"
	The -meas argument can be repeated to export multiple measure arrays to the same output file.
	Measures are exported at the base intersection.
-array \"*arraySpec*\"	Specifies the array to export.
	arraySpec must be quoted, and the format is \"dbName|arrayName cellFormat naValue naFormat\"
	■ dbName can be a path to the database (relative paths are relative to the domain root).
	■ Both cellFormat and naFormat 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 *db array*	Instead of specifying all parameters on the command line, this parameter allows the parameters to be read from an array.
	■ *db* specifies the name of a Gem file where the array of parameters is stored.
	■ *array* specifies the name of an array in the specified database that has the above parameters.

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

Argument	Description
`-dim \"`*dimSpec*`\"`	Specifies the dimension to be exported. ■ `dimSpec` must be quoted, and the format is `\"`*dimName conversion format order*`\"` ■ *conversion* 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. ■ `format` is a `printf-style` format for the position names. See the printf function for more information on the possible values. ■ `order` 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` *always | allna | anyna | arrayna*	Controls whether a line of data is exported based on having NAs in a cell. ■ *always* exports data regardless of whether or not it contains NAs. ■ *allna* does not export a row of data if all columns are NA (default). ■ *anyna* does not export a row of data if any cell contains a NA value. ■ *arrayna* does not export a row of data if the value in the given array name is NA (requires -naArray).
`-naArray` *arrayName*	When *arrayna* is specified using the `-skipNA` parameter, this option specifies the export array that is checked to determine if data is exported.
`-index` *arrayName*	Controls whether arrays are indexed by looking at a specified array. Only export the non-NA cells in the given array and each cell in the other arrays that have the same position names. If another array is at a higher dimension level, translate the given arrays cell index to the other arrays.
`-append`	Specifies that output is appended at end of output file. The default is to overwrite output file.
`-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.

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

Argument	Description
`-range` *start:end*	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.
`-precision` *precisionValue*	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` *max*	Defines the maximum number of processes to run in parallel.
`-useArrayNaValue`	This argument will enable the use the NA value of the array instead of the NA value specified in `measSpec` or `arraySpec`.
`-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.

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 8–4   Arguments Used by the mapdata Utility

Argument	Description
`-d` *SrcPath*	Specifies the path to the source domain.
`-dest` *DestPath*	Specifies the path to the destination domain.
`-db` *dbName*	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` *arrayName*	Apply `mapdata` only on the given array. The database in which the array resides must be specified with the `-db` argument.

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.

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 8–5   Arguments Used by the updateArray Utility

Argument	Description
`-destArray` *dbPath.arrayName*	Required argument to specify the destination array where the data will be copied.
	dbPath is relative to `destDomain`.
`-srcArray` *dbPath.arrayName*	Optional argument. Default is no source array.
	Note: This parameter cannot be used with `-srcScalar scalarCell`.
`-destDomain` *domainPath*	Optional argument.
	Default is current working directory.
`-srcDomain` *domainPath*	Optional argument.
	Default is current working directory.
`-maskDomain` *domainPath*	Optional argument.
	Default is current working directory.

Table 8–5   (Cont.) Arguments Used by the updateArray Utility

Argument	Description
-updateMethod *method*	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 *first:last*	Optional argument. Default is no range. Defines range along innermost dimension of source array.

-destRange *first:last*	Optional argument.
	
	Default is no range.
	
	Defines range along innermost dimension of destination array.
	
	The position names of the innermost dimension are the range value. For example, if the range values is one week, the range should be specified as
	
	`-srcRange WEEK200811011:WEEK200811022`
	`-destRange WEEK200811011:WEEK200811022`

-srcScalar "*TYPE:VALUE*"	Optional argument. Default is NA cell. Format for scalar cell is one of: ■ NUMERIC: numeric value ■ STRING: literal value ■ BOOL: Boolean value ■ NA **Note**: This parameter cannot be used with `-srcArray dbPath.arrayName`.

Scan Domain Data - 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.

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 8–6   Arguments Used by the scanDomain Utility

Argument	Description
-version	Prints the version of the utility.
-d *domainPath*	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.

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

Argument	Description
`-backup`	Optional argument to back up database files before attempting to repair them. **Note:** This argument can only be used with `-repairCorruption`.
`-processes` *maximumNumberOfProcesses*	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` *level*	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.

Informal Position Manager

Informal Position Manager is a domain utility which can be used to maintain informal positions for any dimension in a domain. Currently, it can convert positions from formal to informal.

Usage

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

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

Table 8–7   Arguments Used by the informalPositionMgr Utility

Argument	Description
`-d` *domainPath*	The domain to run the utility.
`-hier` *hierName*	The hierarchy to apply the changes.
`-operation` *informalize*	Convert formal positions to informal positions for the given hierarchy. One or more input files are required. Converting a position to informal causes all its children to become informal recursively down to the root dimension.
`-file` *inputFile* or `-dir` *inputDir*	The input files to use. If the `-file` option is used, only one input file is processed and its file name does not matter. If the `-dir` option is used, all files that match the naming patterns described in the File Format section are processed.
`-retain`	Do not move input files to the processed subdirectory after a successful run.

File Format

The file properties are as follows:

- Input file naming convention
  - hierName.informalize{.extraField}. The extraField is optional and can be appended to allow multiple input files to be processed at one run.
  - For example: prod.informalize.20100220
- Content format
  - Each line is for one position specified by dimension name and external position name. Comma delimited. No header line.
  - dimName, positionName

Error Handling

- Input file does not exist: log error and stop execution.
- Any dimension specified in the input file is not DPM enabled or does not exist: log error and stop execution.
- One or more positions do not exist: log warnings, skip the line, and continue.
- Selecting to convert a position to informal that is already informal: no action will be taken on that position. A warning message is logged

9

Operational Utilities

This chapter details 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 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.

Once 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 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.

2. Create the alert (the expression) for which the alert should be evaluated using the Configuration Tools. This flags the registered measure as an alert so that it is recognized when the alert finder is run.

3. 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.

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 9–1 provides descriptions of the arguments used by the alertmgr utility.

Table 9–1   Arguments Used by the alertmgr Utility

Argument	Description
-d *pathToDomain*	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 *a1 a2…*	Evaluate specific alerts in the domain. *a1 a2 …* must be valid names of alerts that are defined in the domain.
-categories *cat1 cat2 …*	Evaluate all alerts in the domain that are associated with specific categories of alerts. *cat1 cat2 …* 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.

Usage

```
copyDomain -xmlConfigFile filename {OPTIONS}
copyDomain -d pathToSrc {OPTIONS}
copyDomain -version
```

Table 9–2 provides descriptions of the arguments used by the copyDomain utility.

Table 9–2   Arguments Used by the copyDomain Utility

Argument	Description
-d pathToSrcDomain	Specifies the path of the domain to be copied. This argument and -dest should not be used with -xmlConfigFile.
-xmlConfigFile pathToXmlConfigFile.xml	This argument allows copyDomain to copy each sub-domain into user-instructed specific locations.
	This argument should not be used with -d OR
	-dest.
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 dimposlist	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 *positions*	Deprecated. Use `-clone` instead.
-copyWorkbooks *workbookList*	Copies only the specified workbooks to the destination location.
	workbookList 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.
-skipEmptyDir	Do not copy the empty directory located in the source domain.

Table 9–2   (Cont.) Arguments Used by the copyDomain Utility

Argument	Description
-maxProcesses 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 (.gem file) 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.

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.

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 sub-domain 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 sub-domain 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** 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.

Usage

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

Table 9–3 provides descriptions of the arguments used by the moveDomain utility.

Table 9–3   Arguments Used by the moveDomain Utility

Argument	Description
-xmlConfigFile pathToXmlConfigFile.xml	This argument will allow moveDomain to move a sub-domain into user-instructed specific locations based paths specified in an xml file. This argument should not be used with the -d, -srcSubDomain, and -dstSubDomain parameters.
-d pathTomaster	This argument will allow moveDomain to move each sub-domain based on the user-specified paths. Enter the path to the master domain.
-srcSubDomain src	Path of the sub-domain to be moved.
-destSubDomain src	Path where the sub-domain is to be moved.

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
<?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.

Usage

```
domainprop -d pathToDomain -expirePassword {days} {-passwordHistory
{oldPasswordCount}} {-property propertyname=value} {-lockAccount {failedLogins}}
```

Table 9–4 provides descriptions of the arguments used by the domainprop utility.

Table 9–4   Arguments Used by the domainprop Utility

Argument	Description
-d *pathToDomain*	Specifies the domain path.
-property *propertyname=value*	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.
	Note: Do not use this argument with -expirePassword, -passwordHistory, or -lockAccount.
-reportSubDomains	Use this option to show property values for subdomains (should all match).

Available Properties

Table 9–5   Available Properties for the domainprop Utility

Domain Property Name	Type	Description
disable_commit_asap	Boolean	This property is not applicable to the Fusion Client because there is only one commit option.
disable_commit_later	Boolean	This property is not applicable to the Fusion Client because there is only one commit option.
domain_name	String	If this property is set, the domain name is displayed on the About page of the RPAS Client (menu Help >About).
help_path	String	This property is not applicable to the Fusion Client.
insert_measure_disabled	Boolean	If this property is set to TRUE, the "insert measure…" item under the Edit menu is disabled.
		By default, this property is set to FALSE.

Table 9–5   (Cont.) Available Properties for the domainprop Utility

Domain Property Name	Type	Description
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.
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

Usage

```
mace -version
mace -d domainPath -find string
mace -d domainPath -newRule {-ruleName ruleName} {-label ruleLabel}{-processes
numProcesses}
mace -d domainPath -delRule ruleName (-processes numProcesses}
mace -d domainPath -addRule    groupName:ruleName {-processes numProcesses}
mace -d domainPath -removeRule  groupName:ruleName {-processes numProcesses}
mace -d domainPath -newGroup    groupName {-label groupLabel}{-processes
numProcesses}
mace -d domainPath -removeGroup groupName {-processes numProcesses}
mace -d domainPath -addExpression ruleName -expression exprString{-processes
numProcesses}
mace -d domainPath -check -expression exprString
mace -d domainPath -run -group groupName     {-debugRuleEngine}
mace -d domainPath -run -expression expString {-processes numProcesses}{-
debugRuleEngine}
mace -d domainPath -resolve groupName -measures measList {-debugRuleEngine}
mace -d domainPath -transit workbookName -group groupList {-debugRuleEngine}
mace -d domainPath -print -rule  ruleList
mace -d domainPath -print -group groupList
mace -d domainPath -print -allGroups
mace -d domainPath -purgeRules {-processes numProcesses}
mace -d domainPath -removeAllRuleData {-processes numProcesses}
mace -d domainPath -validate calc    -ruleGroup groupName
mace -d domainPath -validate general -ruleGroup groupName
mace -d domainPath -validate refresh -ruleGroup groupName -calcRuleGroup calc
```

Table 9–6 provides descriptions of the arguments used by the mace utility.

Table 9–6   Arguments Used by the mace Utility

Argument	Description
-d *domainPath*	Specifies the domain in which to load the measure.
-find *string*	Use this argument to search all expressions for the specified string, printing all the rules and rule groups that have these expressions.
-newRule {-ruleName *ruleName}*	Use this argument to create a new empty rule. If desired, use the *-ruleName* argument to specify a name for the rule.
-label {ruleLabel | groupLabel	Use this argument to specify the label of the rule with the -newRule argument or label of the group with the -newGroup argument.
-processes *numProcesses*	Use this argument to specify the number of child processes to be run in parallel.
-delRule *ruleName*	Use this argument to remove the specified rule.
-addRule *groupName:ruleName*	Use this argument to add the specified rule to the group specified by *groupName*.
-removeRule *groupName:ruleName*	Use this argument to remove the specified *ruleName* from the group specified by *groupName*.
-newGroup *groupName*	Use this argument to create a new rule group with the specified name.
-removeGroup *groupName*	Use this argument to remove the specified group and non-shared rules in it.
-addExpression *ruleName*	Use this argument to add an expression to the specified rule. -expression should be used with this argument.

Table 9–6   (Cont.) Arguments Used by the mace Utility

Argument	Description
-check	Use this argument to validate the specified expression. -expression should be used with this argument.
-run	Use this argument to evaluate the specified expression or rule group. -expression should be used with this argument.
-resolve *groupName*	Use this argument to order (does not evaluate) expressions within rule group. Requires a comma-separated list of edited measures.
-transit *workbookName*	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.
-print {ruleList | groupList | true}	Use this argument to print all the specified rules and rule groups. The ruleList is a comma-separated list of rule names. The groupList is a comma-separated list of group names. If true is supplied for either ruleList or groupList, all rules or rule groups are printed.
-purgeRules	Use this argument to remove all rules not contained in any rule groups.
-removeAllRuleData	Use this argument to remove all rule groups and all rules.
-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 exprString	Use the argument to specify the expression. This argument is used in conjunction with the -addExpression, -check, and -run arguments.
-group groupName	Use this argument to specify the rule group to evaluate using the -run argument.
-measures measureList	Use this argument to specify the measures to resolve.
-group groupList	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 ruleList	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.

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.

Update Auto-Workbook Build Tasks When Using PNI Functionality

If PNI positions in a domain are changed, any auto-workbook build tasks in the workbook batch queue for that domain need to be removed and recreated. If this is not done, the positions in the domain hierarchies that were changed will not be included in the workbook. The following steps illustrate this situation:

1. An auto-workbook build task is in the workbook batch queue.

2. Changes are made to the domain hierarchies. For example, PNI positions are deleted and then reloaded.

3. An auto-workbook build task is created using the wbbatch utility.

4. The workbook will not include the PNI positions that were deleted and then reloaded.

Usage

```
wbbatch -version
wbbatch -d pathToDomain -build queueIndex
wbbatch -d pathToDomain -refresh workbookName
wbbatch -d pathToDomain -commit workbookName
wbbatch -d pathToDomain -scheduleRefresh workbookName [-category categoryName]
wbbatch -d pathToDomain -unscheduleRefresh workbookName
wbbatch -d pathToDomain -scheduleCommit workbookName
wbbatch -d pathToDomain -unscheduleCommit workbookName
wbbatch -d pathToDomain -startQueue [all|build|refresh|commit] [-processes max]
[-categories catName1,catName2]
wbbatch -d pathToDomain -printQueue [all|build|refresh|commit]
wbbatch -d pathToDomain -listCategories
wbbatch -d pathToDomain -addCategories catName1:Label1,catName2:Label2
wbbatch -d pathToDomain -deleteCategories catName1,catName2
wbbatch -d pathToDomain -changeCategoryLabels catName1:NewLabel1,
catName2:NewLabel2
wbbatch -d pathToDomain -queue build -updateCategories queueIndex1:newCatName1,
queueIndex2:newCatName2
wbbatch -d pathToDomain -queue refresh -updateCategories
workbookName1:newCatName1, workbookName2:newCatName2
```

Table 9–7 describes the arguments used by the wbbatch utility.

> **Note:** Fusion Client users cannot submit workbooks into the commit later queue; therefore, the commit arguments in the wbbatch utility do not apply.

