Oracle® Retail Predictive Application Server
Release Notes
Release 11.2.2
January 2006

ORACLE

Customer Support

- https://metalink.oracle.com

When contacting Customer Support, please provide:

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

Release Notes

The following section describes the RPAS 11.2.2 patch. This includes issues that have been addressed in this patch, enhancements and new functionality, changes to existing functionality, and known issues. This information covers the RPAS server, RPAS client, and Configuration Tools components of RPAS 11.2.

The client build number (found in Help → About) for RPAS 11.2.2 is 1.257.

Installing 11.2.2 Server and Tools on Windows Environments

Individuals that are installing the server and configuration tools components of RPAS 11.2.2 on Windows machines (NT, 2000, XP) have two options to install the components:

- using an InstallShield package
- using the rsp_manager

Oracle recommends using the InstallShield package if you are installing RPAS 11.2.2 for the first time and are not upgrading from a previous version of RPAS 11.x. The InstallShield package will install the RPAS server and tools on the local machine and set (and overwrite) all necessary environment variables.

Oracle recommends using the RSP Manager if you are upgrading from a previous version of RPAS 11.x and you wish to have this utility upgrade one or more domains. Use of the RSP Manager is also advised for users who run multiple versions of RPAS 11 as the RSP Manager allows you to manually manage the environment variables.

Enhancements and Changes

New Aggregation Method - "hybrid"

'Hybrid' is a new aggregation method that allows the aggregation of a measure to be calculated differently up different hierarchies. Therefore, it is approximately functionally equivalent to the aggregate procedure. Measures that use the hybrid aggregation type cannot be manipulated above their base intersection because there is no mechanism to spread changes, but it may be used with the old modifier where measures that use the aggregate procedure cannot. The hierarchy aggregation parameters are specified in a new measure property ─ '**Agg Spec**.' See the RPAS Configuration Guide for more information.

Solution Extensions

Solution Extensions offer special functionality that may be configured in its own solution or embedded within any RPAS solution. With this patch, the following Solution Extensions are available for configuration as part of the RPAS platform:

- Forecast Procedure
- LostSale Function
- Cluster Procedure
- ASO Space Function
- ASO Assort Function

See the RPAS 11.2.2 Solution Extension Configuration Guide for information on configuring these extensions using the RPAS Configuration Tools.

Changes to the utility that is used to move data from one domain to another – 'mapdata'

Specifically, mapdata copies data from an existing domain, database, or array to a new domain, database or array.

The usage of the RPAS utility mapdata was modified in this patch. Changes include:

Previous Argument	Current Argument	Function
–srcDomain	–d	Specifies the path of the domain from which data is being moved.
–destDomain	–dest	Specifies the path to the domain to which data is being moved

The complete usage of this utility is as follows:

```
mapdata -d SrcPath -dest destPath [-db dbName [-array arrayName]]

-d      srcPath   --> path to source domain
-dest   destPath  --> path to destination domain
-db     dbName    --> apply mapData only on the given database
-array  arrayName --> apply mapData only on the given array.
(Requires -db option)
```

Run mapData without -db and -array arguments to map the entire domain.

Miscellaneous Changes to the Configuration Tools

- **Position Format ─** There were some changes to the Position Format, which is the date/time format used for the names of positions in the root dimension of the CLND (Calendar) hierarchy.

  Changes include:

  — a new title ─ "Position Format for Calendar's Root Dimension"

  — special characters are not allowed

  — system now validates that there is a dimension named DAY in the CLND hierarchy, but DAY does not have to be root

- **Rule Group Simulator ─** The rule group simulator enables the verification of the interaction between measures from within the configuration tools.

  The rule group simulator no longer fails with a spurious error saying that the measures should have an agg type of recalc when there is a rule that uses the "aggregate" procedure in the rule group.

Changes to the Utility that is used to Load Data into the Domain – 'loadmeasure'

Loading Multiple Measures from a Single Input File

The loadmeasure utility was updated to allow multiple measures to be loaded from a single input file in a single call to the utility. Previously, each measure had to be specified in individual calls to the utility. This change was made primarily for performance reasons. When loading a measure into the master domain of a Global Domain environment, the input file is split by the loadmeasure utility so that the data can be loaded into the individual local domains (if the data resides in the local domains). This enhancement ensures that an input file is only read and split once when loading into multiple measures.

The complete usage of the loadmeasure utility is as follows:

Loading Measure Data – loadmeasure

The **loadmeasure** utility is used to load data into the domain. You must specify the measure name(s) and the path to the domain that contains the measure(s).

In a Global Domain environment, **loadmeasure** is centralized and can only be called in the master domain. The 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 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.

Usage:

```
loadmeasure –d pathToDomain –measure
measureName{,measureName,…} {–logdirectory directoryName} {–
applyloads} {–loglevel level}
```

Argument	Description
–d *pathToDomain*	Specifies the domain in which to load the measure.
–measure *measureName*	Specifies the name of the measures to load. The name must be lowercase. If more than one measure is given, all the measures must have the same input file.
–logdirectory *directoryName*	Specifies the location of the output error log. The default location is pathToDomain/scripts/err.
–applyloads	Use this argument to apply any staged loads for the named measure.
–loglevel *level*	Use this argument to set the logger verbosity level. Possible values: all, profile, debug, information, warning, error, or none.
–processes *count*	For running the utility in parallel. Specifies the maximum number of child loadmeasure processes to be running concurrently across the local domains in a Global Domain environment. If this argument is omitted or is less than two, 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, so you will actually have count + 1 processes.

The loadMeasure utility also provides functionality to automatically run scripts before and after the utility is executed. These are referred to as pre- 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 are run on the domain in which the utility has been called prior to the execution of the utility. Likewise, after the utility has completed running, it checks for the existence of scripts named "post<*measurename*>.sh" and executes them immediately. In the case of multiple measures, only the preprocessing script for the first listed measure will have any affect on the data.

Rerun the alert finder

The Alert Finder checks the current data against the list of alerts, and it builds a list of those alerts. It counts the number of times the alert is triggered and stores this information for a particular user. The alert finder is typically run as part of a batch process performed when system use is minimal.

Customers using the alert functionality in RPAS must run the alert finder (alertmgr utility with the –findAlerts argument) after applying this patch. This is required because the RPAS client was fixed in this patch to be able to distinguish between two scenarios:

1. When no alert hits are found
2. When the alert threshold (the maximum number of alerts) is reached

Creating a Subset of a Domain Environment

The RPAS utility **copyDomain** is used to perform a number of tasks related to moving and copying domain environments. This utility was enhanced with a new argument (**–clone**), which is used to copy a subset of a domain environment. When using this parameter, you can specify individual positions along multiple hierarchies. These positions and their corresponding data will be copied to the new domain.

Usage:

```
copydomain –clone dimposlist
```

Where *dimposlist* is a list of positions in the format:

```
dim1,pos1,…,posn:dim2,pos1,…,posn:…
```

Notes:

Only one dimension for each hierarchy can be specified

The –clone argument is supported in both global domain environments and simple domains

Existing hierarchy input files can be backed up in the same directory if –skipInput is not specified

The –copyWorkbooks argument can be used with –clone; if copyWorkbooks is not specified then all workbooks are copied

Resolved Issues

The following section contains defects that have been fixed in this patch for the RPAS server, RPAS client, and Configuration Tools.

> **Note:** Due to time constraints, the following defects with asterisks (*) had not been tested at the time this document was prepared.

