

Oracle® Retail Demand Forecasting

User Guide for the RPAS Fusion Client

Release 13.4.1

May 2013

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

Primary Author: Melissa Artley

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

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

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

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

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

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Licensing Note: This media pack includes a Restricted Use license for Oracle Retail Predictive Application Server (RPAS) - Enterprise Engine to support Oracle® Retail Demand Forecasting only.

Value-Added Reseller (VAR) Language

Oracle Retail VAR Applications

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

(i) the **MicroStrategy** Components developed and licensed by MicroStrategy Services Corporation (MicroStrategy) of McLean, Virginia to Oracle and imbedded in the MicroStrategy for Oracle Retail Data Warehouse and MicroStrategy for Oracle Retail Planning & Optimization applications.

(ii) the **Wavelink** component developed and licensed by Wavelink Corporation (Wavelink) of Kirkland, Washington, to Oracle and imbedded in Oracle Retail Mobile Store Inventory Management.

(iii) the software component known as **Access Via**TM licensed by Access Via of Seattle, Washington, and imbedded in Oracle Retail Signs and Oracle Retail Labels and Tags.

(iv) the software component known as **Adobe Flex**TM licensed by Adobe Systems Incorporated of San Jose, California, and imbedded in Oracle Retail Promotion Planning & Optimization application.

You acknowledge and confirm that Oracle grants you use of only the object code of the VAR Applications. Oracle will not deliver source code to the VAR Applications to you. Notwithstanding any other term or condition of the agreement and this ordering document, you shall not cause or permit alteration of any VAR Applications. For purposes of this section, "alteration" refers to all alterations, translations, upgrades, enhancements, customizations or modifications of all or any portion of the VAR Applications including all reconfigurations, reassembly or reverse assembly, re-engineering or reverse engineering and recompilations or reverse compilations of the VAR Applications or any derivatives of the VAR Applications. You acknowledge that it shall be a breach of the agreement to utilize the relationship, and/or confidential information of the VAR Applications for purposes of competitive discovery.

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

Contents

Send Us Your Comments	xxiii
Preface	xxv
Audience	xxv
Documentation Accessibility	xxv
Related Documents	xxv
Supplemental Documentation	xxvi
Customer Support	xxvi
Review Patch Documentation	xxvi
Oracle Retail Documentation on the Oracle Technology Network	xxvi
Conventions	xxvii
1 Introduction	
Forecasting Challenges and RDF Solutions	1-1
Selecting the Best Forecasting Method	1-2
Overcoming Data Sparsity Through Source Level Forecasting	1-2
Forecasting Demand for New Products and Locations	1-3
Managing Forecasting Results Through Automated Exception Reporting	1-3
Incorporating the Effects of Promotions and Other Event-Based Challenges on Demand	1-4
Oracle Retail Demand Forecasting Architecture	1-4
The Oracle Retail Predictive Application Server and RDF	1-4
Global Domain versus Simple Domain Environment	1-5
Oracle Retail Demand Forecasting Workbook Groups	1-6
Forecast	1-6
Promote (Promotional Forecasting)	1-7
Cross Promo Effect Module (CPEM)	1-7
Curve	1-8
Grade	1-9
RDF Solution and Business Process Overview	1-9
RDF and the Oracle Retail Enterprise	1-10
RDF Primary Workflow	1-10
2 Navigation	
Before you Begin	2-1
Logging into RDF	2-1

Taskflow Pane	2-3
Access-based Visibility	2-6
Switching Between Multiple Tasks	2-6
Opening a Workbook	2-6
Deleting a Workbook	2-7
Renaming a Workbook	2-8
Understanding the Workbook Wizard Window	2-9
Accessing Online Help	2-11
Contents	2-12
Figures and Tables	2-12
Index.....	2-13
Search.....	2-13
Logging Out of RDF	2-14