Table 9–7 Arguments Used by the wbbatch Utility

Argument	Description
-d pathToDomain	Specifies the domain containing the workbooks.
-build *queueIndex*	Runs workbook build for provided *queueIndex*.
-refresh *workbookName*	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.
-commit *workbookName*	This argument is not applicable to the Fusion Client because it does not allow entries into the commit later queue.
-processes *count*	Used with either -build or -refresh to build or refresh workbooks in the auto-workbook queue in parallel using the specified number of parallel processes.
-scheduleRefresh	Schedules a workbook to be refreshed later by adding it to the workbook refresh batch queue. If the -category option is specified, the scheduled workbook will be in that category. Otherwise, it will be in the default category.
-unscheduleRefresh *workbookName*	Removes a workbook from the workbook refresh batch queue.
-scheduleCommit *workbookName*	This argument is not applicable to the Fusion Client because it does not allow entries into the commit later queue.
-unscheduleCommit *workbookName*	This argument is not applicable to the Fusion Client because it does not allow entries into the commit later 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.
	Since the commit later queue is not applicable to the Fusion Client, the commit queue will always be empty.
-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.
	Since the commit later queue is not applicable to the Fusion Client, the commit queue will always be empty.
-listCategories	Lists both the name and label for all categories.

Table 9–7 (Cont.) Arguments Used by the wbbatch Utility

Argument	Description
-addCategories *Name1:Label1, Name2:Label2, Name3:Label3*	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 *catName1, catName2*	Deletes a category by specifying the name of that category. Multiple categories can be deleted if separated by a comma.
-changeCategoryLabels *Name1:NewLabel1, Name2:NewLabel2*	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 *queueIndex1:newCatName1, queueIndex2:newCatName2* Or: -updateCategories *wbName1:newCatName1, wbName2:newCatName2*	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:
	wbbatch -d pathToDomain -queue build -updateCategories *queueIndex1:newCatName1, queueIndex2:newCatName2*
	Refresh example:
	wbbatch -d pathToDomain -queue refresh -updateCategories *workbookName1:newCatName1, workbookName2:newCatName2*

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.

Usage

```
wbmgr -version
wbmgr -d pathToDomain -list -all
wbmgr -d pathToDomain -list -user userName
wbmgr -d pathToDomain -print  -wbList wb1,wb2,...
wbmgr -d pathToDomain -remove -all
wbmgr -d pathToDomain -remove -user userName
wbmgr -d pathToDomain -remove -user userName -wbList wb1,wb2,...
```

Table 9–8 provides descriptions of the arguments used by the wbmgr utility.

Table 9–8   wbmgr Utility Arguments

Argument	Description
-d *pathToDomain*	Specifies the domain that contains the workbooks.
-list -all	Lists all workbooks in the domain.
-list -user *userName*	Lists all workbooks belonging to the user.
-print -wbList wb1,wb2,...	Prints detailed information about workbooks in the list.
-remove -all	Removes all workbooks from the domain.
-remove -user *userName*	Removes all workbooks from the domain belonging to the specified user.
-remove -user *userName* -wbList *wb1,wb2*	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.

Usage

```
regmeasure -version
regmeasure -d pathToDomain -add measureName -type typeName(-baseint
baseIntersection|-scalar) {-label labelString} {-db dataDbPath}{-navalue naValue}
{-defagg aggType} {-defspread spreadType}{-allowedaggs "aggType1 aggType2"}{-
refreshable (true|false)} {-insertable (true|false)}{-basestate (read|write)} {-
aggstate (read|write)}{-stageonly (true|false)} {-filename fileName}{-loadint
loadIntersectionString} {-clearint clearIntersectionString}{-loadstokeep
loadsToKeep} {-start fieldStart} {-width fieldWidth}{-loadagg loadAgg} {-range
range} {-purgeage purgeAge} {-viewtype viewType}{-syncwith syncWith} {-description
descriptionString} {-picklist}{-materialized (persistent|display)}{-lowerbound
measurename} {-upperbound measurename}{-attr attrName -attrpos attrPosName} {-
scriptname scriptName}{-specialval
action:specval:behavior,action:specval:behavior,...} {-fnhbi}{-hybridaggspec
hiername:aggop,hiername:aggop,...}{-periodstartvalue (true|false)}

regmeasure -d pathToDomain -modify measureName {-label labelString}{-defagg
aggType} {-defspread spreadType} {-allowedaggs "aggType1 aggType2..."}{-
refreshable (true|false)} {-insertable (true|false)}{-basestate (read|write)} {-
aggstate (read|write)}{-stageonly (true|false)} {-filename fileName}{-clearint
clearIntersectionString}{-loadstokeep loadsToKeep} {-start fieldStart} {-width
fieldWidth}{-loadagg loadAgg} {-range rangeString} {-purgeage purgeAge|-
clearPurgeAge}{-viewtype viewType} {-syncwith syncWith} {-description
descriptionString}{-picklist|-nopicklist} {-materialized (persistent|display)}{-
lowerbound measureName} {-upperbound measureName}{-attr attrName -attrpos
attrPosName} {-scriptname scriptName}{-specialval
action:specval:behavior,action:specval:behavior,...}{-hybridaggspec
hiername:aggOp,hiername:aggOp,...}{-periodstartvalue (true|false)}

regmeasure -d pathToDomain -remove measureName
```

Table 9–9 provides descriptions of the arguments used by the regmeasure utility.

Table 9–9   regmeasure Utility Arguments

Argument	Description
-d *pathToDomain*	Specifies the path to the domain. A valid domain path must be specified.
-add *measureName*	Adds a measure with the specified name. Set the values for the measure by using the required arguments and any of the optional arguments.
-type *typeName*	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 *baseIntersection* | -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 *labelString*	Specifies the measure label. If not specified, it defaults to the measure name specified for the -add option.
-db *dataDbPath*	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 *naValue*	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 *aggType*	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 *RPAS Configuration Tools User Guide*. 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 *spreadType*	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 *RPAS Configuration Tools User Guide*. If not specified, it defaults to the measure type's default spread method: Ratio for int and real, and Replicate for string, date, and boolean.

Table 9–9   (Cont.) regmeasure Utility Arguments

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 inlcude 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 9–9   (Cont.) regmeasure Utility Arguments

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 'mmddyyyy:mmddyyyy', where the first date is the starting date of the range and the second date is the ending date of the range. 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 9–9   (Cont.) regmeasure Utility Arguments

Argument	Description
-viewtype viewtype	Specifies the view type of this measure on the RPAS Client. The valid values are: 0 for none, 1 for view_only, 2 for sync_first_lag, 3 for sync_lead_last, 4 for sync_first, and 5 for sync_last. If not specified, it defaults to none. If the view type starts with "sync", the measure is called a 'Virtual Measure'.
	A measure of sync_first_lag type must have two sync measure names specified by the
	-syncwith option. The first syncwith measure name is a 'Period Start Value' type of measure, like opening stock. Measure data at the beginning period of the calendar is synchronized with this period start value kind of measure. The subsequent measure data is synchronized with the other measure data but lagged one period.
	A measure of sync_lead_last type must have two sync measure names specified by the
	-syncwith option. The first measure is a 'Period End Value' type of measure. Measure data at the last period of the calendar is synchronized with this period end value. Measure data of previous periods is synchronized with the other measure lead one period data.
	A measure of sync_first type must have one measure name specified by the -syncwith option. The data of the beginning period is synchronized with this syncwith measure.
	A measure of sync_last type must have one measure name specified by the -syncwith option. The data of the ending period is synchronized with this syncwith measure.
	Measures of view_only type are
	non-persistent. View only measures can only be used in workbooks. Their measure data is calculated during the Fetch process using a calc expression usually specified in the workbook's calc rule group.
-syncwith syncWith	Specifies the measures that the measure must be synchronized with. This option must be specified if the measure is not a virtual measure.
	For sync_first_lag and sync_lead_last measures, the syncwith option must have two measure names separated by a comma. The first measure is used to synchronize the data at the first or the last calendar period. The second measure is used to synchronize data at other periods.
	For sync_first and sync_last measures, the syncwith option must be specified with a single measure name that will be used to synchronize the first or last calendar period.
-description descriptionString	Specifies the description of the measure.
-picklist | -nopicklist	Specifies whether the measure is displayed as a picklist in the Client. The actual value of the picklist is specified by the -range option of the measure.
	-nopicklist is only available with the
	-modify option. It means the measure should not be displayed as a picklist measure in the RPAS Client.

Table 9–9   (Cont.) regmeasure Utility Arguments

Argument	Description
-materialized (persistent | display)	Specifies whether the measure is persistent or display only on the RPAS Server side.
	Persistent measures must have a valid database and arrays to store the measure data.
	Display only measures do not have permanent data arrays associated with it. The data for a Display only measure must be calculated on the fly. As a result, Display only measures can not be used on the RHS of any expression. Display Only measures can still be used on the LHS of a calc expression used in a workbook, in which case a temporary array will be created in the workbook to hold the temporary data for the Display measure.
-lowerbound measureName	Specifies a measure name that defines the lower bound for each cell of the measure. The difference between the -lowerbound and
	-range options is that the -range option specifies a single scalar as the lower bound for all cells of the measure, but the lower bound value specified by the -lowerbound option can be different from cell to cell.
-upperbound measureName	Specifies a measure name that defines the upper bound for each cell of the measure. The difference between the -upperbound and
	-range options is that the -range option specifies a single scalar as the upper bound for all cells of the measure, but the upper bound value specified by the -upperbound option can be different cell to cell.
-attr attrName	Specifies the measure attribute name. If not specified, it defaults to no attribute is assigned to the measure.
	Note: If this option is specified, the -attrpos option must also be specified.
-attrpos attrPosName	Specifies the measure attribute position name. Combined with the -attr option, the measure attribute provides a way to group measures together based on measure attributes.
	Note: If this option is specified, the -attr option must also be specified.
-scriptname scriptname	Specifies a shell script that must be executed as part of a specific event. Currently, the only script that is handled is to give the option of selecting a hierarchy position name as the content of a string measure. In other words, when a user clicks in a cell, the user is presented with a hierarchy dimension single-tree pop-up. The format for this is 'SingleSelect(HIER="<HIER>", DIM="<DIM>") where <HIER> and <DIM> should be replaced with the actual names of the hierarchy and dimension for which the single-tree pop-up should be created.
-specialval action:specval:behavior, action:specval:behavior,...	Specifies a list of measure special values in the form of "Action:SpecialValue:Behavior,…". The special values are stored in the domain's meta data database.
	For Action, the only action supported is: "DISPLAY".
	The only SpecialValue supported is "NAVAL".
	For Behavior, "NULL" means translate any na cell to a blank cell for display. "CELLVALUE" means no translation, just display the navalue as a regular value.

Table 9–9   (Cont.) regmeasure Utility Arguments

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 ParserException 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.

Usage

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

Table 9–10 provides descriptions of the arguments used by the regTokenMeasure utility.

Table 9–10   regTokenMeasure Utility Arguments

Argument	Description
-d *pathToDomain*	Specifies the path to the domain. A valid domain path must be specified.
-add *tokenMeasure=valueMeasure*	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 tokenMeasure=*valueMeasure*	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.

10

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.

The domain path (`-d`) is required for all commands except `-expectedversion`.

Usage

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

Table 10–1 provides descriptions of the arguments used by the `domaininfo` utility.

Table 10–1   Arguments Used by the domaininfo Utility

Argument	Description
`-d`	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.
`-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).

Table 10–1   (Cont.) Arguments Used by the domaininfo Utility

Argument	Description
-xnames	Lists dimensions which use external names.
-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.
-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.
-showrelativepaths	When listing subdomains, indicates if paths are relative. Only relevant in combination with -listsubdomains or -all.
-masterdomaininfo	Lists the master domain path and partition dims for subdomains.
-subdomain *dim,pos*	Indicates to which local domain the specified position belongs. The position can be at or below the partition level.
-all	Displays all of the above information about the domain.
-version	Displays the version of this utility.

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.

Usage

```
checkDomain -d pathToDomain -type expectedType {-q}
```

Table 10–2 provides descriptions of the arguments used by the checkDomain utility.

Table 10–2   Arguments Used by the checkDomain Utility

Argument	Description
-d *pathToDomain*	Path to the domain that needs to be validated.
-type *expectedType*	Expected type of domain: simple, master, or sub.
-q	Quiet mode. Do not display progress messages.

When `checkDomain` is run on a simple domain the following two items get validated:

- The domain directory exists
- It is a simple domain

If `checkDomain` is run on a global domain, it verifies the following:

- The global domain exists
- The global domain is a master domain
- The global domain checks all of the sub-domains for:
    - The sub-domain directory exists and is a sub domain type
    - 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 - rpasversion Utility

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

Usage

```
rpasversion -l pathToLibrary
```

List Contents of a Database - listDb Utility

Use `listDb` to list the basic information of all arrays contained in the databases provided.

Usage

```
listDb pathToDb*
listDb -row -db pathToDb*
listDb -row -pageUsage -db pathToDb*
listDb -row -standardOptions -db pathToDb*
listDb -standardOptions -db pathToDb*
listDb -version
```

Table 10–3 provides descriptions of the arguments used by the `listDb` utility.

Table 10–3   Arguments Used by the listDb Utility

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

Printing Data from Arrays - printArray Utility

Use `printArray` to print the contents of an array.

Usage

```
printArray -array db.array -specs {-maxpos num}
printArray -array db.array {-cell "dim1:pos1,dim2:pos2,..."
{-format "formatString"}
printArray -array db.array -slice "dim1:pos1,dim2:pos2,..."{-format
"formatString"} {-cellsperrow num} {-noposnames}
printArray -array db.array -allpopulatedcells {-format "formatString"}{-
cellsperrow num} {-noposnames}
```

Table 10–4 provides descriptions of the arguments used by the `printArray` utility.

Table 10–4   Arguments Used by the printArray Utility

Argument	Description
`-array` *db.array*	Specifies the array to print.
	Specify the full path to the database containing the array.
	Required for all commands except `-version`.
	db is a full or relative path to a database. Do not specify the .gem suffix. If no other commands are included, the array defaults to `-allpopulatedcells` with cells per row 1. The `-allpopulatedcells` command is still available, but now functions as a useful default action. The `-noposnames`, `-cellsperrow`, and `-format` parameters may still be specified when relying on the implicit `-allpopulatedcells` behavior.
`-specs`	Prints the specifications of the array and positions along each dimension.
`-popcount`	Outputs only the popcount of the specified array. Useful to shell script writers to get the `popcount` value into a shell script variable. For example, `export POPCOUNT=`printArray -array hmaint.dim_year -popcount``
`-cell` *CELLSPEC*	Prints a specific cell value from the array. Must not contain spaces. Must identify a single of 1-D slice. Specify using the format `dim1:pos1,dim2:pos2,...`
`-cellplain` *CELLSPEC*	Outputs a specific cell value with no space padding.
	Useful for scripts when capturing cell values into shell variables. Must not contain spaces. Must identify a single of 1-D slice. Specify using the format `dim1:pos1,dim2:pos2,...`
`-slice` *CELLSPEC*	Prints a one-dimensional slice from the array. Must not contain spaces. Must identify a single of 1-D slice. Specify using the format `dim1:pos1,dim2:pos2,...`
`-allpopulatedcells`	Print all populated cells including the navalue of the array.

Table 10–4   (Cont.) Arguments Used by the printArray Utility

Argument	Description
-format *"fmtstr"*	If -format is specified, any cells with numeric values are interpreted as dates.
	fmtstr (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
-cellsperrow *num*	For multi-cell output commands (-slice and -allpopulatedcells), indicates how many cells should be printed on each line.
-noposnames	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.

Usage

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

Table 10–5 provides descriptions of the arguments used by the printMeasure utility.

Table 10–5   Arguments Used by the printMeasure Utility

Argument	Description
-d *pathToDomain*	Specifies the domain that contains the measure to print. Requires the -m parameter.
-m *measure*	Specifies the measure to print.
-wb *workbookName*	Specifies the workbook associated with the measure to print.
	If -wb is not used, the domain measure information is printed. Requires the -m parameter.
-list	Returns a list of all registered measures in the domain. This argument does not require -d domainPath.
-listHBIMeasures	In a global domain, printMeasure returns a list of all measures registered at or above the partition dimension.
-specs	Returns the list of measure properties. Requires the -m parameter.
-listDataIntersections	Returns the base intersection of the measure
-printData *aggType.intersection*	Prints the nobs and nods format of the measure array at the specified intersection and agg type.