Resolved Issue	Defect
A fix was made so that when Synchronized Page Scrolling is enabled, position queries are executed in other worksheets as slices are being changed in main worksheet.	*4857674
A fix was made to client to ensure that a workbook's format is saved after you create a percent-to-parent measure.	*4857510
The RPAS Server was updated so that when an extended measure is created for percent-to-parent and the measures label is also overridden in the workbook, the label override will carry over to the percent-to-parent measure.	*4856297
A update was made to the ConstParseNode class member function that returns the value of the node to either return a valid value or throw an error describing the reason the value could not be computed.	*4856101
An update was made to loadHier to remove usage information pertaining to '-purgeAge,' which serves no useful purpose.	4845382
A fix was made to the RPAS code to allow calculations to be performed on rules that contain two expressions. Problems occurred when changing the measures affecting the second expression in the rule.	4845647
The RPAS Server was updated so that measures that are not assigned to a database will not come up in the measure analysis template.	4845961
The RPAS Server was updated so that when running loadmeasure with a .rpl file extension that has 0kb, it will update the measure information by replacing or clearing out all of the measure data information.	4846775
The RPAS Configuration Tools were updated so that when building a domain using the updateStyles option, the measures that do not have styles assigned to them will not be affected.	4847370
A fix was made to RPAS code regarding WorkbookTemplate.cpp and WorkbookTemplateGroup.cpp to use ConstLogicalIterators instead of using LogicalIterator. Therefore, databases are only required to be in read mode.	4848708
The RPAS Configuration Tools were updated so that you no longer get a null pointer exception occurring when searching for a measure on refresh.	4848717
The RPAS Configuration Tools were updated so that in the Workbook tool within the Configuration Tools, edited measures are not removed when switching from the workbook tool to the rules or measures tool.	4849709

Resolved Issue	Defect
The RPAS Client was updated so that the tree icons correctly line up with the text in workbooks.	4850599
The RPAS Client was updated so that when the manual calculation setting is selected, the client will not perform an automatic calculation when selecting various menu functions and canceling them.	4850980
A workaround was implemented in the client so that the correct values are being returned in the Fetch call after the calculation cycle is initiated on the Forecast Administration/Final Level Parameters worksheet.	4853548
The RPAS Configuration Tools were updated so that when running a patchInstall on a domain, it will not error when the 'single hier select' (scriptname) property measure is modified.	4853616
The RPAS Configuration Tools were updated so that a configuration can now be opened if it contains a measure import entry for a measure that was removed from the configuration (in its native solution).	4854929
A fix was made to client to activate the Save workbook functionality after a refresh or any custom menu rule group has been initiated.	4855377
The RPAS Client was updated so that position queries will be executed when opening a saved workbook.	4856294
The RPAS Configuration tools and Server were updated so that when running an upgrade on the domain, the measure formatting that was saved will be kept.	4856934
The RPAS Configuration Tools was updated so that it does not error when running the simulator on a load group that contains a Recalc measure.	4857219
The RPAS Configuration tools were updated so that when you have two solutions open and there are errors in the task list for both solutions; when you close one of the solutions, it will now remove that solution's errors from the task list.	4858093
The RPAS Server was updated so that when security is enabled on the product hierarchy, it will not cause errors when attempting to build workbooks on local domains.	4858094
A fix was made to 'Print Preview/Print' to correct the code from erroneously initializing pointers to 0 instead of the value pointed to.	4871032
The RPAS Client was updated so that you no longer receive unexpected exceptions or error messages when selecting the print preview option for a workbook.	4874289 & 4871320
The RPAS Client has been updated so that you no longer receive an error when attempting to insert a measure into a workbook of a local domain environment.	4874364
A fix was made to the RPAS code to correct a problem with regmeasure API crashing while attempting to update a measure for special values if the measure is a non-stored HBI measure.	*4851265

Resolved Issue	Defect
A fix was made to enable you to make edits and calculations within saved workbooks.	*4850461
The RPAS Server was updated so that when running the loadmeasure utility, you no longer get an exception on forced non-higher based intersection measures.	*4848493
A fix was made to enable users to commit data successfully within the Hierarchy Maintenance workbook on upgraded domains.	*4847925
The RPAS Configuration Tools were updated so that an external measure will show as valid or invalid based on whether the 'real' instance of the measure (the one that is not imported) is valid or not as a dimension attribute.	*4845465
This issue is related to a current limitation of the RPAS calculation engine that occurs under rare circumstances with non-conforming recalc measures. The issue can be addressed with a change to the configuration. Additional information about this limitation is in the section below, Known Issues → Non-conforming recalc measures.	*4845982

Known Issues

Issues that cannot currently be addressed

The following issues exist in either RPAS or the Configuration Tools and either cannot or are not likely to be addressed in the near future.

Function validation – There is no validation for the "propspread" and "passthrough" functions in the Configuration Tools. Specifically, the Tools cannot validate that the correct number of arguments have been supplied. Both should have a predictable number of arguments based on the number of measures being calculated (for example, propspread should have one more argument than measures being calculated). See the RPAS Rule Functions Reference Guide for the syntax of these functions.

Under specific circumstances, you currently cannot resize the workbook and worksheet windows in Workbook Designer of the Configuration Tools. Using the current release of the Tools, once you resize the window that displays the list of workbooks and worksheets in the Workbook Designer, you are not able expand the window beyond a very small area when you are in a particular view (depends on several miscellaneous factors). There is a technical issue with Java that currently prevents this from easily being addressed (defect 377487).

Non-conforming recalc measures

Under complicated and rare circumstances, incorrect values can be calculated when an expression uses a recalc measure above its base intersection. The problem does not usually arise when a recalc measure is used above its base intersection because it depends on the rule group configuration and the exact changes made by the user. More detail is given in the RPAS Configuration Guide. The configuration tools now issue the following warning when a recalc measure is used above its base intersection. The issue can be avoided by replacing the recalc measure in the expression with the expression used to calculate it.

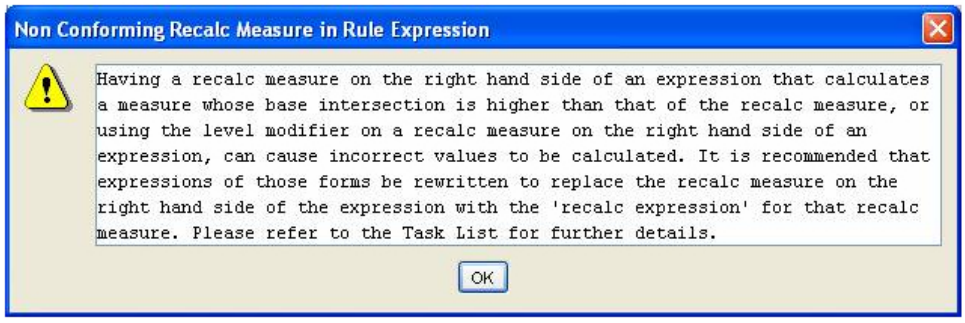

Notable Outstanding Defects

Known Issues	Defect
Specific versions of Java 1.4.1 on AIX are causing the configuration copy stored in the domain to get corrupted during the final steps of the install or patch process.	N/A (old defect # 393335)
Users running the RPAS server on Windows/NT are not currently able to use the relative time period functionality in the wizard process.	N/A (old defect # 397581)
Automatic calculations are issued by the RPAS client when a calculate is pending, and the user reorients the window by dragging and dropping or enters the “Show/Hide” dialog.	N/A (old defect # 397734)

Previous Patch – RPAS 11.2.1

Enhancements and Changes

Add new local domains and move partitions between local domains

RPAS has implemented two enhancements for global domain environments. Both of these activities are accomplished using the RPAS utility **reconfigGlobalDomainPartitions**.

Ability to add new local domains to existing global domain environments

Ability to move existing positions at the partition level (and their data) between local domains

Both operations can be executed by using their respective argument in a call to the utility (-add and –move) or by using an XML input file with the specified command. The usage of the utility for this operation is included in the section **reconfigGlobalDomainPartitions** (active link) in the "Hierarchy administration and position repartitioning" section of this document. The usage is also available in the RPAS Administration Guide.