11

Internationalization

Internationalization is the process of creating software that can be translated more easily. Changes to the code are not specific to any particular market. This section describes configuration settings and features of the software that ensure that the base application can handle multiple languages.

Oracle Retail applications have been internationalized to support multiple languages.

Translation

Translation is the process of interpreting and adapting text from one language into another. Although the code itself is not translated, components of the application that are translated include the following:

- Graphical user interface (GUI)

- Error messages

The following components are not translated:

- Documentation (online help, release notes, installation guide, user guide, operations guide)

- Batch programs and messages

- Log files

- Configuration tools

- Reports

- Demonstration data

- Training materials

The user interface has been translated into the following languages:

> **Note:** In Table 11–1, the language identifier is used for position labels. For more information, see the Position Label Translation section. The Windows Language ID is in the foundation.ini file. For more information, see the Translation Administration section.

Table 11–1   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 Position Label Translation.

Every product, location, and calendar position can be presented in multiple languages, as can messages presented through the client. However, before translated strings can be viewed in the client, the following processes must be followed to setup the environment to support multiple languages.

1. Build the domain(s) with the Multi-Language setting enable in the Configuration Tool properties.

2. Change the Web browser locale settings to reflect the relevant language settings in the application user interface using the following steps:

    a. Launch Internet Explorer.

    b. In the Tools menu, click **Internet Options**. The Internet Options window appears.

    c. In the Internet Options window, on the General tab, click **Languages**. The Language Preference window appears.

    d. In the Language Preference window, click **Add** in the Language Preference area to add the relevant language.

     e. Select the language you added, and use the **Move Up** or **Move Down** buttons to set up the order of preference.

     f. On the Language Preference window, click **OK**.

     g. On the Internet Options window, click **OK**.

3. Log on to the RPAS Fusion Client, and access the Translation Administration workbook to set up the translated text for the measure labels, workbook template names, template group names, and other domain-related elements:

     a. On the Activity Taskflow pane, click the **Expand** icon next to Administration. The Administration panel appears.

     b. In the Administration panel, click the **Expand** icon next to Administration.

4. Under Administration, select the **Create New Workbook** icon next to Translation Administration. For more information, see Translation Administration Workbook.

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.

12

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. Within the RPAS Fusion Client, the user has a single commit option available which will run a Commit ASAP process. The process to commit a workbook can be found in the *RPAS User Guide for the Fusion Client*.

Notes:

- If a user attempts to commit a workbook that already has a process in the queue, the original processes will be removed from the queue. That means that there can only ever be one pending commit in the queue for a given workbook/user/template name combination.

- Data within the workbooks can be committed without saving the workbook. When users click Commit, and the workbook has not been saved by then, the RPAS Fusion Client provides an option to Save, Commit, or perform both actions on the workbook.

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 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.

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 12–1   showWorkbookQueues Utility Arguments

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

Usage

```
configCommitAsap -d pathToDomain [-maxProcs numProcs]
[-deadline time] [-display]
```

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

Table 12–2   Arguments Used by the `configCommitAsap` *Utility*

Argument	Description
`version`	Prints the RPAS version, revision, and build information of the utility.
`-maxProcs numProcs`	Sets the maximum number of concurrent commit processes where `numProcs` is an integer greater than 0. Workbooks can be committed in parallel if they do not require access to the same measure databases. If they do share databases, they will be committed sequentially.
`-deadline time`	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.
`-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 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 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 clients.

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 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 that is currently in the commit process.

- **Success** directory - Contains one file per submitted Commit that has successfully completed its commit process.

- **Failed** directory - Contains one file per submitted Commit that either had a failure during its commit process or could not be committed prior to the deadline.

- **Unknown** directory - If the Commit process detects a corrupted queue file, a message gets logged and the file gets moved into the unknown directory.

13

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.

Table 13–1 shows examples of the formatting.

Table 13–1   CVS File Format

Data	Formatted Data
Item 001	Item 001
Item 001, Soda	"Item 001, Soda"
"Large Screen" TV	"Large Screen"" TV"
Item 002, "Generic Brand" Cereal	"Item 002, ""Generic Brand"" Cereal"

RPAS Utilities Logging Options

RPAS has a number of applications used to control or process data. Currently there are no unified methods for logging output, controlling the level of logging, or directing logging to a particular file. Instead each utility has its own methods, although many are similar. The current behavior for each utility follows.

Log Levels

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 13–2   Utilities that Can Use the -loglevel Option

■ alertmgr	■ regfunction
■ checkDomain	■ regmeasattr
■ checkParents	■ regmeasure
■ 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 APPNAMEYYYYMMDDHHMIbXX, 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:

- defrag
- exportData
- loadmeasure
- reshapeArrays
- reconfigGlobalDomainPartitions
- updatedpmpositionstatus
- copyDomain
- wbbatch

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_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; see the entry for reshapeArrays as well. Any log messages generated by reshapeArrays will go to the log file reshapeArrays.log in the current working directory. 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

Themace utility uses standard logging. The -debugRuleEngine option logs some messages to the file mace.log in the current working directory.

positionBufferMgr

Uses standard logging. The reshapeArrays process is spawned as a child. See its entry for details.

reconfigGlobalDomainPartitions

Uses standard logging. The reshapeArrays and loadHier processes may be spawned as children. See their entries 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 regServerYYYYMMDDHHMIbXX, 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.

reshapeArrays

The reshapeArrays utility is used to make arrays conform to the current hierarchies in the domain. Any positions added to dimensions as a result of running loadHier are added to arrays that have this dimension. Any positions removed from dimensions are removed from the arrays that have this dimension. reshapeArrays updates the arrays to reflect these changes made in the hierarchies.

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 Korn shell, 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 13–3   Standard Arguments for RPAS Utilities

Argument	Description
-version	Use this argument to get the version information of the utility (for instance, RPAS 11.2.0). 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

Users that enter workbooks into the commit ASAP queue, while an exclusive lock is in place, will have the entries processed once 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)

- 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 13–1   Domain Access Control
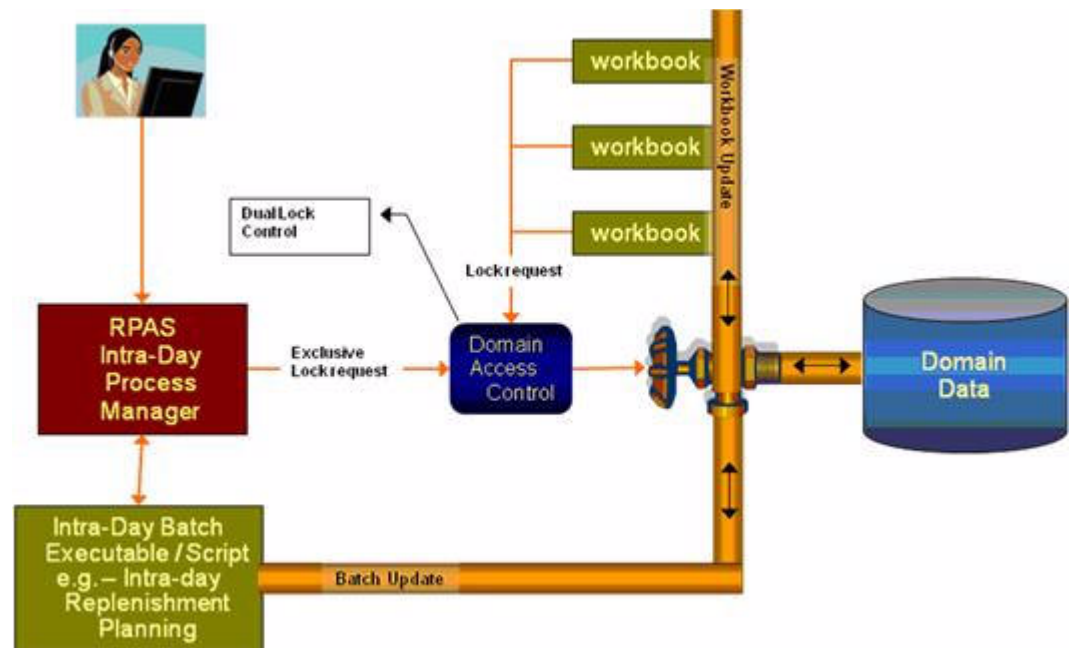
This figure shows the process control that is in place with the locking schema. In this case, an administrator requests exclusive access to a domain in order to run a batch job. This requires an exclusive lock on the domain in order for the job to run. Once the lock is received, no other workbook operations 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 for details on that procedure.

Usage

```
ride -d domain -process pname|-script sname -args args {-message messageString} {-
timeout minutes} {-wait minutes} {-partitions pos1,pos2,…} {-masterInBatch}
```

Table 13–4 provides descriptions of the arguments used by the ride utility.

Table 13–4   Arguments Used by the ride Utility

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

Table 13–4   (Cont.) Arguments Used by the ride Utility

Argument	Description
-message *messageString*	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 *minutes*	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 *minutes*	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 *pos1,pos2,…*	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 13–5   Running the ride Utility Specifying Only the Master Domain

Workbook Operation		Master Domain	Local Domain 1	Local Domain 2	Local Domain 3
Build	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Blocked	Blocked
Refresh	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Blocked	Blocked
Commit Now	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Blocked	Blocked
Custom Menu (not ride concurrent)	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Blocked	Blocked

Table 13–5   (Cont.) Running the ride Utility Specifying Only the Master Domain

Workbook Operation		Master Domain	Local Domain 1	Local Domain 2	Local Domain 3
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 13–6   Running the ride Utility Specifying One Local Domain

Workbook Operation		Master Domain	Local Domain 1	Local Domain 2	Local Domain 3
Build	HBI Measures	Blocked	Blocked	Allowed	Allowed
	No HBI Measures	Blocked	Blocked	Allowed	Allowed
Refresh	HBI Measures	Blocked	Blocked	Allowed	Allowed
	No HBI Measures	Blocked	Blocked	Allowed	Allowed
Commit Now	HBI Measures	Blocked	Blocked	Allowed	Allowed
	No HBI Measures	Blocked	Blocked	Allowed	Allowed
Custom Menu (not ride concurrent)	HBI Measures	Blocked	Blocked	Allowed	Allowed
	No HBI Measures	Blocked	Blocked	Allowed	Allowed
Insert Measure	HBI Measures	Blocked	Blocked	Allowed	Allowed
	No HBI Measures	Blocked	Blocked	Allowed	Allowed
DPM		NA	Blocked	Allowed	Allowed

Scenario 3

This scenario 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 13–7   Running the ride Utility Specifying One Local Domain and masterInBatch

Workbook Operation		Master Domain	Local Domain 1	Local Domain 2	Local Domain 3
Build	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Allowed	Allowed
Refresh	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Allowed	Allowed
Commit Now	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Allowed	Allowed
Custom Menu (not ride concurrent)	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Allowed	Allowed
Insert Measure	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Allowed	Allowed
DPM		NA	Blocked	Blocked	Blocked

Scenario 4

This scenario 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 13–8   Running the ride Utility Specifying Two Local Domains

Workbook Operation		Master Domain	Local Domain 1	Local Domain 2	Local Domain 3
Build	HBI Measures	Blocked	Blocked	Blocked	Allowed
	No HBI Measures	Blocked	Blocked	Blocked	Allowed
Refresh	HBI Measures	Blocked	Blocked	Blocked	Allowed
	No HBI Measures	Blocked	Blocked	Blocked	Allowed
Commit Now	HBI Measures	Blocked	Blocked	Blocked	Allowed
	No HBI Measures	Blocked	Blocked	Blocked	Allowed
Custom Menu (not ride concurrent)	HBI Measures	Blocked	Blocked	Blocked	Allowed
	No HBI Measures	Blocked	Blocked	Blocked	Allowed
Insert Measure	HBI Measures	Blocked	Blocked	Blocked	Allowed
	No HBI Measures	Blocked	Blocked	Blocked	Allowed
DPM		NA	Blocked	Blocked	Allowed

Scenario 5

This scenario 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 13–9   Running the ride Utility Specifying Two Local Domains and the masterInBatch Option

Workbook Operation		Master Domain	Local Domain 1	Local Domain 2	Local Domain 3
Build	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Blocked	Allowed
Refresh	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Blocked	Allowed
Commit Now	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Blocked	Allowed
Custom Menu (not ride concurrent)	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Blocked	Allowed
Insert Measure	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Blocked	Blocked	Allowed
DPM		NA	Blocked	Blocked	Blocked

Scenario 6

This scenario is an example of running ride to lock the master domain.

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

Table 13–10   Running the ride Utility to Lock the Master Domain

Workbook Operation		Master Domain	Local Domain 1	Local Domain 2	Local Domain 3
Build	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Allowed	Allowed	Allowed
Refresh	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Allowed	Allowed	Allowed
Commit Now	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Allowed	Allowed	Allowed
Custom Menu (not ride concurrent)	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Allowed	Allowed	Allowed
Insert Measure	HBI Measures	Blocked	Blocked	Blocked	Blocked
	No HBI Measures	Blocked	Allowed	Allowed	Allowed

Table 13–10   (Cont.) Running the ride Utility to Lock the Master Domain

Workbook Operation	Master Domain	Local Domain 1	Local Domain 2	Local Domain 3
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.

Usage

```
domainStatus -d domain -autoRefresh refreshPeriod
```

Table 13–11 provides descriptions of the arguments used by the domainStatus utility.

Table 13–11   Arguments Used by the domainStatus Utility

Argument	Description
-d *domain*	Refers to a simple or master domain. When master domain is specified, all local domains in the global domain environment are locked as well as the master domain.
-autoRefresh *refreshPeriod*	Refreshes the lock status information every number of seconds specified by the refreshPeriod.

14

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 14–1 Console Manager Window Opened for First Time

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

Figure 14–2   Console Manager Window Showing Console Root Directory
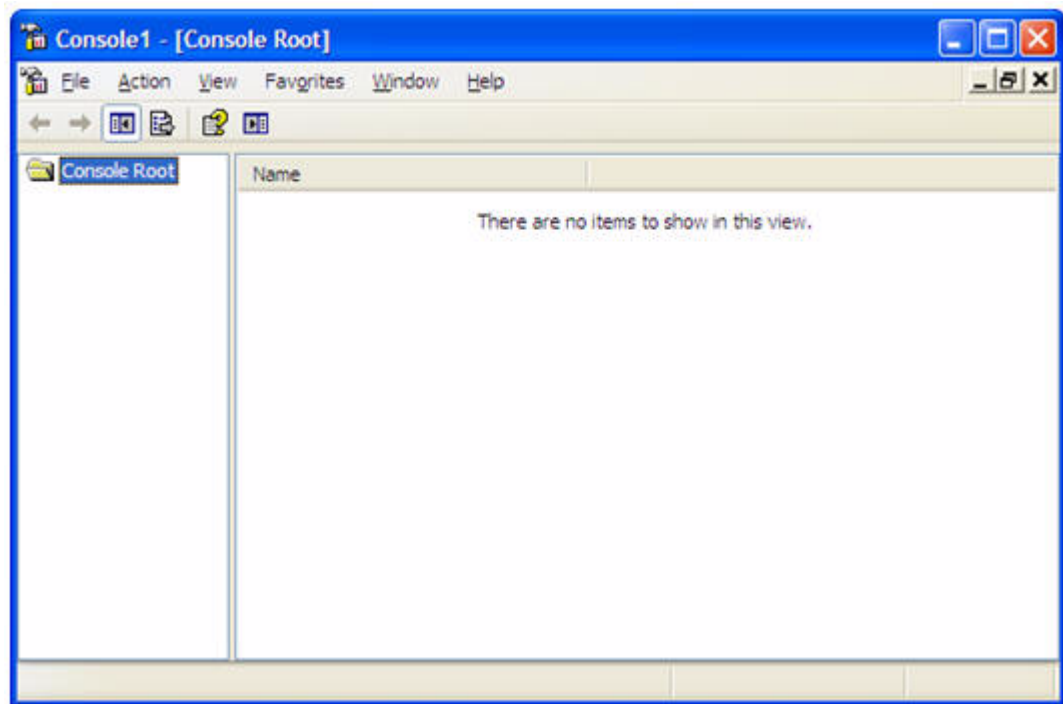
To connect to a remote host, the Management Console is installed on a Windows PC and the Agent service is installed on a remote server.