Resolved Issues

Resolved issue	Defect
The utility for loading measures (loadmeasure) was fixed to properly handle data of different types. Previously, an entire file was being rejected for string and date measures when there was a single invalid record. Now these invalid records are ignored and the load proceeds. Additionally, integer measures were loading values with decimals. Now these values are rounded to the nearest integer when they are stored.	394766
An issue was fixed in the configuration tools measure tool where measure components were being truncated when switching the display between measure name and measure label (set in File / Preferences / Measure Tool Options).	396411
The utility for copying domains (copyDomain) no longer fails on the AIX operating system when there is not a valid global domain configuration file (globaldomainconfig.xml).	396529
The hierarchy manager utility (hierarchyMgr, new to 11.2) was fixed so that the –list argument functions properly.	396896

Resolved issue	Defect
The “install” directory contains a copy of the configuration and is used by the configuration tools when building and patching domains. This directory was moved into the domain directory so that it is included when copying a domain.	396982
An issue was fixed where the English translation of message labels that are loaded via a special translation measure (r_msglabel.ovr.english) were not being loaded into the domain during the domain build process.	397112
An issue with the utility for copying domains was fixed that occurred when copying a global domain environment with absolute paths to the same location. The domain environment was becoming unusable with this operation.	397118
The buttons in the dialog windows “Add Rule” and “Copy Selected Rules” were reordered to display in the following sequence: OK – Back – Next – Cancel.	397338
Two issues with the rule tool were fixed in the configuration tools. First, there was a problem when copying a rule into another rule group. When editing that rule, the original rule was getting updated instead of the copy. Second, the find/replace functionality was not working properly.	397339
Previously, configurers were getting an “out of memory” error message when trying to patch a domain when there was a very large task list (greater than 6 MB). The task list is no longer loaded along with the configuration when patching a domain. The system was unable to patch domains when the task list reached a certain size. The task list is not involved in the domain patch process so it has been removed from that operation and is populated when entering the tools.	397515
When patching a domain the system will no longer unregister a measure that was removed from a solution if a measure of that same name exists in another solution.	397535
Position queries in the z-axis (slice) of a workbook are now properly executed when opening a workbook.	397555
An issue was addressed in the workbook tool of the configuration tools. Edited measures were being removed from workbooks after they had been added and going to another area of the tools. Upon reentry to the workbook, the edited measure had been removed. These measures are no longer removed.	397616

Resolved issue	Defect
An issue was addressed in the configuration tools. This issue prevented the tools window from being resized if the previous session ended with the tools full screen.	n/a

RPAS 11.2 Release

Release Overview

Introduction

This document provides information about the changes and enhancements in RPAS 11.2 and covers the RPAS server, the RPAS client, and Configuration Tools. Each feature contains the following information in varying levels of detail:

- Functional or technical definition of the feature and the underlying requirements the feature meets
- How to set up and configure the feature
- How to use the feature in the RPAS client
- Notes, assumptions, and limitations about the feature

Documentation

The RPAS 11.2 GA documentation is currently only available in draft form. All of the RPAS 11.2 documents will be updated and re-released. This is tentatively scheduled for the middle of November 2005.

The following documents are included with this release and will be updated in an upcoming RPAS 11.2 patch:

- RPAS 11.2 Installation Guide – instructions for installing the RPAS server, the client, and Configuration Tools. Oracle recommends that existing customers running RPAS 11.1 upgrade to 11.2 using the version of RPAS available on ROCS. Using ROCS, the upgrade process is primarily automated and will be easier for existing customers.
- RPAS 11.2 Administration Guide – a primarily technical document that describes how to administer an RPAS environment and server, including how to use the RPAS standard administrative workbook templates and the suite of standard RPAS utilities.
- RPAS 11.2 Users Guide – provides an overview of RPAS concepts and how to use the RPAS client.
- RPAS 11.2 Configuration Guide – describes how to configure a solution using the RPAS Configuration Tools.
- RPAS 11.2 Calculation Engine Users Guide – describes how the RPAS calculation engine works, including key concepts and basics about aggregation, spreading, rules, expressions, and so on.
- RPAS 11.2 Rule Functions Reference Guide – contains the detailed syntax and usage of standard RPAS functions, procedures, and modifiers; which are mechanisms of the calculation engine that are used when writing expressions.

Impact and Implications of Upgrading

The following should be noted about upgrading to RPAS 11.2. Some of these points are important notes while others are unavoidable implications of upgrading due to an architectural change or the implementation of a new feature.

Like any RPAS patch, workbooks must be rebuilt after upgrading to reflect underlying changes to the domain. You should be reminded to commit workbooks if necessary. Administrators are prompted to confirm the deletion of workbooks when the upgrade process begins.

Limitation for Hierarchy and Dimension Names – must force same names

- RPAS has a limitation that does not allow a hierarchy to contain a dimension with the same name. For example, the PROD hierarchy cannot contain a dimension named PROD. It is imposed as a "soft" limitation in that RPAS will allow this condition but requires an administrator to force it.
- This limitation is imposed to address a specific Internationalization issue in the RPAS Translation Administration workbook template (used for modifying labels in multi-lingual environments). The names of hierarchies and dimensions are stored as positions, and the template cannot handle these identical position names.
- Although this scenario is allowed, existing configurations will be checked by the Configuration Tools to detect this condition, and you will be warned in the RPAS client if a workbook is built with this condition (including the Translation Administration template).

Scheduled Enhancements for RPAS 11.2 Patches

The following enhancements and changes are scheduled for an upcoming patch of RPAS 11.2. The release date was not determined at the time this document was written.

Enhancement to "find next alert" – RPAS will introduce a new way to find the next alert within an alert workbook. The current functionality was built to find the next alert based on a single sorted list, which assumes that you are building workbooks with the positions they wish to see. The enhancement will allow you to find the next alert based on the current "slice" that is in the z-axis of a window.

11.2 Summary

The following chapter summarizes the features and changes that are available in RPAS 11.2.2.

Functional Enhancements

- Measure formatting controlled from the tools – Ability to define measure formatting styles in the Configuration Tools and assign those styles to measures on windows. Provides styling upgrade path when patching domains or rebuilding workbook templates.
- Alert configuration support via the Configuration Tools – Ability to configure alerts in the Tools.
- Rule group simulator – For a chosen Rule Group, allows you to preview the simulated effects on other measures of editing measures, effectively providing the ability to run "what-if" scenarios by flagging measures as changed and viewing which measures will be calculated, protected, and the calculation paths, without building a domain.
- Rule pooling – Ability in the Configuration Tools to use named rules in multiple Rule Groups. In previous versions, 'identical' rules are duplicated, leading to an explosion in rules, and potential inconsistencies.
- Measure pooling – Ability to define and reuse measures so that they are not duplicated within a solution/configuration.
- Configuration task list – Ability to see a comprehensive list of configuration errors in the Tools, the nature of those errors, and some actions to resolve.
- Printing/reporting capabilities in the Tools – Ability to generate a number of reports about configurations from within the Tools and output them in electronic format. Standard reports include measures, rules, workbook templates, and translation files.
- Function helper – Enhancement to the Rule Tool that will provide information about the syntax and usage of standard RPAS functions when writing expressions.
- Single hierarchical selection widget – A widget/window with a single hierarchical tree structure where you can use hierarchy navigation to drill down the appropriate branch of the hierarchy to the leaf nodes, and select a single position. The widget can be launched from a cell in the UI, and the selected position is pasted into the cell.
- Context sensitive dynamic picklists – The contents of a picklist measure can be based on the contents of another measure, which can vary based on the current context so that different contexts within the same workbook can have different picklist contents.
- Relative calendar in wizards – extended use of position queries to define a relative time position to use in wizards to build workbooks.

Technical Enhancements

- Ability to add or remove positions without having to reshape – ability for new positions to be added to a domain and be immediately available for use without having to reshape the data structures in the domain (off-line utility name "reshapearrays").
- Position repartitioning (a.k.a. reclassification) – Automatically move positions (when required due to restructuring) between local domains in a global domain environment.
- Enhanced hierarchy load process – in conjunction with the above two features, RPAS has enhanced the hierarchy load process. This new process allows for a single call to the loadhier utility in the master domain, after which the utility handles the movement of positions below the partition level. It then updates the data structures in the domain without having to manually run the old reshaping process (reshapearrays utility).
- Automatic "reshaping" – The process of reshaping arrays (using reshapearrays utility) has been incorporated into the hierarchy load and rebuffering processes. Data in all domains of an environment is reshaped, if necessary, after loading a hierarchy that has changed or as necessary when rebuffering a domain. As such, the rehapearrays utility was deprecated.
- Add new local domains – Ability to add new local domains to an existing global domain environment.
- Move partitions between local domains – Ability to move positions at the partition level between local domains, including all relevant data for those positions.
- Centralization and parallelization of data load and export – In global domain environments, loading and exporting measure data was centralized and can be run in parallel processes. The measure load process now also includes the splitting process. This allows for a single input file to be loaded from the master domain, and RPAS will split the data and load it into the local domain where the data resides.
- New installation process for RPAS server and Tools on Windows – the RPAS 11.2 server and Configuration Tools can now be installed on Windows using an InstallShield installation process. This process also sets the required system environment variables.
- User interface for the RPAS installer – user interface that can be used to build and patch domains in a Windows environment instead of running the RPAS installer utility from a command line. The Installer UI can be launched from a start-up menu selection that is created when installing the Tools. The parameters for the UI Installer are the same as the traditional installer run from the command line.
- Remove a user-defined dimension from a domain.
- Run a subset of alerts – previously, the "alert finder" (-find argument of the alertmgr utility) evaluated all alert rules in the domain. The utility was updated to allow a subset of alerts (either a single rule or alert rule group) to be evaluated.

Minor Changes and Enhancements

The following is a list of minor changes and enhancements made to RPAS and the Tools that are available in RPAS 11.2.

RPAS Administrative Templates

- Security Administration – modified workbook template limits worksheets such that undefined values will appear as blank cells. Technically the system stores a value of 1 billion by default, but you will no longer see this value and will only see a value if one is defined (other than 1 billion).

RPAS Client

- Added "Commit ASAP" button to dialog window when closing a workbook.

Technical

- Ability to force the purge routine when loading measure data (**-forcePurge** for **loadmeasure** utility).
- RPAS server component is compiled using HP's ACC compiler on HPUX. Using the ACC compiler results in performance gains for certain off-line operations.

Tools

- Measures and measure components can now be displayed by the name or label in various locations of the Tools. This property is set in the *Tools Preferences* of the File menu, and can be set for measures and measure components.

> **Note:** If "label" is selected and no label is defined for a given measure, the system will display a blank value in the Tools.

Hierarchy Administration and Position Repartitioning

Overview

There are several new features in RPAS 11.2 that simplify the administration of global domain environments. The following new functionality is included in this release:

- Enhanced hierarchy load process that removes a number of operations that previously required additional utilities.
- Ability to add or remove positions in a domain and use new positions without having to "reshape" the data in the domain. New positions are mapped to what were previously 'dummy positions.'
- Position repartitioning – Ability to automatically move positions and data when a position has been moved from one local domain to another. This functionality is commonly referred to as "reclassification" by many RPAS customers. Positions at the partition level must be moved using an RPAS utility.
- Ability to move partitions between local domains.
- Ability to add new local domains to existing global domain environments. Previously, the number of domains had to be determined before the initial build of the environment

In RPAS 11.1, it was possible to have position names longer than 24 characters. This was accomplished with the implementation of phase 1 of Position Name Indirection (PNI) where there is effectively a mapping of internal and external position names.

RPAS 11.2 introduces phase 2 of PNI and provides the ability for a domain to contain dummy positions. These dummy positions are held as a buffer. As new positions are added to a domain in the hierarchy load process they are mapped to the dummy positions, so the domain does not need to be "reshaped," which is a time-consuming process in large domain environments.

Dummy Positions

Phase 2 of the PNI implementation provides the ability for a domain to contain dummy positions. 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. For example, imagine 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.

New positions are added to a domain through the normal hierarchy load process (RPAS utility **loadhier**). 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. Thus deletes can also 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 "re-buffering." The buffer can be updated automatically, but unpredictably, when required (based on the lower and upper bounds that are defined for a PNI dimension), but it is recommended that the rebuffering process is scheduled to ensure that batch process windows are predictable. Details of this process are discussed in later sections.

The buffer of dummy positions is defined by a high and low setting for the buffer size that is a percentage of total positions for each PNI-enabled dimension.

Position Repartitioning

Position repartitioning allows for positions to move 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:** Moving positions at the partitioning level is not automatic and is accomplished with the RPAS utility **reconfigGlobalDomainPartitions**.

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, and all corresponding data to LocalDomain2. This process is often referred to as "reclassification" by RPAS customers. RPAS is referring 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.

All of the activities and processes associated with PR will be automatically handled by RPAS and will not require manual intervention by the system administrator.

> **Note:** Position repartitioning relies upon dummy positions. Thus all dimensions at or below a level that may be repartitioned must have dummy positions enabled.

Configuration and Administration

There are a number of decisions that must be made regarding the use of dummy positions and the position repartitioning process. Dimensions will need to have dummy positions enabled for the following processes:

- Adding or deleting positions without having to reshape – customers wishing to add positions (most commonly adding new SKU's in the product hierarchy) will need to enable dummy positions for the dimensions.
- Using the automated position repartitioning processes for positions below the partition level.

All dimensions with dummy positions enabled need to have a buffer defined that represents the targeted number of dummy positions. There are high and low values for the target buffer size, which can be defined in the Tools and/or set using an RPAS utility. These buffers are used when "rebuffering" a domain.

Administrators need to determine how they wish to rebuffer the domains. Rebuffering can be scheduled and run manually using the Position Buffer Manager utility or can be handled automatically as part of the hierarchy load process.

Enabling a Dimension for Dummy Positions

PNI buffering is enabled in a domain per dimension. Dimensions are enabled using the Configuration Tools both before and after a domain is built. Two properties were added to the Hierarchy Tool, and these properties can be set for each dimension. 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, then running **loadhier**. The syntax for **loadhier** remains unchanged. 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 to other domains will be automatically processed.

Configuring and Scheduling the Rebuffering Process

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 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 time critical batch process. The automated rebuffering would then only be used as a backup and would only run in unusual circumstances.

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 local domains on a cyclical basis (such as a few each 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 depends on a number of criteria. The goal will typically be 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.

Hierarchy Load Process

From an administrative perspective, the hierarchy load process is unchanged from RPAS 11.1. All hierarchy files are loaded only into the master domain and contain all positions for a given hierarchy. If a position below the partition level was added, deleted, or reassigned; additional background processes are executed to perform the required activities. Administrators will only need to use the **loadhier** utility to perform these activities.

RPAS will know if a position in the hierarchy data files is new and will perform a number of activities to "enable" that position in the domain. If an existing position was reassigned such that it requires movement to a different local domain, RPAS will perform a number of operations to move the position and its corresponding children and data from the previous domain to the domain to which it is being reassigned.

Utility Usage

positionBufferMgr

The position buffer manager is used to set a number of parameters when using dummy positions in a domain. See the above write-up for information about using and administering dummy positions.

Usage:

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

— -rebuffer Adjust all dimensions of the provided hierarchy to have the specified number of positions based on the configuration of the dimension.

— -report Report the count status of positions for all dimensions of the provided hierarchy.

Notes:

If a hierarchy is not provided, the -rebuffer or -report argument will occur on all non-calendar hierarchies in the domain.