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 14–3   Add/Remove Snap-in Dialog Box
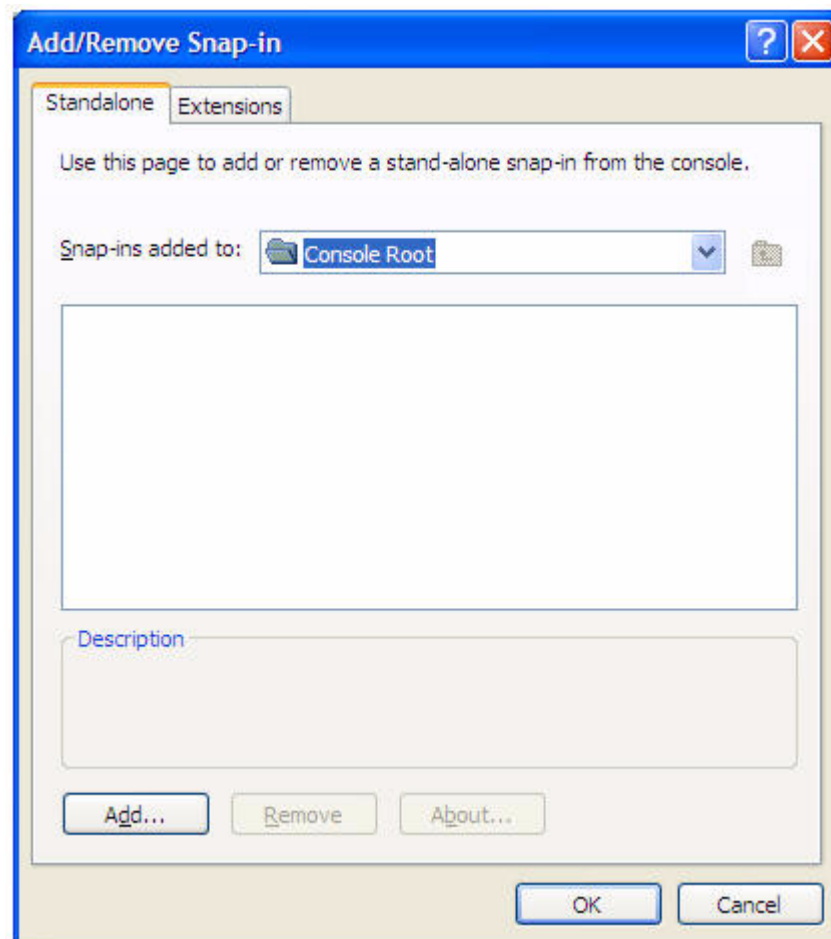
4. To add a new snap-in, click **Add**.

Figure 14–4   Add Standalone Snap-in Dialog Box
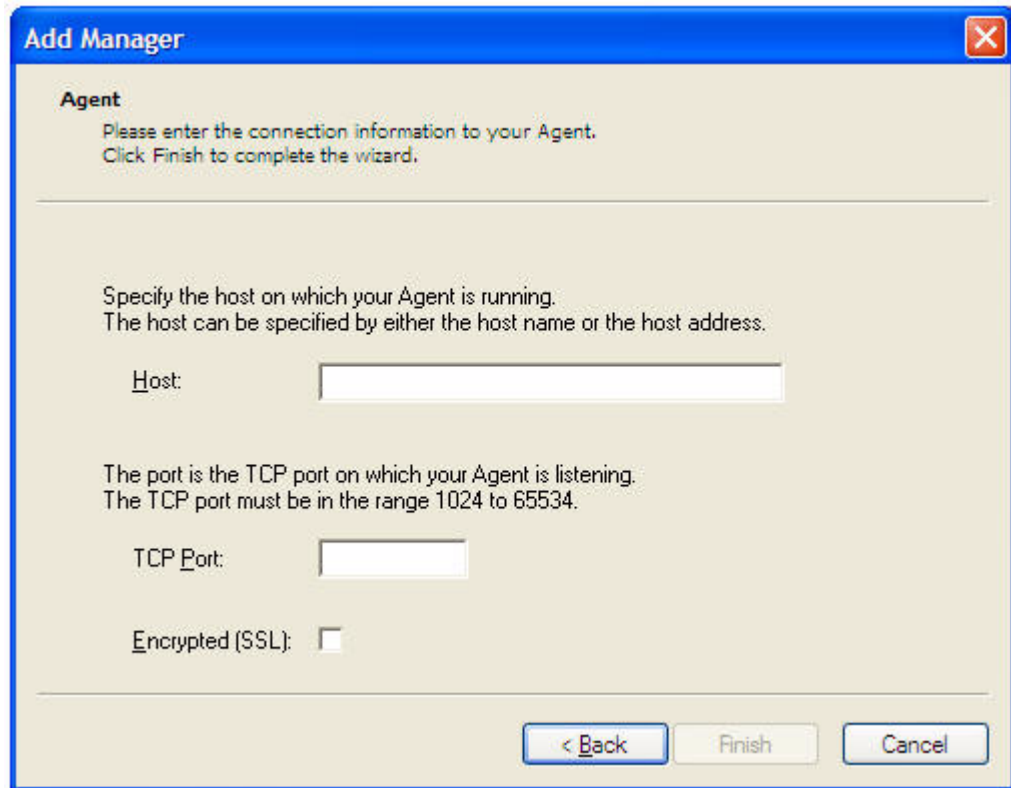
5. Select **Oracle Retail RPAS ODBC Server Manager** and click **Add**.

Figure 14–5   Add Manager Dialog Box Showing Selection of the Mode for a New Manager
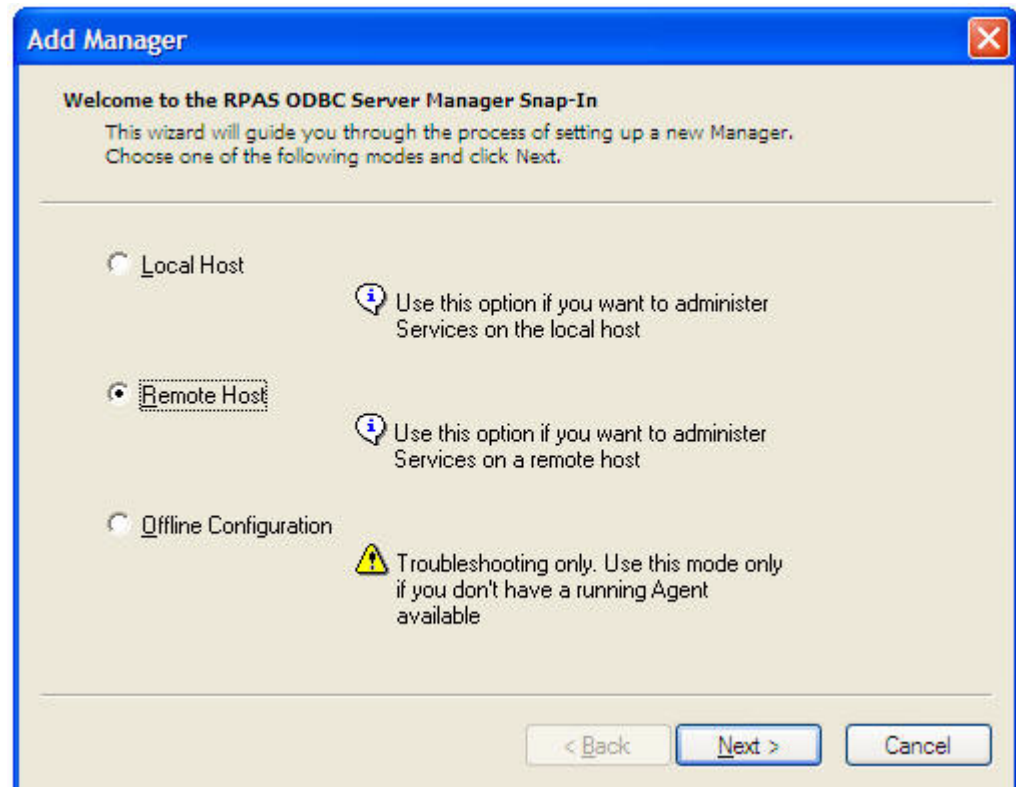
6. Select Local Host or Remote Host depending on the location of the server you want to manage and then click **Next**.

Figure 14–6   Add Manager Dialog Box Showing Connection Information to be Added for the Agent

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 14–7   Console Manager Window Showing Expanded Console Root Directory
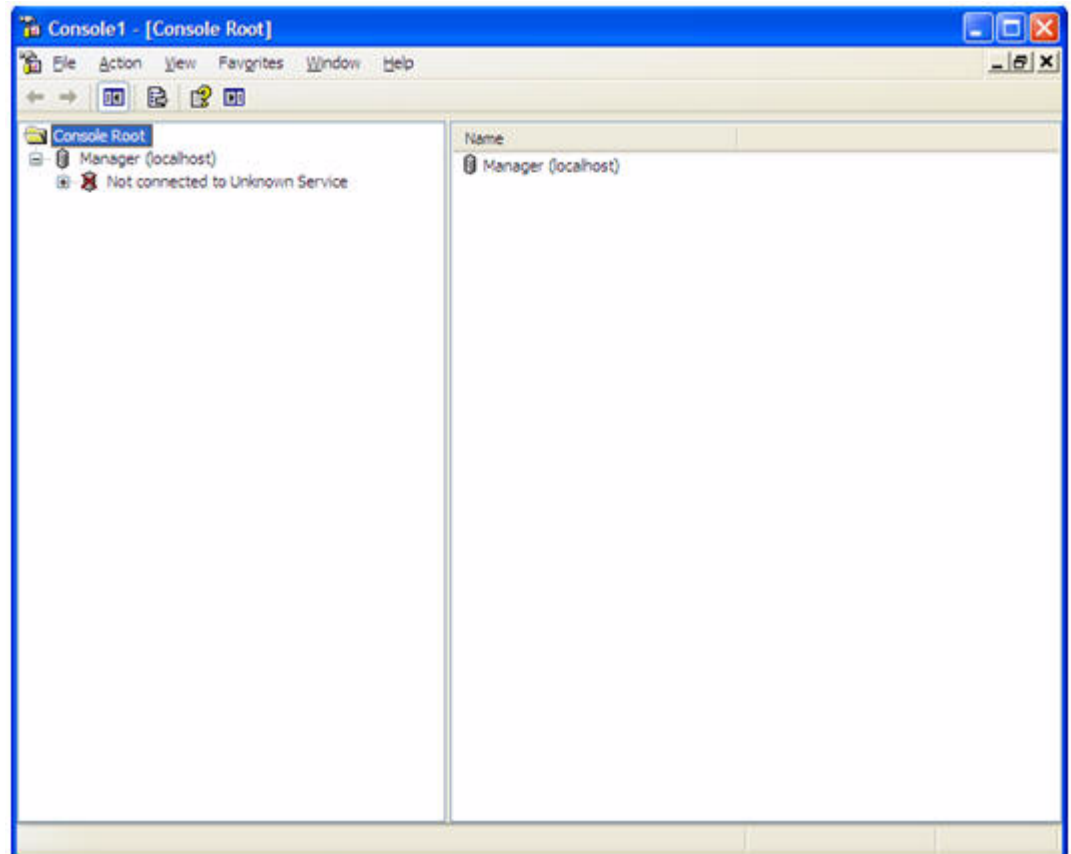
2. To connect to the Agent service, click **+** in front of "not connected to Unknown Service". If the Logon to Service dialog box is displayed, log on using the user name and password of a user who can administer the service on the local or remote server.

3. Expand Services by clicking **+**.

Figure 14–8   Console Manager Window Showing Expanded Connected to RPAS_ODBC_Agent Directory
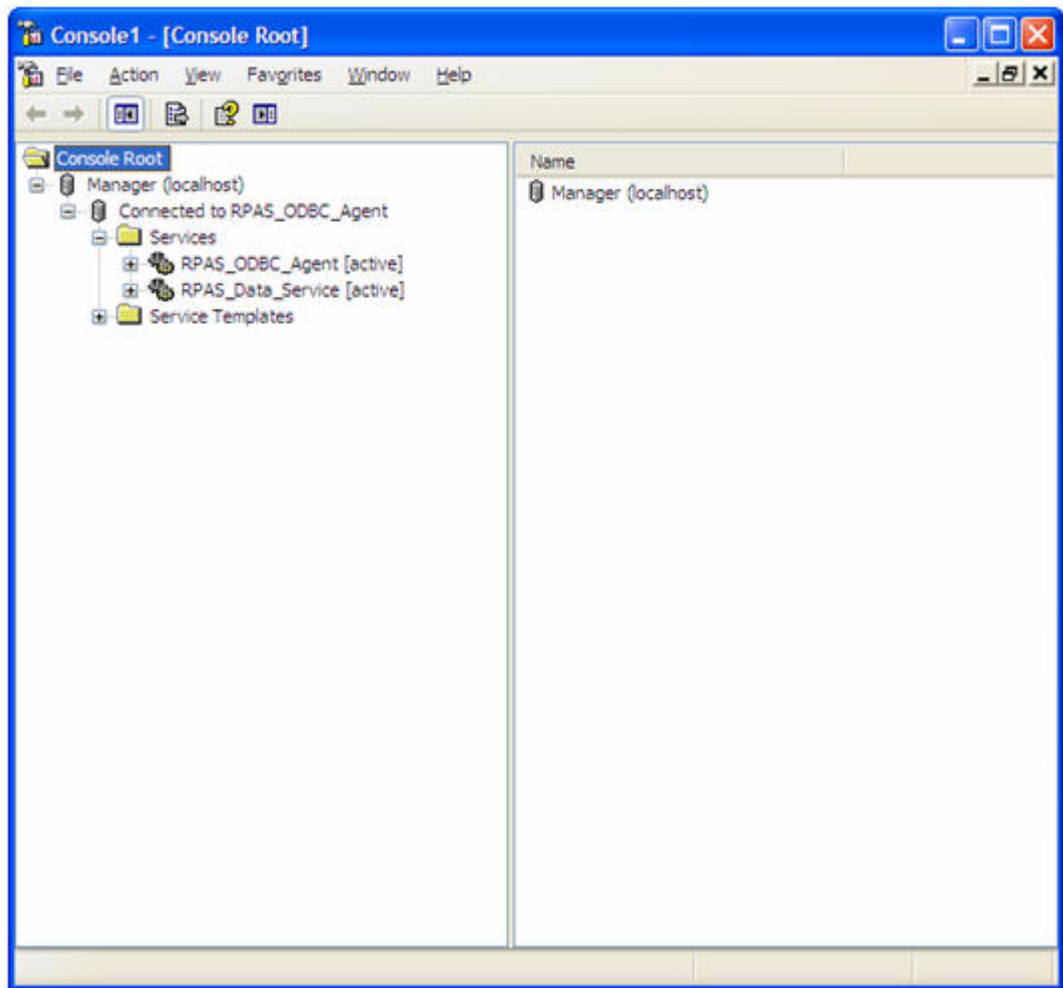
4. Set up the environment variables.