If a -partitionPositions is provided, the -rebuffer or -report command will work on just the domains that contain those positions.

To get this usage text, use -?, -help or –usage.

To get the version of this utility, use –version.

To set the logger verbosity level, use -loglevel with values of: all, profile, debug, information, warning, error, or none.

To disable timestamp header, use -noheader.

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:

2. dimensionMgr –d pathToDomain –dim dimensionName [COMMAND]

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	Command that displays the properties of the specified dimension.
–width *widthVal*	Command that 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 and cannot be decreased.
–bufPctMin *minVal*	Command that sets the minimum percent of unused positions.
–bufPctMax *maxVal*	Command that sets the maximum percent of unused positions.

Notes:

Multiple command arguments are allowed.

Buffer Minimum and 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 Mininum of 5 and Maximum of 20 could have between 10 and 40 extra buffer positions at any given time.

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

To get this usage text, use -?, -help or –usage.

To get the version of this utility, use –version.

To set the logger verbosity level, use -loglevel with values of: all, profile, debug, information, warning, error, or none.

To disable timestamp header, use -noheader.

reconfigGlobalDomainPartitions

This utility is not new to 11.2, but it has been changed to support the enhancements in 11.2. As such, the usage has been provided below. 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 new positions at the partition-level

Add new local domains to an existing global domain environment

Remove existing positions at the partition level

Remove local domains (this is automatic if all partition-level positions in a local domain are removed or moved)

Move existing positions at the partition level (and their data) between local domains

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 and specify the local domain to which the positions 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 **reconfigGlobalDomainPartitoins** 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
```

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.

Argument	Description
–sub *pathToSubDomain*	Specifies the path to the local domain to which positions are being added or the destination local domain for positions being moved. 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, which includes 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.

Using an input file

When using the –input argument, the file must be in a particular format and must contain the following:

"add" or "move" commands

The path to each local domain to which positions are being added or the destination for positions being moved

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.

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.
- 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.

Functional Enhancements

Single Hierarchical Selection Widget

Feature Overview and Scope

A new 'single hierarchy selection widget' was introduced to allow the selection of a single position from a potentially large number of positions for a measure. The selection is made in a window with a single hierarchical tree structure where you can use hierarchy navigation to drill down the appropriate branch of the hierarchy to the leaf nodes and select a single position. You will be able to view either position names or descriptions in this widget.

The hierarchy selection widget can be launched from a cell (by using a new control in the cell) for a normal string measure, and the cell will be filled with the (external) name of the position selected. The measure will not be manipulable other than through launching the widget to select a position.

Usage

In the Measure Tool, you declare that the measure is using the single hierarchy selection widget and the hierarchy and dimension from which the position will be selected.

You can launch the single hierarchy selection widget from a cell for the measure in the RPAS client. The widget contains the entire hierarchy that the position to be selected belongs in, down to the defined dimension, as in the scope of the workbook, which will normally be a subset of the hierarchy as a whole.

When the widget is launched, the position that is the current value of the cell it is launched from will be selected in the widget. This will show just the hierarchical rollup in the active window from which the widget was launched. If the hierarchy is not present in the active window from which the widget was launched, a default hierarchical rollup will be used. You may change the roll up being used using functionality similar to that on a spreadsheet window. When you select a position, any previously selected position in the widget is unselected. That is, there can only be a single position selected at any time.

> **Note:** The value pasted into the cell is the (external) name of the selected position.

Since the measure is an ordinary string measure, not a picklist, the UI displays the contents of the cell (for instance, the name). If it is required that the user see the label (description) of the position, a further measure will be required that holds the label of the position, calculated using the attribute function.

The following example is of such an approach where the data is stored in a measure named "labelmeas" with data at the SKU dimension. The complete definition and syntax of these functions is available in the RPAS Rule Functions Reference Guide.

```
destmeas=if(labelmeas==””,”none chosen”,attribute(“label”,
[prod].[SKU_],index([prod].[SKU_],lablemeas)))
```

Relative Time Period Position Queries in Wizards

RPAS 11.2 provides the ability to select positions from the calendar hierarchy based on a relative offset from the current date (as defined by the system or overwritten by an RPAS environment variable). This functionality is available in all 2-tree wizard processes on the wizard page where the calendar hierarchy selections are made. This offset is evaluated every time the wizard process is executed, which allows for the calendar positions to "roll" across time.

You will see this functionality in the calendar 2-tree of all wizard processes. You can specify a start period and end period that is calculated as an offset from the current period, and all periods between the 'start' and 'end' are selected. Use a negative value to specify a period before the current time and a positive value to specify a period after the current time period. The end period should be a value that corresponds to a time period after the start period; however, if it actually corresponds to a time period before the start period, the system will switch the values around, so that the section of the calendar horizon between the two values will be selected.

The start and end periods can be saved as a position query. This is done in the same manner as explicit user selections could be saved in RPAS 11.1. These position queries can be saved for an individual user or for all users (world access), and can be used (loaded) in other wizard processes, and used in autoworkbook builds.

Context Sensitive Dynamic Picklists

Overview

RPAS 11.2 introduces the ability to have picklists where the values displayed for a given cell/intersection can vary based on the context of the current selection. These enhanced "dynamic" picklists follow the same concept as the dynamic picklists that exist in RPAS 11.1, but 'look up' the picklist values dynamically each time the picklist is opened, rather than looking up the values dynamically once during workbook build, but then having static picklists inside the workbook. The feature allows for the measure in which the valid picklist values are held to be at any level in the hierarchies and the values are effectively "looked up" using normal 'non-conforming measure' handling.

Since the valid values for a picklist for a cell are referenced from a measure dynamically, it is possible, if required, for the valid values of picklists to change during the life of the workbook. This is as a result of calculations or end-user edits. The value used will always be as at the last calculate, so direct or indirect (through calculation) edits to the picklist value measure are ignored when a calculation is pending.

Usage

In the Configuration Tools, the configurer will declare whether a picklist measure will be a dynamic picklist or a context sensitive picklist and supply the name of the measure that holds the valid picklist values.

From the perspective of the end user, the use of context sensitive picklists will be identical to dynamic or static picklists.

Configurable Measure Formatting

To ease the effort of defining formatting options for measures, the Tools now provide the ability to define styles that can be assigned to measures. These styles have an inheritance structure such that each style inherits the settings of its "parent" similar to the way styles work in Microsoft Office products.

Styles are created in the Configuration Tools and are defined for a project. These styles are then assigned to measures in the Measure Tool (as an inheritable property), and can be further overwritten within a workbook window.

> **Note:** A defined style for a given measure will overwrite any existing formatting in a given workbook.

The following formatting can be defined in a style within the Tools. These settings will be applied to the label and/or each cell of the measure (as applicable):

- Prefix – String to display as a prefix to values in cells. For example, enter "$" to indicate the monetary nature of a sales measure.
- Suffix – String to display as a suffix to values in cells. For example, enter "%" to indicate the value is a percentage.
- Scale factor – Integer value to act as a scaling factor for the measure. When you enter data for a measure that contains a scaling factor, the value you enter is multiplied by the measure's scaling factor to obtain an internal value that is recognized by the server in data calculations.
- Precision – Formatting option to specify the number of digits to be displayed to the right of the decimal.
- Separator – Specifies the character to use as a separator between every three digits of a value in cells (usually either a comma or period).
- Text Font – Font in which the text of the measure's label and/or cells will be displayed.
- Text Style – Text can be displayed as regular, bold, italic, or both bold and italic.
- Text Size – Size of the text.
- Text Color – Color of text for the selected measure's label and/or cells.
- Background – Background color for the selected measure's label and/or cells.
- Alignment – Specifies left, center, or right alignment for the selected measure's cells.
- Border Style – Specifies the type of line (solid, dotted, or dashed lines) and to which edges to apply the border.
- Border Color – Specifies the color of the specified border for cells.

> **Note:** The measure formatting background color for a measure takes priority in the RPAS client over the 'read/write' background color. Thus if styles are used through the tools, the 'read/write' color will not be seen. If a specific read/write color is desired, set that as the background color of the default style, and do not override it for any other styles. On the other hand, the 'read only' background color, takes priority over the measure formatting background color, so that 'protection processing' will be visible.

- When building or patching a domain, you can decide whether or not to apply the styles that have been defined.

Configurable Alerts

RPAS 11.2 provides the ability to configure alerts in the Tools. In RPAS 11.1 and prior versions, alerts had to be configured using a combination of the measure registration utility (regmeasure) and the alert manager utility (alertmgr).

Alerts are an exception management tool for users. An alert is a measure that evaluates a business rule (returning a value of true or false). RPAS then notifies you of the "true" conditions and allows you to build workbooks to resolve the scenario that drove the alert.

Alert measures are first defined for the domain. These measures are of type Boolean, which means they have a value of true or false. These measures are configured as regular measures in the Measure Tool.

Next, a rule (expression) or rule group is associated with a given alert to define the business rule or rules that are used to evaluate the alert. These rules and/or rule groups are defined in the Rule Tool like regular rules. They are then associated with the alert measure in the Measure Tool by clicking on the *Alert Expression* measure property and selecting the defined rule or rule group to use for evaluation.

Once the registration process is complete, the alert utility is run to "find" the alerts in the domain. See the RPAS 11.2 Administration Guide for additional information. After the alert finder has been run, you can view the identified alerts in the Alert Manager in the RPAS client.

Rule Pooling

RPAS has introduced the concept of "rule pooling" in the Configuration Tools that allows a named rule to be defined only once and reused in multiple rule groups. Rules will be created in a "pool" within the Tools and managed by unique names. These rules can then be added to rule groups, and subsequent changes to a given rule are then reflected in all instances where that rule is used.

To help minimize the possibility of a duplicate rule, the Configuration Tool warns you when attempting to create a rule that has the same measures as an existing rule, regardless of the ordering or the calculation. You then have the ability to use an existing rule or create the new one.

To summarize, the following capabilities have been provided:

- Create and name a unique rule that is added to a pool.
- Add a rule from the pool to a rule group.
- Edit a rule and have changes reflected in all instances of that rule.
- Remove a rule from a rule group (does not delete from the pool).
- Delete a rule from a rule group (deletes the rule from the pool and all rule groups where it is used).
- Validation – warns you when creating a rule that has the same measures as an existing rule.
- Provide the capability of visually inspecting the contents of the pool with the ability to search by rule name starting string or rule measure content. This also includes displaying the list of rule groups that use a given rule.

Measure Pooling

RPAS has introduced the concept of "measure pooling" in the Configuration Tools that prevents the use of duplicate measures with different characteristics or measure properties. This is particularly valuable when wanting to 'share' a measure between solutions in the same domain. For example, making a financial plan target visible in an item planning solution.

Rule Group Simulator

Background

The RPAS calculation engine is powerful, complex, and configurable. The 'rule group' approach means that there are many potential calculation paths. However, during any configuration exercise, there is a significant design verification cost to ensure that the behavior is 'as would be expected' by an end user. In the absence of any rule group simulation functionality in the configuration tools, verification of the behavior of a rule group (between measures, not the calculations themselves) can only be undertaken by physically building domains and workbooks from the configuration, and then testing. With a large and complex workbook, the time spent configuring, building domains and workbooks, testing and reconfiguring, and going around the cycle again may be significant.

Overview

RPAS 11.2 introduces the Rule Group Simulator, which is integrated into the workbook tool. When launched from the workbook tool, the simulator uses all of the measures used in the rule set in the workbook, which may be more than those mentioned in the rule group.

> **Note:** The rule group simulator is not currently able to simulate the expressions that will be evaluated as the result of a rule group transition, or simulate the calculations that will follow if a rule group is evaluated in 'full' mode, such as when evaluated from the mace utility or the evaluation of the load rule group when a workbook is built.

Users of the rule group simulator are expected to understand the calculation cycle, especially with respect to measure protection processing and the process that determines which expressions will be evaluated. For more details see the calculation engine user guide.

The feature is provided in a separate window with two areas:

- a measure table
- a 'tree window' with upstream and downstream panes

The measure table shows all the measures in the scope of the simulation. The first column of the table gives the measure name and (by color coding and a 'tool tip') the status of the measure. All measures can be shown, or the list of measures can be filtered according to status (such as only those that will be calculated, those that will not be calculated, those that are read only, or those that are on a specific worksheet), or by name (those that include a specific string). A search feature that allows you to locate a specific measure in the table is also available. Measures that are currently manipulable can have their status toggled (to or from 'edited') by clicking on them, and the simulator immediately updates all statuses, calculations, and trees.

The second column of the table shows the ‘calculation’ for each measure. For those measures that would be calculated if you issued a ‘calculate’ with the current collection of edited measures, the rule name and expression that would be used to calculate the measure is shown. For measures that would not be calculated, this column displays the status of the measure. Selecting a measure in this column of the table makes it the ‘active measure’ that is used as the focus of the tree windows.

The tree window shows (in separate panes) the upstream and downstream measure relationships for the measure with focus. A new measure with focus may be selected by clicking on the table’s calculation column for the new measure or by clicking on the measure in one of the tree panes. If the measure with focus is calculated, the upstream pane shows the measures that affect its calculation and all measures that those measures are dependent upon (calculated from). The downstream pane shows measures that are dependent upon (calculated from) the measure with focus (if there are any).

Usage

The simulator can be invoked from the Workbook Tool by first selecting a Rule Group from the Workbook Properties table. Right click the mouse to bring up the context menu and select “Simulate Rule Group” to bring up the Simulation dialog.

The status of each measure, given the current set of ‘changes,’ is shown in the simulator measure table through color coding, and tool tips. The color scheme is as follows:

Color	Meaning
Yellow	Edited. If it is a recalc measure, it will be calculated by indirect spreading of another measure through a mapping rule and recalculation at aggregated levels. If the measure has another aggregation type, it will be calculated by spreading and aggregation.
Pale Grey	Editable. However, although the measure is not forced, and thus is still editable, it will be calculated through the calculation engine having to select an expression in an affected rule.
White	Editable. Will not be calculated, and thus will not change at all.
Pale Blue	Protected by protection processing. Although the measure is protected (usually this will be because it is the measure on the LHS of the only expression in a rule), it is not ‘forced’ because none of the RHS measures are changed, so it does not need to be calculated, and it will not change at all.
Mid-Blue	Protected by protection processing. Is ‘forced’ and so will be calculated.

Color	Meaning
Dark Blue	Read Only. The measure is set as being read only in the measure properties, and will not change at all. A measure that is read only, but is going to be calculated will be shown as mid-blue as that status takes priority over read only.