> **Note:**   This step only applies to RPAS_Data_Service running on a Windows platform.

The PATH environment variable must be configured to include the following paths:

■   Path to the RPAS home library

■   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 14–9   Console Manager Window for Configuration
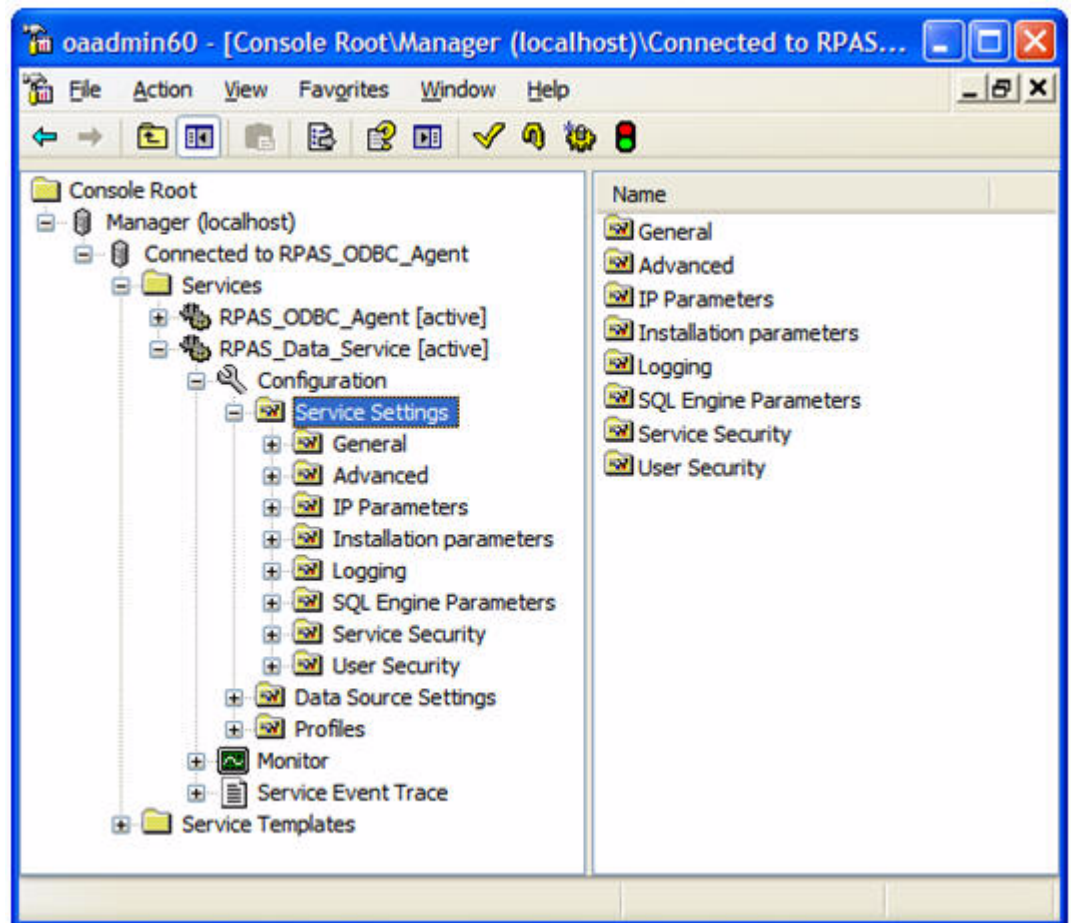
b. Right-click on the blank space of the right panel.

c. Select **New/Attribute** in the menu.

d. Select **ServiceEnvironmentVariable** from the list for Attribute. In the Value field, enter the following:

PATH={pathToRpasHomeLib};{pathToODBCServerBin};{pathToODBCServerIPBin}

In the sample shown in 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 14–10   New Attribute Dialog Box
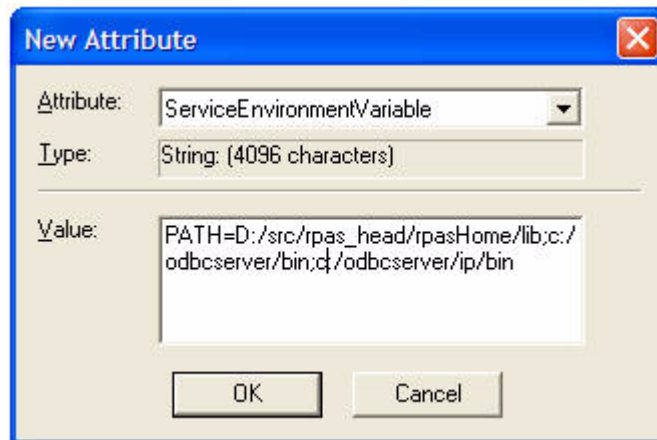
e. Click **OK**.

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.

6. Under gdom, select **IP Parameters**.

Figure 14–11 Console Manager Window with IP Parameters Selected Under gdom
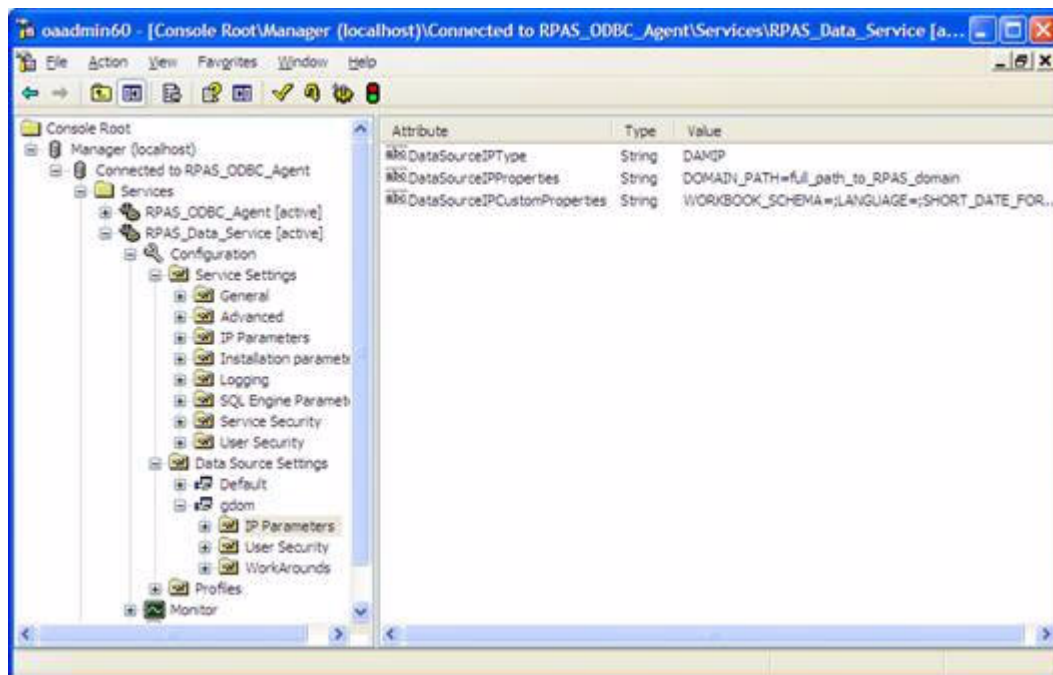
7. On the right panel of the window, double click **DataSourceIPProperties** attribute. The following window appears.

Figure 14–12 DataSourceIPProperties Properties Dialog Box
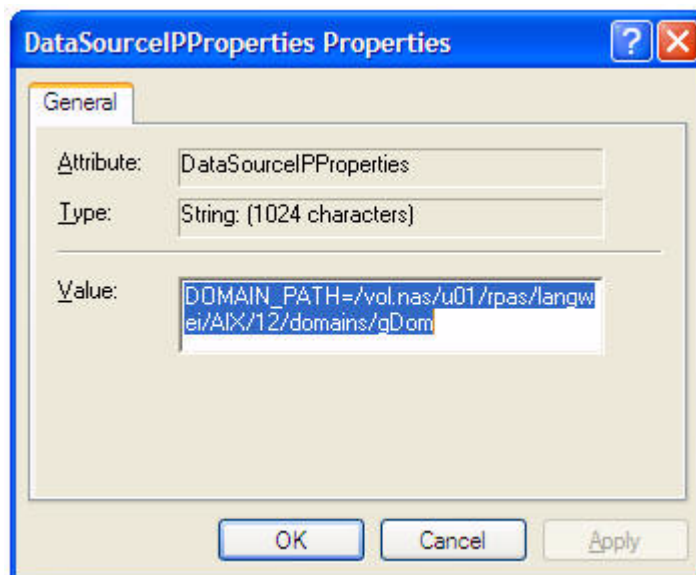
Make sure the Value field has the keyword **DOMAIN_PATH=,** and the string after the "=" sign is the absolute path to the RPAS domain you want connect to on the server side.

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 14–13   DataSourceIPCustomProperties Properties Dialog Box
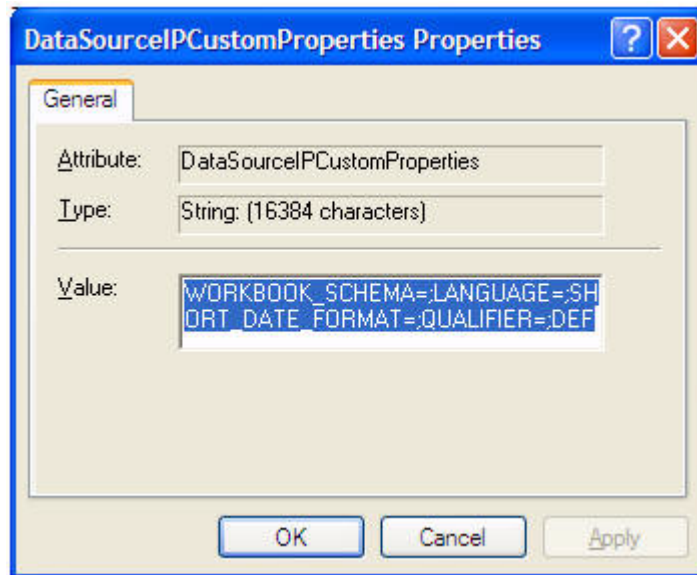
Copy the value from the DataSourceIPCustomProperties of the sample gdom data source to make sure all pre-registered properties are included.

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 *<installdir>*\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 14–14   ODBC Data Source Administrator Dialog Box
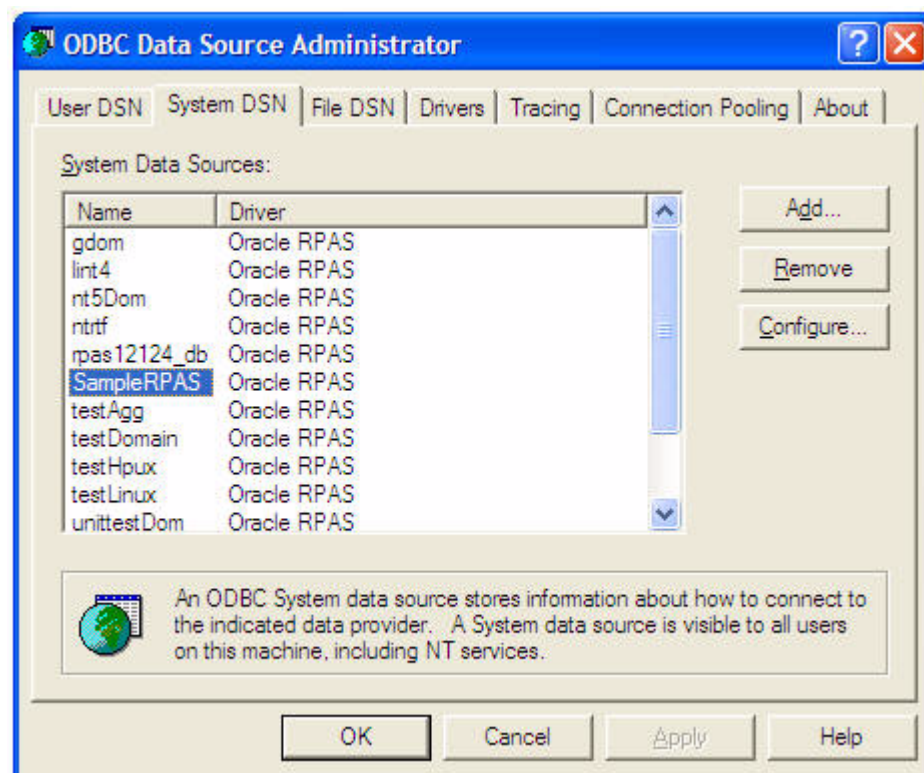
3. In the Create New Data source window, select **Oracle RPAS**.

Figure 14–15   Create New Data Source Dialog Box
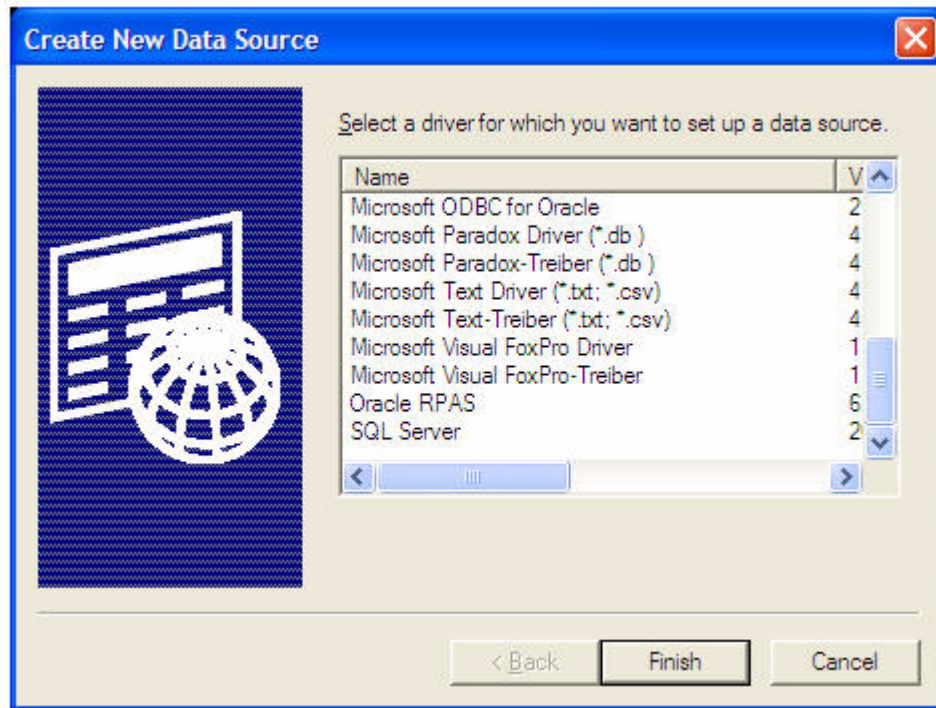
4. In the Oracle Retail RPAS ODBC Driver Setup window, enter the following information:

- Name and description of the ODBC data source.

- In the Service Host field, enter the name of the server. If connecting to a service running on a local host, enter **localhost** or the name of the local host server. If the Agent service is running on a remote server, enter the name of the remote host server.

- In the Service Port field, enter the port number that the data service is listening on.

   > **Note:** This is the port number of the data service and not the Agent service port number.