Note: The status of a measure encapsulates two concepts that are not as closely linked as may appear at first sight: whether or not the measure is manipulable (shades of blue = no, white/grey/yellow = yes) and whether or not the measure will be calculated. It is perfectly possible for a measure to currently be manipulable, but, it would be calculated if a calculate were issued. Similarly, it is possible for a measure to be protected by protection processing that would not be calculated if a calculate were issued.

Measures that are currently manipulable and not already changed (white and grey) may have their status changed to edited (yellow) by clicking on their names in the first column of the measure table. All other measure statuses, calculations and the tree panes are updated. Similarly, a measure that is currently edited (yellow) may have that edit 'removed' by clicking on the measure in the measure table. The measure reverts to another status (white or grey), and all other measure statuses, calculations, and the tree panes will be updated. In addition, a reset button is provided so that you can start a new simulation from a clean slate with no measures set as changed.

The simulator provides a mechanism for filtering the measures. Using a drop down box, the user may opt to show:

- All measures
- Only those that will be calculated (arranged according to order of calculation)
- Only those that will not be calculated
- Read only measures

Furthermore, there is a "Contains String" filter that can be used in conjunction with a search text field. Selecting this filter and typing a string to filter by updates the table to display only measures that contain the filter string. If you select the "All" filter and types, a search string on the text field, the first measure found containing the search string becomes the selected measure. A "By Worksheet" filter, in conjunction with a drop down box, allows you to view only the measures in a given worksheet.

The tree panes show measure dependencies through expressions that will be evaluated based on a single measure that has focus. The measure with focus can be selected by clicking in the calculation column of the measure table for the measure, or by clicking on the measure name in a tree pane. When the focus changes, the tree panes are redrawn as appropriate to the new focus, and the measure table scrolls so the measure with focus is shown.

If the measure with focus is calculated, the ‘upstream’ pane shows the measures that it is dependent on (calculated from, directly and indirectly). This is shown by a ‘tree’ structure with the measures used to calculate an individual measure showing as ‘children’ of that measure. If such children are also calculated, they will be shown with the measures they depend on, and so on, recursively. Thus the ‘upstream’ pane fully explodes all the measure relationships that affect the measure with focus. If the measure with focus is not calculated, the ‘upstream’ pane will display that measure as a leaf by itself, which indicates that it is not dependent on any other measures.

The downstream pane shows measures that are dependent upon (calculated from) the measure with focus (if there are any). Measure relationships (expressions) are again shown in a parent-child ‘tree’ structure. There may be several measures that are immediately dependent on the measure with focus (that is, the measure with focus may be on the RHS of several expressions that will be evaluated), so these relationships cannot all be shown at the same time in a simple tree structure. A ‘VCR control’ allows the selection from several immediate ‘dependent measures’ by paging through their trees in the downstream pane. Thus the ‘downstream’ pane does not (indeed cannot) fully explode all of the measure relationships that are affected by the measure with focus. The full explosion of detail can only be viewed by ‘walking down the tree,’ changing the focus to the desired ‘dependent measure,’ and selecting the sub branches through the VCR control.

The simulator displays multiple result expressions by indicating this beside the rule’s name in the second table column. The second table column also displays the entire expression showing all of the left hand measures that comprise the multiple results. If a measure that has focus is one of the multiple result measures, it will be shown in the tree displays as MeasA [+MeasB][+MeasC] where MeasA is the measure in focus and MeasB and MeasC are the other multiple result measures. Measures that are calculated in a ‘cycle,’ which typically includes BOP and EOP inventory values, are indicated with an ‘*’ at the end of the measure name. The tree displays will show only one instance of a cyclic relationship between two measures.

Configuration Task List

RPAS 11.2 has introduced the concept of a “task list” that provides a single location for all configuration issues encountered when configuring a solution. When configuring a solution, it is common to run into situations where a given configuration setting is invalid. This can be for a wide range of causes and can often be difficult to troubleshoot. The Task List is intended to ease this troubleshooting process by providing a location where all outstanding configuration-related issues can be seen.

The task list is a partition in the bottom half of the main window of the Configuration Tools and appears by default when opening the Tools. The window can effectively be minimized by clicking on the “Task List” title bar of the window, or shrunk or enlarged by dragging the divider between the top and bottom windows.

The Task List can identify two types of configuration issues: local errors and remote warnings. A local error is one that is caused when actually making an edit, such as a syntax error or a duplicate name. A remote warning is the result of an edit where the edit itself is not necessarily invalid, but where the edit may cause another configuration setting, not found in the currently visible window, to become invalid.

The configuration issues checked by the Task List are based on validation within the Tools and validation provided by the interface between the RPAS server and the Tools.

The number of possible configuration errors is large and is therefore not explicitly listed in this document.

Reporting Configuration Information

The Configuration Tools provide the ability to generate reports about a configuration in RPAS 11.2. These reports contain a variety of information about a configuration and are output to a text file that can be read into other applications for documentation or inspection purposes.

The Report Generator is run from the *Utilities* menu of the Configuration Tools. Once inside the Report Generator, you must select from the following reports. Note that additional reports will be added in future releases.

Measures

Data interface

Measure descriptions for translation purposes

Measure labels for translation purposes

Rules

Rule group labels for translation purposes

Workbooks

Workbook group labels for translation purposes

Workbook labels for translation purposes

RPAS messages for translation purposes

Miscellaneous Changes to Existing Tools

Rules

Previously, adding a Rule to a Rule Group resulted in the Rule being added at the end of the list in the GUI. The Rule is now inserted at the currently selected location.

Up/Down buttons are now available on the toolbar to make moving Rules easier. Previously, if you wanted to move a rule several places up or down, you would have to click once on a small arrow beside the Rule name, and then mouse over to the new location of the arrow for the next click. The tool bar buttons still require multiple clicks, but there is no longer a need to move the mouse before each click.

Autoscrolling is now available whenever moving Rules in a small window.

Measures

If the database field is populated for a measure, the Tools will prepend these values with "data\" in the configuration file (cfg).

Workbooks

When creating position queries in the tool, you can now bring up all the Boolean measures with more than one dimension in their intersection. There is no longer a need for the measure to be in the Viewable Profile before it can be selectable as a position query measure. Validation was added to prevent selection of a measure whose intersection dimensions did not comprise a subset of the sheet intersection dimensions. Validation was also added to flag a previously defined position query measure whenever a change in sheet intersection makes the measure invalid for the given sheet.

Technical Enhancements

Centralized and Parallelized Measure Loading and Export

The RPAS utilities **loadmeasure** and **exportdata** have been updated so that they only need to be run in the master domain of a Global Domain environment.

The **loadmeasure** utility has also been enhanced to incorporate the process of splitting input files so that the data is loaded into the appropriate sub-domain when the measure loading utility is called. Previously, the splitting process was completed using the **splitmeasuredata** utility, which was deprecated.

Both utilities have a new argument to run multiple processes in parallel across local domains. The –processes argument should be included in the call to the utility and determines the number of additional child processes that are run in addition to the main process (in the master domain). The number of processes should be less than or equal to the number of processors on the server.

Loadmeasure takes an input file and splits it into multiple files based on the mapping of positions to their respective domains. Data for positions that belong to the master domain (for instance, for measures with a higher base intersection than the partition level) will be loaded in the master domain.

> **Note:** It is possible to load a single input file or to split the data into multiple files before running **loadmeasure**.

exportdata allows some arguments to be specified in a file in the domain instead of the command-line via the '-params' argument (and entering the arguments individually). If run in a Global Domain environment, the array specified by this argument must be in the master domain, not the local domains.

> **Note:** The –array argument is not supported in Global Domain environments.

New Utilities

The following list of utilities new to 11.2 is documented in a different section of this document:

dimensionMgr – included in Hierarchy Administration and Position Repartitioning

Managing Settings for Hierarchies – hierarchyMgr

The hierarchy manager utility can be used to set the following properties:

Enabling/disabling position-level security for a dimension in a hierarchy (other than calendar)

Allowing the use the same name for hierarchies and dimensions

- Syntax: **hierarchyMgr** –d domainPath –h hierarchyName [COMMAND]

Enhancements to Existing Utilities

Ability to define a named rule using calc engine utility – mace

RPAS added the ability to add a named rule to the domain using the calculation engine utility **mace**.

Syntax: **mace** -d domainPath -newRule {-ruleName ruleName}

Filtering out values when exporting data – exportData

An enhancement was made to the RPAS utility **exportData** (new parameter **–precision**) that allows you to ignore (not export) values that differ from the NA value less than a specified value.

> **Note:** This feature was also released in a patch to RPAS 11.1.

The purpose of this enhancement is for the export utility not to export values that vary only slightly from the NA value, most commonly very small values close to zero. For example, consider a numeric measure with a NA value of zero. If the precision value is specified as 0.01, any number smaller than that is ignored, so 0.0034 would be ignored while 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.

Evaluating (finding) a subset of alerts – alertmgr

The alert manager utility (alertmgr) was enhanced to allow the evaluation of a subset of alerts. The utility can be used to evaluate one or more individual alerts, or one or more alert categories (assigned when configuring the alerts).

Syntax: **alertmgr** –d domainPath –findAlerts {-alerts “a1 a2 …” | -categories “cat1 cat2 …}

Force the purging of data when loading measures – loadmeasure

A new parameter was introduced (**–forcePurge**) for the **loadmeasure** utility that forces data to be purged when running the utility. Using this flag when running **loadmeasure** will ensure that the purge routine is run even when no new data is loaded. This allows an administrator to "touch" an empty .ovr file and then run **loadmeasure** to just purge the measure without loading anything.

Deprecated Utilities

The following utilities are no longer available in RPAS 11.2. They are either no longer required or have been replaced by other utilities.

reshapearrays – The reshaping process was incorporated into the hierarchy load process. A new utility **defrag** is now provided to be used when it is necessary to defragment or clean up a domain.

inithier – The inithier utility was used for enabling or disabling position-level security and for extending the width of positions for a dimension (to allow for position names longer than 24 characters).

- Position-level security is now set using the **hierarchyMgr** utility
- The width of position names can be extended using the **dimensionMgr** utility

splitmeasuredata – In RPAS 11.1, the process of splitting the data of an input file for loading into local domains was completed using the splitmeasuredata utility. This process was added to the **loadmeasure** utility so that a single input file can be used to load data into the master domain of a global domain environment.

VisualStudio.net 2003

The server component of RPAS 11.2 is compiled using Visual Studio .net 2003 (version 7). Any developer using Visual Studio to develop extensions to RPAS will need to upgrade to this new version to develop and compile the extensions. This change is only relevant for RPAS 11.2.x.

Miscellaneous Technical Information

RpasInstall Logging

A verbosity setting was added to rpasInstall (-verbose). Currently, this setting only relates to information output by createRpasDomain. The default is that the parameter is not set, which means only errors and warnings will be output. If -verbose is provided, the utility will output information in addition to warnings and errors.

RPAS Lock Timeout Variable

RPAS introduced a new environmental variable named “RPAS_LOCK_TIMEOUT.” This variable was added to allow an administrator to set how long they want to wait before timing out when there are lock contention issues. This most commonly occurs when two users attempt to simultaneously commit/save the same data.

This environment variable is set on the back end to the number of milliseconds to wait for a file lock before returning a lock contention error. The default value is 60000 milliseconds (one minute). 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
DomainDaemon -port 55123  -start -debug &
```