- In the Service Data Source field, enter the name of the service data source that has been confided 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

The following table lists the available custom properties (all optional):

Table 14–1   Available Custom Properties

Property Name	Description
LANGUAGE	Name of the language you use. The RPAS ODBC/JDBC driver is multi-language enabled. If data is in any language other than English, the LANGUAGE property should be set to the name of that language. If multiple languages are used in the domain, set this property to the name of the language other than English. For example, if some position names are in English and some are in Japanese, then LANGUAGE should be set to Japanese.
	The default is English.
WORKBOOK_SCHEMA	Name of the workbook you wish to connect to. If not set, the driver connects to the domain.
SHORT_DATE_FORMAT	Valid short date format used in RPAS.
DEFAULT_SCHEMA	Default schema name if the table name in 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 14–16   Oracle Retail RPAS ODBC Driver Setup Dialog Box
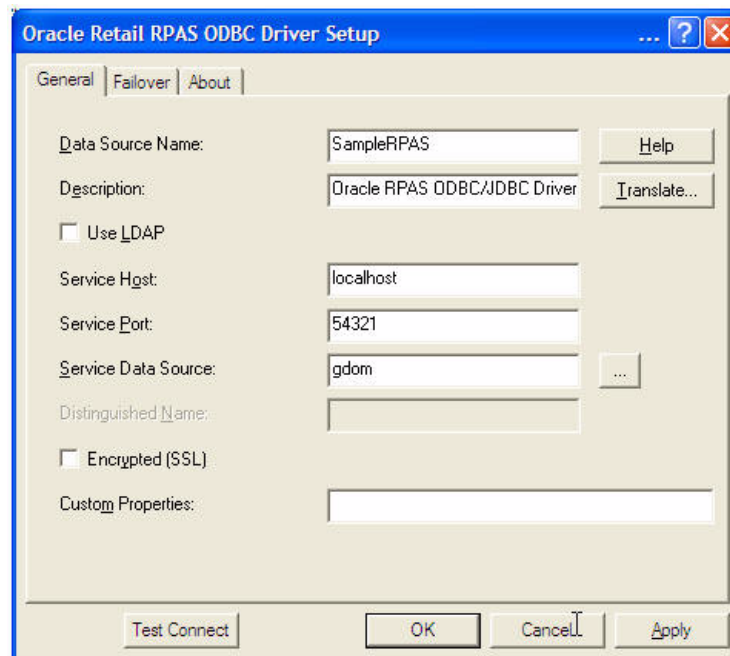
5. Click **Test Connect**. If the user security of the gdom Data source setting has been set to DBMSLogon, the Logon to Data Source dialog is displayed. Enter the data source user name and password configured for the data source.

Figure 14–17   Logon to Data Source Dialog Box
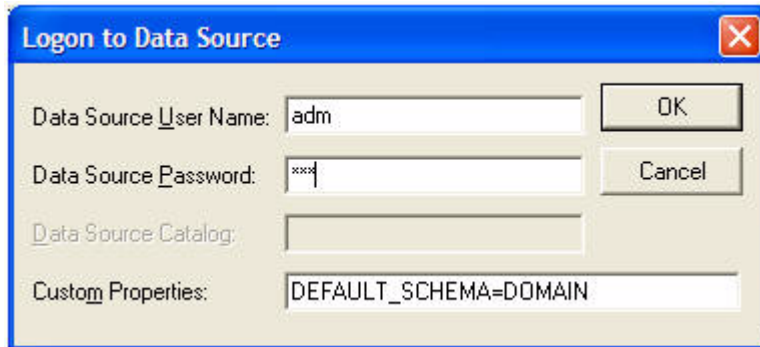
If the connection is successful, the following dialog box is displayed.

Figure 14–18   Oracle Retail RPAS ODBC Driver Setup Dialog Box for Successful Connection

6. Click **OK**.

Starting the RPAS ODBC Server Process

The RPAS ODBC Agent and Data Services should have automatically started after successful completion of the server installation.

The RPAS ODBC Data Service should be stopped and restarted using the Management Console.

1. To stop a data service, right-click the service name and then click **Stop** in the menu.

2. Right-click RPAS_Data_Service and then click **Start RPAS_Data_Service** in the menu.

Figure 14–19   Console Manager Window with Data Service Selected
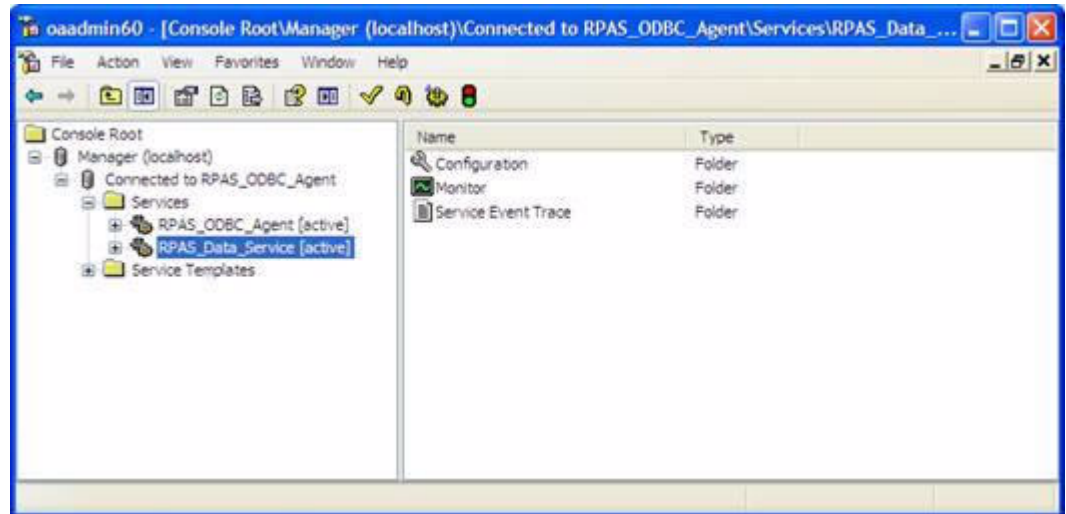
Testing the Connection Using Interactive SQL

Once 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/odbctrac.so and InstallDir to the full path of the odbcclient directory.

   b. Edit the [ODBC Data Sources] section: Add an entry for the new data source you are creating. The entry has the following format:

      ```
      MyRPASDataSource= Oracle RPAS ODBC Driver
      ```

c. Create a new [MyRPASDataSource] section: The [SampleRPAS] section can be copied and modified. In the new [MyRPASDataSource] section:

- Set Driver to the full path of odbcclient/lib/ivoa22.so.

- Set Host to the name or IP address of the server.

- Set Port to the port number the RPAS_Data_Service listens at. This is not the port number used by the RPAS ODBC Agent.

- 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 14–20   Console Manager Window with gdom Selected
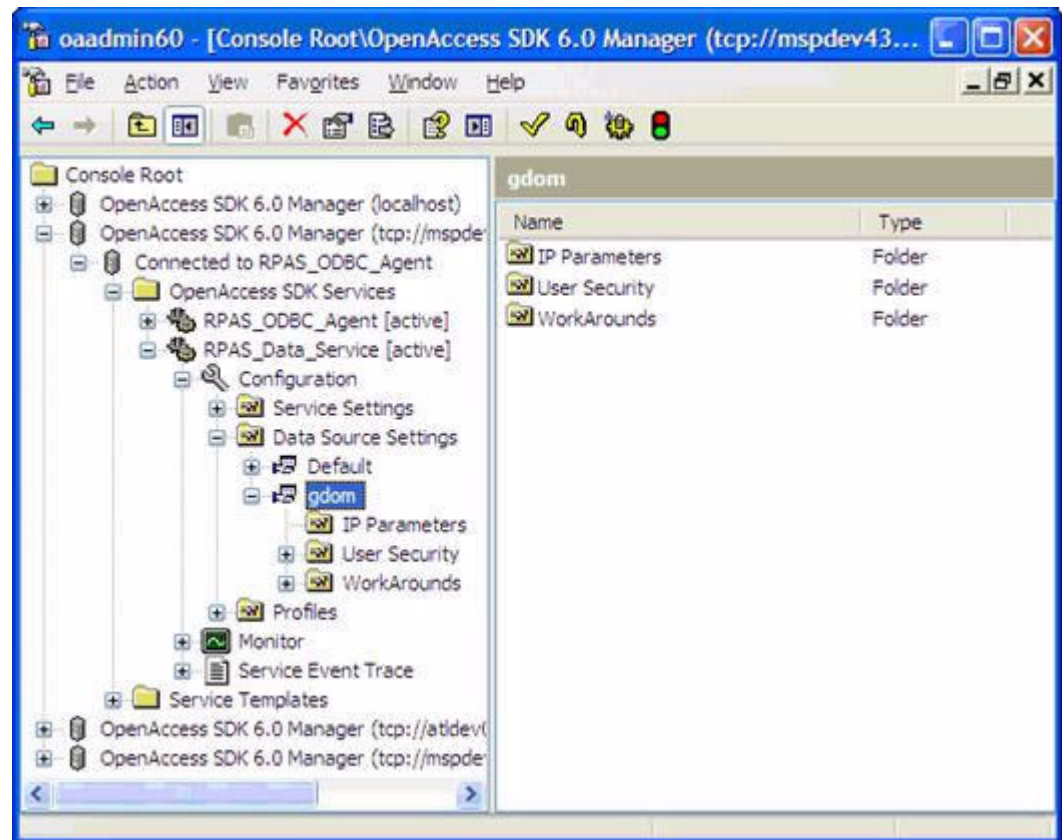
Testing the Connection

Once 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*.

Once the JDBC driver is installed on your system, you need to update the CLASSPATH environment variable. This variable ensures that the JDBC client can access the appropriate Java classes needed to connect to the database.

Updating Environment Variables for the JDBC Driver on Windows

1. Open **System** in the Control Panel. The System Properties window appears.

2. On the **Advanced** tab, click **Environment Variables**. The Environment Variables dialog appears.

Figure 14–21 Environment Variables Dialog Box
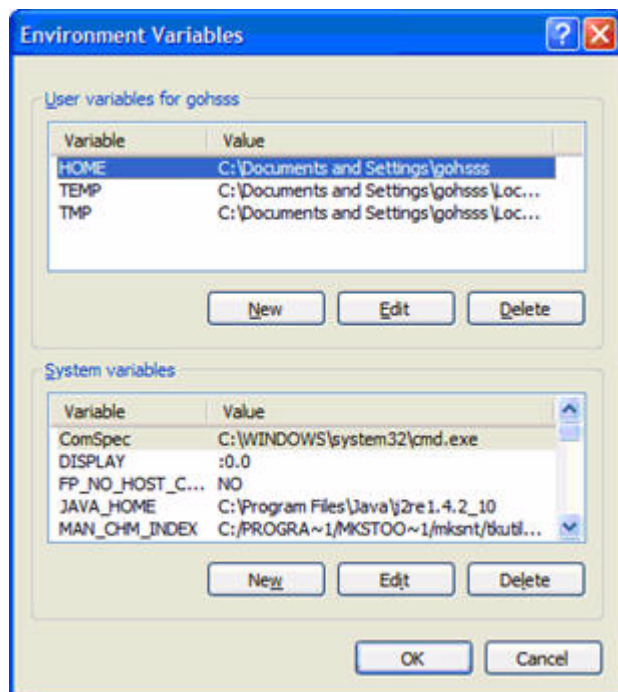
3. Select the **CLASSPATH** from the **System variables** list and click **Edit**. The Edit System Variable dialog box appears.

Figure 14–22   Edit System Variable Dialog Box
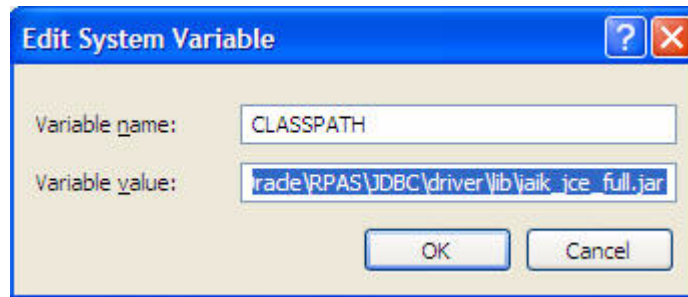
4. Add the current working directory ".", `driver_home/driver/lib/ORjc.jar`, `driver_home/driver/ORssl14.jar`, and `driver_home/driver/iaik_jce_full.jar` to the CLASSPATH environment variable and click **OK**.

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.

Once 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.

The following table lists the available attributes:

Table 14–2   Attributes Available

Key and Value	Description
log=System.out	Redirects logging to the Java output standard, System.out.
log=(file)*filename*	Redirects logging to the file specified by *filename*. For example, C:\temp\spy.log
linelimit=*numberofchars*	The maximum number of characters, specified by *numberofchars*, that Spy will log on one line. When set to no (the default), there is no maximum limit on the number of characters.
logLobs={yes ǀ no}	Specifies whether Spy logs activity on Blob / Clob. The initial default is no.
logIS={yes ǀ no ǀ 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 (the default), Spy does not log activity on InputStreams.
logTName={yes ǀ no}	Specifies whether Spy logs the name of the current thread. When set to no (the default), Spy does not log the name of the current thread.

Table 14–2   (Cont.) Attributes Available

Key and Value	Description
timestamp={yes ∣ no}	Specifies whether a timestamp should be included on each line of the Spy log.
	When set to no (the 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.jdbcspy.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/oaik_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

```java
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

Domain data queries can be performed in batch mode or intraday mode, but with intraday mode there are some limitations. If you perform a data query in intraday mode, you should restrict the number of users working with the domain, the number of users running reports, and the level of dimension detail within the reports. Detailed reports created while many users are using the domain and running reports may take longer to create and are more likely to cause concurrency issues with user operations such as committing data.

These intraday limitations do not affect data queries that are run in batch mode.

Workbook Queries

All workbook queries must be performed against saved workbooks. The workbook can be open or closed. The user has to 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 14–23   Database Diagram for All Metadata Tables in a Domain or in Each Workbook

> **Note:**   The MD_WORKBOOK_SCHEMAS table is not included in workbooks.

Fact and Dimension Tables

The following figure shows an example of the structure of fact and dimension tables and the relationships between them. A fact table represents an intersection where one or more measures' data is stored. Each measure is represented by a column in the table. Additionally, each dimension on the intersection is represented by a column. A record in the fact table is uniquely identified by a unique combination of position names for the intersecting dimensions.

A dimension table represents a dimension. It includes a column to list all position names, their labels, and their rollup mapping to each dimension at higher levels in the hierarchy.

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 14–24   Example of Star - De-normalized Schema to Represent Facts and Dimensions in RPAS
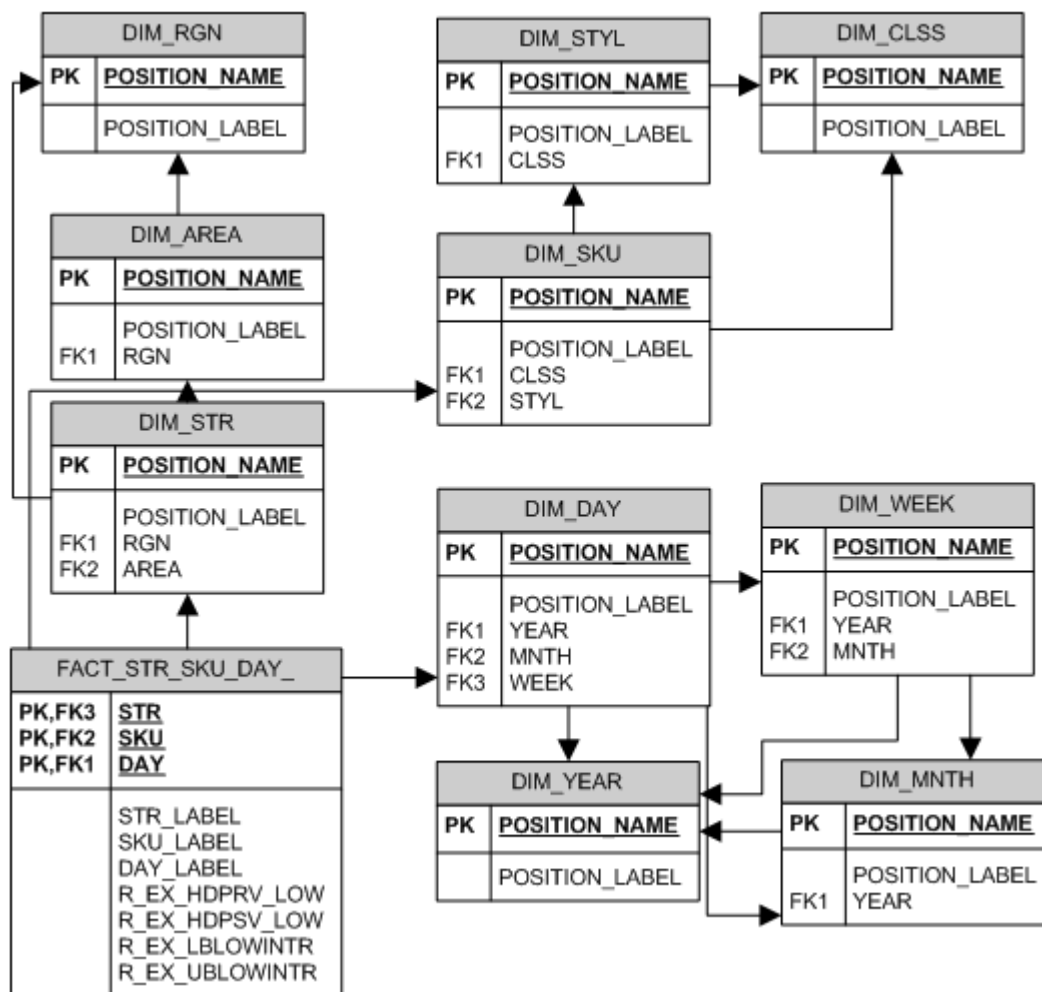
At connection time, all intersections at which any measure is stored at its base level are available as fact tables within the database. Additional aggregate level intersections may be made available in the database by specifying them in a custom connection property. These fact tables are a part of the 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 odbctrac.so that comes with RPAS odbcclient. Set InstallDir to the RPAS odbcclient installation directory.

3. In the [ODBC Data Sources] section, insert an entry for RPAS domain.

    Example:

    ```
    rpas_domain=This is the name of the data source for RPAS.
    ```
    The name here (rpas_domain) should be the same as the data source name configured in the RPAS ODBC Server.

4. Create a section in the file for the rpas_domain. The following example is subject to changes. Refer to the [SampleRPAS] section in odbc.ini or odbc64.ini under ODBC client home directory for all up-to-date settings.

    Example:

    ```
    [rpas_domain]
    Driver= absolute_path_to_odbc_client/lib[64]/ivoa22.so
    Description=Oracle Retail RPAS ODBC Driver
    Host=<RPAS ODBC Server host>
    Port=<odbc_data_service_port>
    ServerDataSource=<data_source_name>
    UseLDAP=0
    DistinguishedName=
    Encrypted=0
    LoadBalancing=0
    AlternateServers=
    ConnectionRetryCount=0
    ConnectionRetryDelay=3
    CustomProperties=
    ```
    Save your changes to the file.

Connecting OBIEE to an RPAS Domain

To connect OBIEE to a predefined DSN for an RPAS Domain:

1. Make sure the following Windows services are running:

   ■ Oracle BI Java Host

   ■ Oracle BI Server

2. Start the OBIEE Administration Tool. Select **Start**, **All Programs**, **Oracle Business Intelligence**, and then **Administration**.

3. From the File menu, select **open - online**. A window appears to 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**. Once 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.

7. Click each table that you want to import or link, and then click **OK**. If you are linking a table that does not have an index that uniquely identifies each record, then 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:

1. Start JDeveloper.

2. On the JDeveloper left panel, select the **Connections** tab.

3. Right-click on **Databases**, and select **New Database Connection**.

4. On the first screen of the Create New Database Connection wizard, enter a connection name and select **Third Party JDBC driver** for Connection Type.

5. On the second screen, enter the user name and password, and then click **Next**.

6. On the third screen, click **New** to add the driver. It opens up the following dialog. You need to find the jar files ORjc.jar, ORssl14.jar, iaik_jce_full.jar, and ORy.jar and their paths (they are made available from the installation of the JDBC Driver). Then, create a library group "RPAS JDBC Driver" with those four jar files.

In the Driver Class field, enter**com.oracle.ard.jdbc.openaccess.OpenAccessDriver**.

Then enter in the URL field: **jdbc:RPAS://{host_name}:{port_number};ServerDataSource={data_source_name}** , where **host_name** is the host name or IP address of the ODBC server, **port_number** is the RPAS_Data_Service port number and **data_source_name** is the Data Source Name which is configured on the ODBC server.

Figure 14–25   Register JDBC Driver Dialog Box
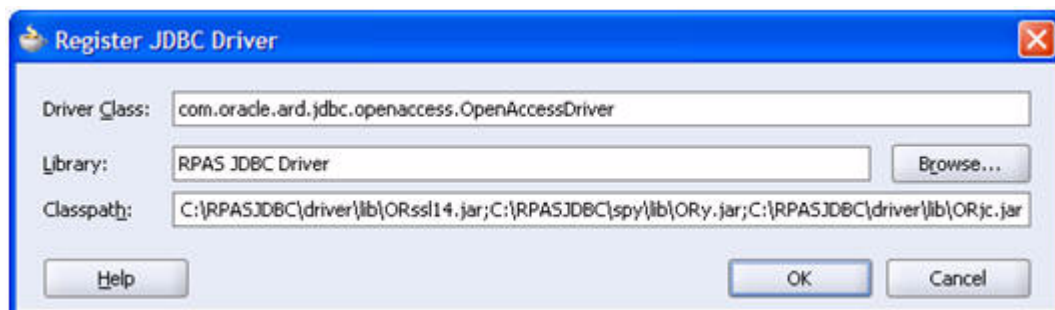
7. Follow the instructions to finish creating the connection.

Once 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
```
Once 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 14–3   Numeric Functions

Function	Description
ABS(*numeric_exp*)	Returns the absolute value of numeric_exp. For example:
	SELECT ABS(-1.0), ABS(0.0), ABS(1.0) FROM emp WHERE empno = 1;
	This returns 3 result columns with values 1, 0, and 1.
ACOS(*float_exp*)	Returns the arccosine of float_exp as an angle, expressed in radians. For example:
	SELECT ACOS(-1) FROM emp WHERE empno = 1;
	This returns 3.14159.
ASIN(*float_exp*)	Returns the arcsine of float_exp as an angle, expressed in radians. For example:
	SELECT ASIN(-1.0) FROM emp WHERE empno = 1;
	This returns -1.57079.
ATAN(*float_exp*)	Returns the arctangent of float_exp as an angle, expressed in radians. For example:
	SELECT ATAN(45.0) FROM emp WHERE empno = 1;
	This returns 1.54857.
ATAN2(*float_exp1, float_exp2*)	Returns the arctangent of the x and y coordinates, specified by float_exp1 and float_exp2, respectively, as an angle, expressed in radians. For example:
	SELECT ATAN2(35.175, 129.44) FROM emp WHERE empno = 1;
	This returns 0.2653399.
CEILING(*numeric_exp*)	Returns the smallest integer greater than or equal to numeric_exp. The return value is of the same data type as the input parameter. For example:
	SELECT CEILING(123.45), CEILING(-123.45), CEILING(0.0) FROM
	emp WHERE empno = 1;
	This returns 124, -123 and 0.
COS(*float_exp*)	Returns the cosine of *float_exp*, where *float_exp* is an angle expressed in radians. For example:
	SELECT COS(14.78) FROM emp WHERE empno = 1;
	This returns -0.59946542.
COT(*float_exp*)	Returns the cotangent of *float_exp*, where *float_exp* is an angle expressed in radians. For example:
	SELECT COT(124.78) FROM emp WHERE empno = 1;
	This returns -0.82045588.
DEGREES(*numeric_exp*)	Returns the number of degrees converted from *numeric_exp* radians. For example:
	SELECT DEGREES(3.143) FROM emp WHERE empno = 1;
	This returns 180.0806.

Table 14–3   (Cont.) Numeric Functions

Function	Description
EXP(*float_exp*)	Returns the exponential value of *float_exp*. For example: SELECT EXP(378.615) FROM emp WHERE empno = 1; This returns 2.69404760606322E+164
FLOOR(*numeric_exp*)	Returns the largest integer less than or equal to *numeric_exp*. The return value is of the same data type as the input parameter. For example: SELECT FLOOR(123.45), FLOOR(-123.45) FROM emp WHERE empno = 1; This returns 123 and -124.
LOG(*float_exp*)	Returns the natural logarithm of *float_exp*. For example: SELECT LOG(5.175643) FROM emp WHERE empno = 1; This returns 1.64396358.
LOG10(*float_exp*)	Returns the base 10 logarithm of *float_exp*. For example: SELECT LOG10(145.175643) FROM emp WHERE empno = 1; This returns 2.161893758. SELECT LOG10(0), LOG10(-1), LOG10(1) FROM emp WHERE empno = 1; This returns -1.#INF, -1.#IND and 0
MOD(*integer_exp1*, *integer_exp2*)	Returns the remainder (modulus) of *integer_exp1* divided by *integer_exp2*. For example: SELECT mod(empno, 2) FROM emp WHERE empno = 11; This returns 1.
PI( )	Returns the constant value of pi as a floating-point value. For example: SELECT PI() FROM emp WHERE empno = 1; This returns 3.14159265358979.
POWER(*numeric_exp*, *integer_exp*)	Returns the value of *numeric_exp* to the power of *integer_exp*. For example: SELECT POWER(2, -5), POWER(2, 5) FROM emp WHERE empno = 1; This returns 0, 32.
RADIANS(*numeric_exp*)	Returns the number of radians converted from *numeric_exp* degrees. For example: SELECT RADIANS(45.0) FROM emp WHERE empno = 1; This returns 0.785398.
RAND([*integer_exp*])	Returns a random floating-point value using *integer_exp* as the optional seed value. For example: SELECT RAND(0) FROM emp WHERE empno = 1; This returns 38.

Table 14–3   (Cont.) Numeric Functions

Function	Description
ROUND(*numeric_exp, integer_exp*)	Returns *numeric_exp* rounded to *integer_exp* places right of the decimal point. If *integer_exp* is negative, *numeric_exp* is rounded to |*integer_exp*| places to the left of the decimal point. For example: SELECT ROUND(123.344, 2), ROUND(123.345, 2) FROM emp WHERE empno = 1; This returns 123.34 and 123.35. SELECT ROUND(748.58, -1), ROUND(748.58, -2), ROUND(748.58, 3), FROM emp WHERE empno = 1; This returns 750, 700 and 1000.
SIGN(*numeric_exp*)	Returns the positive (+1), zero (0), or negative (-1) sign of the given expression. For example: SELECT SIGN(empno) FROM emp WHERE empno = 11; This returns 1. SELECT SIGN(-1 * empno), SIGN(0) FROM emp WHERE empno = 1; This returns two result columns with values -1 and 0.
SIN(*float_exp*)	Returns the sine of *float_exp*, where *float_exp* is an angle expressed in radians. For example: SELECT SIN(1.570796) FROM emp WHERE empno = 11; This returns 0.999999.
SQRT(*float_exp*)	Returns the square root of *float_exp*. For example: SELECT SQRT(45.35) FROM emp WHERE empno = 11; This returns 6.7342.
TAN(*float_exp*)	Returns the tangent of float_exp, where *float_exp* is an angle expressed in radians. For example: SELECT TAN(0.785398) FROM emp WHERE empno = 11; This returns 0.999999.
TRUNCATE(*numeric_exp, integer_exp*)	Returns *numeric_exp* truncated to *integer_exp* places right of the decimal point. If *integer_exp* is negative, *numeric_exp* is truncated to |*integer_exp*| places to the left of the decimal point.
NCHAR(code)	Returns the Unicode character that has the specified code as a SQL_WCHAR value. The valued of code should be between 0 and 65535. Example: "SELECT NCHAR(945)" returns the character d.

String Functions

Table 14–4   String Functions

Function	Description
ASCII(*string_exp*)	Returns the ASCII code value of the leftmost character of string_exp as an integer. For example:
	SELECT ASCII(ename) FROM emp WHERE ename = 'Adam';
	This returns 65 which is the ASCII code of A.
BIT_LENGTH(*string_exp*)	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(*code*)	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(*string_exp*) CHARACTER_LENGTH(*string_exp*)	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(*string_exp1, string_exp2*)	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'.

Table 14–4   (Cont.) String Functions

Function	Description
INSERT(*string_exp1, start, length, string_exp2*)	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(string_exp) LOWER(string_exp)	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'.
LEFT(*string_exp, count*)	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:
	SELECT LEFT(ename, 2) FROM emp WHERE ename = 'John';
	This returns 'jo'.
	SELECT LEFT(ename, 2) FROM emp WHERE ename = N'John';
	If ename is a column of wide character data type, this returns wide character string N'jo'.
LENGTH(*string_exp*)	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:
	SELECT LENGTH('John '), LENGTH('John') FROM emp;
	This returns 4 for both result columns as trailing blanks are excluded.
	SELECT LENGTH(N'John '), LENGTH(N'John') FROM emp;
	This returns 7 for the first result column and 4 for the second result column. Trailing blanks are not checked in wide character implementation.

Table 14–4   (Cont.) String Functions

Function	Description
LOCATE(*string_exp1*, *string_exp2*[, *start*])	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:
	SELECT LOCATE('h', 'John', 1) FROM emp;
	This returns 3 as 'h' is the found at the third position.
	SELECT LOCATE(N'h', N'John', 1) FROM emp;
	This returns 3 as N'h' is the found at the third position.
LTRIM(*string_exp*)	Returns the characters of string_exp, with leading blanks removed. For example:
	SELECT LTRIM(' ABC') FROM emp;
	This returns 'ABC'.
OCTET_LENGTH(*string_exp*)	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:
	SELECT OCTET_LENGTH(ename) FROM emp WHERE ename =
	'John';
	This returns 4.
POSITION(*character_exp1 character_exp2*)	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 character_exp1 and character_exp2 are wide character strings, returns the position of the first wide character expression in the second wide character expression. If character_exp2 is a wide character string, the result is computed by treating both arguments as wide character strings. For example:
	SELECT POSITION('abc', '1234abc def') FROM emp;
	This returns 5.
	SELECT POSITION(N'abc', N'1234abc def') FROM emp;
	This returns 5.