New Application Programming Interfaces

Disabling wizard page controls

RPAS now provides the ability to programmatically disable wizard page controls in custom wizards.

The class WizardPage has a new method enable(String, bool) that allows an application developer to enable or disable a control through a C++ API. The first parameter is the control name, and the second parameter is 'true' for enabled and 'false' for disabled.

Date control for wizard pages

Wizard pages now support date control. Below is a list of types with definitions and examples.

Date picker control. The parameters are identical to that of addEdit(...)

```
bool addDatePicker(
      const String& controlName,
      const String& styles,
      long xStart, long yStart,
      long width, long height);
```

Use the following to specify the initial value of the control. If this method is not called, the wizard page framework will set the initial value to the previously selected value. If there is not a previously selected value, the control defaults to the current date.

```
bool setDate(
        const String& controlName,
        const SimpleDate& date);
```

The following method allows the date picker to be ranged. The dates specified are inclusive with respect to the range.

```
bool setDateRange(
        const String& controlName,
        const SimpleDate& startDate,   // inclusive
        const SimpleDate& endDate);    // inclusive
```

The method below allows a wizard page to get the value "picked" by the GUI.

```
bool getDate(
        const String& controlName,
        SimpleDate& date) const;
```

Hiding positions in a worksheet

A new API was added that allows an application developer to hide positions in a worksheet. The following information describes this new feature and how it can be utilized.

Note: This feature was original introduced in RPAS 11.1.6.

A new interface “addWindowPositions” was added into WorkbookTemplate class. This function can be used together with preexisting function “setVisibleWindowPositions” to selectively hide window positions for newly built workbooks.

Application teams may use this function in their template to specify all window positions first and then to set the visible position by invoking “setVisibleWindowPositions.” All unmentioned window positions will be hidden.

```
void addWindowPositions(const String& windowName,
                        const String& dimName,
                        const String::SetT& positionNames);
void setVisibleWindowPositions(const String& windowName,
                        const String& dimName,
                        const String::SetT& positionNames);
```

Below are some answers to some anticipated questions about this feature:

- What happens if the positions already exist when I call addWindowPositions?
  - Existing positions will not be added again, and retain their previous state. Basically, this function creates the window position array (arrayName = winName + dimName) if it does not exist and add positions if they are not already in.
- When I call addWindowPositions, are the positions hidden or visible by default?
  - The positions added by “addWindowPositions” are visible by default. However, calling “setVisibleWindowPositions” will set mentioned positions visible and all other unmentioned existing positions hidden.
- If you have already saved some styles, what is the effect of calling these two functions? Does it just affect visibility of the positions, or does it override some other kind of format too? (for instance, width or something else)
  - It just affects visibility of the positions.

Frequently Asked Questions

Below is a list of several key questions that have been asked about RPAS 11.2:

Q: Is there a way to defer the reshaping process until after loading all hierarchy files when not using dummy positions (reshaping is handled by the rebuffering process when using dummy positions)? We want to avoid reshaping the domain until all hierarchy files have been loaded.

A: Yes, there is an argument for loadhier named "-loadall" that will load all hierarchy input files (with a .dat file extension) that are in the "input" directory and reshape the data after processing all the files.

Q: When configuring styles in Tools, what happens when overwriting certain settings in the client? Will they be overwritten? If so can we change?

A: Measure formatting is only overwritten for the measures that are changed. So existing formatting for measures will remain as it was until the formatting is set in the Tools and the domain is patched. All other formatting will remain the same as only measure formats are configurable in the Tools right now.

Q: What is the performance impact of using dummy positions?

A: The use of dummy positions does have an impact on performance, and it appears to be related to the number of dimensions with dummy positions more so than the number of dummy positions themselves. There is no generic advice available as the exact impact depends on so many different factors of an environment.

Q: What is a good strategy for determining the number of dummy positions that should be in the environment?

A: The simplest answer is "enough positions to last between rebuffering processes." Rebuffering has effectively replaced the reshaping process, so rebuffering is going to require a similar amount of time. Administrators will want to determine how long their rebuffer process is going to take, and figure out how often that means they can afford to rebuffer the environment (one or more domains nightly, weekly, monthly, and so on). One can then determine the time in between rebuffering processes for each domain. Using that information, they should determine the maximum amount of position growth that would be required, and size the buffer accordingly.

Q: Are dummy positions available in workbooks so that a workbook can be refreshed without having to be rebuilt?

A: Dummy positions are only available in the domain and are only for new positions added to the domain. Workbooks still must be rebuilt to reflect changes to hierarchies.

Q: I want to use the single hierarchial selection widget for picklists but want to display the position label to the user rather than the position name. How can this be done?

A: Below is an example of such an approach where the name of the position is stored in a measure named “labelmeas” (that is the measure that uses the Single Select Widget) with data at the SKU dimension. The complete definition and syntax of these functions is available in the RPAS Rule Functions Reference Guide. Note that this only updates after a calc.

```
destmeas=if(labelmeas==””,”none chosen”,attribute(“label”,
[prod].[SKU_],index([prod].[SKU_],lablemeas)))
```