Table 14–4   (Cont.) String Functions

Function	Description
REPEAT(*string_exp, count*)	Returns a character string composed of string_exp repeated count times.
	If string_exp is wide character string, the return value is a wide character string. For example:
	SELECT REPEAT(ename, 2) FROM emp WHERE ename = 'John';
	This returns 'JohnJohn'
	SELECT REPEAT(ename, 2) FROM emp WHERE ename = N'John';
	If ename is a column of wide character data type, this returns N'JohnJohn'.
REPLACE(*string_exp1, string_exp2, string_exp3*)	Search string_exp1 for occurrences of string_exp2, and replace with string_exp3.
	If string_exp1 is wide character string, the return value is a wide character string. For example:
	SELECT REPLACE(address, 'San Francisco', 'SFO') FROM emp
	where address = '100 Vanness, San Francisco';
	This returns '100 Vanness, SFO'.
	SELECT REPLACE(address, N'San Francisco', N'SFO') FROM
	emp WHERE address = N'100 Vanness, San Francisco';
	If address is a column of wide character data type, this returns N'100 Vanness, SFO'.
RIGHT(*string_exp, count*)	Returns the right-most 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:
	SELECT RIGHT(ename, 2) FROM emp WHERE ename = 'John';
	This returns 'hn'.
	SELECT RIGHT(ename, 2) FROM emp WHERE ename = N'John';
	If ename is a column of wide character data type, this returns N'hn'.
RTRIM(*string_exp*)	Returns the characters of string_exp with trailing blanks removed. For example:
	SELECT RTRIM('abc ') FROM emp;
	This returns 'abc'.
SPACE(*count*)	Returns a character string consisting of count spaces. For example:
	SELECT ename+space(5)+ename FROM emp WHERE ename =
	'John';
	This returns 'John John'.

Table 14–4   (Cont.) String Functions

Function	Description
SUBSTRING(*string_exp, start*, length) SUBSTR(*string_exp*, length)	Returns a character string that is derived from string_exp, beginning at the character position specified by start for length characters. If string_exp is wide character string, the return value is a wide character string. Offset (start and length) 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 ename is a column of wide character data type, this returns N'Joh'
UCASE(*string_exp*) UPPER(*string_exp*)	Returns a string equal to that in string_exp, with all lowercase characters converted to uppercase. For example: SELECT UCASE(ename) FROM emp WHERE ename = 'John'; This returns 'JOHN'.
UNICODE(string_exp)	Returns the Unicode code of the first character of the string_exp as a SQL_INTEGER value. Example: "SELECT UNICODE('deʤ')" returns an integer value of 945.

Time / Date Functions

Table 14–5   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 14–5   (Cont.) Time / Date Functions

Function	Description
DATEADD(*datepart , number ,date*) TIMESTAMPADD(*datepart , number , date*)	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 *date* value. *datepart*: the parameter that specifies on which part of the date to return a new value. Both ODBC notation and SQL Server notation for datepart 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. *number*: the value used to increment the datepart. If value is not an integer, the fractional part of the value is discarded. For example, if you specify **day** for *datepart* and**1.75** for *number*, *date* is incremented by 1. *date:* an expression that returns a date or timestamp value or a character string in a date-time format.

Table 14–5   (Cont.) Time / Date Functions

Function	Description
DATEDIFF(*datepart , startdate , enddate*) TIMESTAMPDIFF(*datepart , startdate , enddate*)	Returns the number of date and time boundaries crossed between two specified dates. *startdate* is subtracted from *enddate*. If *startdate* is later than *enddate*, a negative value is returned. *datepart*: the parameter that specifies on which part of the date to calculate the difference. Both ODBC notation and SQLServer notation for datepart 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 dayofyear, millisecond and factional second specifications. For example: 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'; This returns 2 (assuming that the curdate() returns year 2002).
DAYNAME(*date_exp*)	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 *date_exp*. For example: SELECT DAYNAME('2002-01-01'), DAYNAME('2002-01-02') FROM emp; This returns 'Tuesday' and 'Wednesday'.
DAYOFMONTH(*date_exp*)	Returns the day of the month based on the month field in *date_exp* as an integer value in the range of 1-31. For example: SELECT DAYOFMONTH('2002-01-05') FROM emp; This returns 5.
DAYOFWEEK(*date_exp*)	Returns the day of the week based on the week field in *date_exp* as an integer value in the range of 1-7, where 1 represents Sunday. For example: SELECT DAYOFWEEK('2002-01-05') FROM emp; This returns 7.
DAYOFYEAR(*date_exp*)	Returns the day of the year based on the year field in *date_exp* as an integer value in the range of 1-366. For example: SELECT DAYOFYEAR('2002-01-05') FROM emp; This returns 5.

Table 14–5   (Cont.) Time / Date Functions

Function	Description
HOUR(*time_exp*)	Returns the hour based on the hour field in *time_exp* as an integer value in the range of 0-23. For example: SELECT HOUR('22:20:20') FROM emp; This returns 22.
MINUTE(*time_exp*)	Returns the minute based on the minute field in *time_exp* as an integer value in the range of 0-59. For example: SELECT MINUTE('22:21:20') FROM emp; This returns 21.
MONTH(*date_exp*)	Returns the month based on the month field in *date_exp* as an integer value in the range of 1-12. For example: SELECT MONTH('2002-01-05') FROM emp; This returns 1.
MONTHNAME(*date_exp*)	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 Januar through Dezember for a data source that uses German) for the month portion of *date_exp*. 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(*date_exp*)	Returns the quarter in *date_exp* 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(*time_exp*)	Returns the second based on the second field in *time_exp* as an integer value in the range of 0-59. For example: SELECT SECOND('22:21:20') FROM emp; This returns 20.
WEEK(*date_exp*)	Returns the week of the year based on the week field in *date_exp* as an integer value in the range of 1-53. For example: SELECT WEEK('2002-01-05') FROM emp; This returns 1.
YEAR(*date_exp*)	Returns the year based on the year field in *date_exp* as an integer value. For example: SELECT YEAR('2002-01-01') FROM emp; This returns 2002.

System Functions

Table 14–6   System Functions

Function	Description
DATABASE( )	Returns the name of the database corresponding to the connection handle. (The name of the database is also available by calling SQLGetConnectOption with the SQL_CURRENT_QUALIFIER connection option.)
USER( )	Returns the user name in the DBMS. (The user name is also available using SQLGetInfo by specifying the information type: SQL_USER_NAME.) This can be different than the login name.

Aggregate Functions

Aggregate functions return a single row based on groups of rows, rather than on single rows.

Table 14–7   Aggregate Functions

Function	Description
AVG([ALL | DISTINCT] *expression*)	Returns the average of the values in a group. Null values are ignored.
SUM([ALL | DISTINCT] *expression*)	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] *expression*] | *})	Returns the number of items in a group. COUNT(*) returns the number of items in a group, including NULL values and duplicates. COUNT(ALL *expression*) evaluates *expression* for each row in a group and returns the number of non-null values. COUNT(DISTINCT *expression*) evaluates *expression* for each row in a group and returns the number of unique, non-null values.
MAX([ALL | DISTINCT] *expression*)	Returns the maximum value in the expression.
MIN([ALL | DISTINCT] *expression*)	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 14–8   CAST Function

Keyword	Value
CHAR	255
BINARY	255
BIT	1
DATE	6
DOUBLE	8
FLOAT	8
INTEGER	4
LONGVARBINARY	1000000
LONGVARCHAR	1000000
NUMERIC	34
SMALLINT	2
REAL	4
TIME	6
TIMESTAMP	16
TINYINT	1
VARBINARY	1024
VARCHAR	1024
WLONGVARCHAR	2000000
WVARCHAR	512
WVARCHAR	2048
CHAR, BINARY	255

Table 14–8   (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 14–9   CONVERT Function-Length

Keyword	Length
SQL_CHAR	256
SQL_BINARY	256
SQL_BIT	1
SQL_DATE	6
SQL_DOUBLE	8
SQL_FLOAT	8
SQL_INTEGER	4
SQL_LONGVARBINARY	1000000
SQL_LONGVARCHAR	1000000
SQL_NUMERIC	34
SQL_SMALLINT	2
SQL_REAL	4
SQL_TIME	6
TIMESTAMP	16
TINYINT	1
VARBINARY	1024
VARCHAR	1024
WLONGVARCHAR	2000000
WVARCHAR	512
WVARCHAR	2048

Table 14–10   CONVERT Function-Precision

Keyword	Precision
SQL_BINARY	255
SQL_BIT	1
SQL_CHAR	255
SQL_DATE	10
SQL_DOUBLE	15
SQL_FLOAT	15
SQL_INTEGER	10
SQL_LONGVARBINARY	1000000
SQL_LONGVARCHAR	1000000
SQL_NUMERIC	32
SQL_REAL	7
SQL_SMALLINT	5
SQL_TIME	8
SQL_TINYINT	3
SQL_VARBINARY	1024
SQL_VARCHAR	1024
SQL_WCHAR	255
SQL_WVARCHAR	2048
SQL_WLONGVARCHAR	1000000

Table 14–11   CONVERT Function-Scale

Keyword	Scale
SQL_NUMERIC	5
All other types	0

Unsupported SQL Functions

The SQL engine of the RPAS ODBC Server implements a large portion of the entry level SQL as defined in the X3.135-1992, "Database Language SQL" specification and commercial databases like SQL Server and Oracle. It is compliant with the ODBC minimal grammar specification.

This section describes the un-supported 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 FROMextract-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.

15

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 15–1   JMS System in RPAS

Configuring Subjects and Measures for Monitoring

A subject is defined as a logical grouping of measures. A subject is mapped to one or more measures in one domain, while a measure can belong to one or more subjects. For example, a subject named PlanningMeasures can include all measures associated with planning.

The configuration file of measure change monitoring serves two purposes:

- Defines the mappings between subjects and measures.

- Defines the inclusion filters for monitored measures.

By default, no measures are monitored unless they are specifically included in the configuration.

For each domain, there is one measure change monitor configuration file. It is named MeasureChangeMonitor.properties and must be located under {domainPath}/config/.

> **Note:**   There is only one property file for a global domain. It is placed under the config directory of the master domain.

This configuration file is in standard Java properties file format. Each line defines the relationship between one subject and one measure:

```
Subject.MeasureName=true
```
- `true` means inclusion in monitoring.

- `false` is used to exclude a subject/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 15–1   Configuration Macros

Macro Name	Format	Notes
__EventDateTime__	YYYY-MM-DDThh:mm:ss	Local time of the server.
__SourceURI__	RPAS/JMS/{hostname}	{hostname} is the name of the server where the publisher is running.
__SUBJECT__	String	Subject of the event as defined in MeasureChangeMonitor.properties.
__TYPE__	"MeasureChange"	Type is a constant string.
__DOMAIN__	String	Path to the domain.
__MEASURE__	String	RPAS internal measure name.
__ORIGINUSER__	String	RPAS User ID, only available for workbook commit or DPM. Use "-" if not available.

Following 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 15–2   Arguments Used in the Command Line

Argument	Description
classpath	Should include rpaspmc.jar and jms.jar under $RPAS_HOME/lib and any vendor-specific JMS implementation jar files.
-d *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 15–3   Vendor Neutral Configuration Settings for the JMS Publisher

Setting Name	Description	Required	Value
topic	JMS topic lookup name.	Yes	A string with only alphanumeric characters, underscores, and periods. Maximum length is 100 bytes.
topicConnection Factory	JMS topic connection factory lookup name.	Yes	A string with only alphanumeric characters, underscores, and periods. Maximum length is 100 bytes.
messageTempla te	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.
restartAfterExce ption	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.deliver yMode	Delivery mode.	No	NON_PERSISTENT or PERSISTENT. Default is PERSISTENT
message.priorit y	A priority number for the JMS message.	No	0 to 9. Default is 4.

Table 15–3 (Cont.) Vendor Neutral Configuration Settings for the JMS Publisher

Setting Name	Description	Required	Value
message.timeTo Live	Time to live in milliseconds.	No	0 or any positive integer. Default is 0 (unlimited).

Vendor-Specific Settings

Sun Open Message Queue is used as an example for these settings. For further details, consult the vendor documentation for your JMS implementation.

File Based JNDI Object Store

File based JNDI object store is used primarily for development and testing. It is very easy to set up, but has weak built-in security.

Table 15–4 File Based JNDI Object Store

Property Name	Description	Required	Value
java.naming.factor y.initial	Initial context for JNDI lookup.	Yes	For Open Message Queue, it must be: com.sun.jndi.fscontext.RefFSCont extFactory
java.naming.provi der.url	Directory path to the object store.	Yes	Example: file:///C:/myapp/mqobjs

LDAP Based JNDI Object Store

An LDAP server is the recommended object store for production JMS messaging systems. LDAP servers are designed for use in distributed systems and provide security features that are required in production environments.

Table 15–5 LDAP Based JNDI Object Store

Property Name	Description	Required	Value
java.naming.factor y.initial	Initial context for JNDI lookup.	Yes	For Open Message Queue, it must be: com.sun.jndi.ldap.LdapCtxFactor y
java.naming.provi der.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.securit y.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.

Table 15–5   (Cont.) LDAP Based JNDI Object Store

Property Name	Description	Required	Value
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.
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 15–6   General Settings for the JMS Subscriber

Property Name	Description	Required	Value
topic	JMS topic lookup name.	Yes	A string with only alphanumeric characters, underscores, and periods. Maximum length is 100 bytes.
topicConnectionFactory	JMS topic connection factory lookup name.	Yes	A string with only alphanumeric characters, underscores, and periods. Maximum length is 100 bytes.
logLevel	Logging level. A log level is a cut-off level which means logs with a lower level are filtered out.	No	Specify one of the following values, in low to high order. ■ VERBOSE: All logs. ■ DEBUG: Debug logs. ■ INFO: Informational logs. ■ WARN: Warning logs. ■ ERROR: Error logs. ■ SUPPRESS: No logs. If not specified, default is INFO.
logFile	The path to the log file. Can be a relative or absolute path.	No	If not specified, output to the console.
clientID	JMS Client ID. Only required for a durable subscription. Not recommended for other use.	No	This must be a unique string for all JMS clients using the same connection factory.
durableSubscriptionName	A unique name for a durable subscription.	No	If a name is not provided, the subscription will be a transient subscription. This means messages are not queued up if the connection for the subscriber is lost or the subscriber is not running.

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
```

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*.

A

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

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
-d *domainpath*	Path to the domain.
-s	Set defaults.
-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`
-noheader	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.

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 *domainpath*	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`

B

Appendix: Manually Configuring Single Sign-On

Oracle Single Sign-On (SSO) is configured for the RPAS Fusion Client through the RPAS installer. However, you can manually configure SSO for the Fusion Client.

> **Note:** Manual configuration is not recommended for beginners.

For information about using the RPAS installer to configure SSO, see the *Oracle Retail Predictive Application Server Installation Guide*.

To convert a non-SSO deployment of the RPAS Fusion Client to use SSO, perform the following steps.

1. Locate the Fusion Client ear file called rpas.ear.

2. Make a temporary folder called tmp. In that folder, extract the rpas.war file from the rpas.ear file.

   ```
   mkdir tmp
   cd tmp
   jar xvf ../rpas.ear rpas.war
   ```

3. Extract the web.xml and weblogic.xml files from the rpas.war file:

   ```
   jar xvf rpas.war WEB-INF/web.xml
   jar xvf rpas.war WEB-INF/weblogic.xml
   ```

4. Edit web.xml as follows:

   a. From the beginning of the file, search for the string SSO. The first match is a `security-constraint` XML element. Uncomment this element.

      To uncomment the element, modify the element from this:

      ```
      <!--element start>…</element end-->
      ```

      To this:

      ```
      <element start>...</element end>
      ```

   b. Search for the next occurrence of `SSO`. This match is the `login--config` XML element. Uncomment this element.

   c. Save the web.xml file.

5. Edit weblogic.xml as follows:

   a. From the beginning of the file, search for `SSO`. This locates the `security-role-assignment` XML element. Uncomment this element.

   b. Save the weblogic.xml file.

6. Update rpas.war while in tmp folder:

```
jar uvf rpas.war WEB-INF/web.xml
jar uvf rpas.war WEB-INF/weblogic.xml
```

7. Update rpas.ear while in tmp folder. Delete the tmp folder after you are finished.

```
jar uvf ../rpas.ear rpas.war
cd ..
rm -rf tmp
```

8. Update config.properties in the Fusion Client configroot folder.

   a. Set isSSO=true.

   b. Save.