3 Oracle Retail Demand Forecasting Methods

Forecasting Techniques Used in RDF	3-1
Exponential Smoothing	3-1
Regression Analysis.....	3-1
Bayesian Analysis.....	3-2
Prediction Intervals.....	3-2
Automatic Method Selection	3-2
Source Level Forecasting.....	3-2
Promotional Forecasting	3-2
Time Series (Statistical) Forecasting Methods	3-2
Why Use Statistical Forecasting?	3-3
Average.....	3-4
Bayesian Information Criterion (BIC)	3-4
AutoES Flowchart	3-5
Automatic Forecast Level Selection (AutoSource).....	3-6
The Forecast Level Selection Process	3-7
Determining the Best Source Level Forecast.....	3-8
Status and Scheduling.....	3-8
Using the System-Selected Forecast Level	3-8
Exponential Smoothing (ES) Forecasting Methods.....	3-9
Simple Exponential Smoothing.....	3-9
Croston's Method	3-10
Simple/Intermittent Exponential Smoothing.....	3-11
Holt Exponential Smoothing.....	3-11
Multiplicative Winters Exponential Smoothing	3-11
Oracle Winters.....	3-12
Oracle Winters Decomposition.....	3-13
Winters Standard	3-13
Winters Responsive	3-14
Additive Winters Exponential Smoothing	3-14
Seasonal Exponential Smoothing (SeasonalES)	3-15
Seasonal Regression	3-15
Profile-Based Forecasting	3-16

Forecast Method	3-16
Profile-Based Method and New Items	3-17
Bayesian Forecasting	3-17
Sales Plans versus Historic Data	3-18
Bayesian Algorithm	3-18
LoadPlan Forecasting Method	3-18
Copy Forecasting Method	3-19
Causal (Promotional) Forecasting Method	3-19
The Causal Forecasting Algorithm	3-20
Causal Forecasting Process	3-20
Flowchart	3-21
Causal Forecasting with External Baseline	3-23
Causal Forecasting Using External Baseline Process	3-23
Causal Forecasting for Short-lifecycle Items	3-23
Causal Forecasting at the Daily Level	3-24
Example	3-24
Final Considerations About Causal Forecasting	3-25
4 Forecast Product/Location Cloning Administration Task	
Product/Location Cloning Administration Workbook	4-1
Adjustment Parameters View	4-3
Measures: Adjustment Parameters View	4-3
Location Cloning View	4-4
Measures: Location Cloning View	4-5
Product Cloning View	4-6
Measures: Product Cloning View	4-6
Clone Mask View	4-7
Measures: Clone Mask View	4-7
5 Forecast Administration Task	
Forecast Administration Task Components	5-1
Basic versus Advanced	5-1
Final versus Source Level Forecasts	5-1
Forecasting Methods Available in RDF	5-2
Forecast Administration Workbook	5-5
Forecast Administration Workbook: Basic Settings	5-6
Final Level View - Basic Settings	5-6
Measures: Final Level View - Basic Settings	5-7
Final and Source Level Parameters View - Basic Settings	5-8
Measures: Final and Source Level Parameters View - Basic Settings	5-8
Forecast Administration Workbook: Advanced Settings	5-10
Causal Parameters View - Advanced Settings	5-10
Measures: Causal Parameters View - Advanced Settings	5-11
Final Level Parameters View - Advanced Settings	5-14
Measures: Final Level Parameters View - Advanced Settings	5-15
Final and Source Level Parameters View - Advanced Settings	5-19

Measures: Source Level View - Advanced Settings.....	5-19
Regular Price Parameters View- Advanced Settings.....	5-23
Measures: Regular Price Parameters View- Advanced Settings.....	5-23
Promo Affinity Parameters View.....	5-25
Measures: Promo Affinity Parameters View	5-25
6 Forecast Approval Task	
Forecast Approval Workbook	6-2
Advanced Final Level View.....	6-10
Measures: Advanced Final Level View.....	6-11
Final Level View	6-12
Measures: Final Level View	6-12
Approval View	6-15
Measures: Approval View	6-15
Valid Forecast Run View	6-16
Measures: Valid Forecast Run View.....	6-16
Source Level View	6-16
Measures: Source Level View.....	6-17
Final or Source System Parameters View	6-17
Measures: Final or Source System Parameters View	6-17
7 Forecast Delete Task	
Forecast Delete Wizard.....	7-1
8 Forecast Scorecard Task	
Forecast Scorecard Workbook	8-1
Error Measure View	8-8
Measures: Error Measure View.....	8-9
Actuals versus Forecasts View	8-10
Measures: Actuals versus Forecasts View	8-10
9 Interactive Forecasting Task	
Improve Performance for the Interactive Forecasting Workbook	9-1
Interactive Forecasting Workbook	9-1
Forecasting Parameter View	9-7
Measures: Forecasting Parameter View.....	9-7
Additional Information About the Interactive Forecasting View Parameters.....	9-8
Interactive Forecasting View	9-8
Measures: Interactive Forecasting View	9-8
10 Forecast Like-item, Sister-Store Task	
Forecast Like-Item, Sister-Store Workbook	10-1
Like-Item View	10-3
Measures: Like-Item View	10-3
Sister-Store View	10-4

Measures: Sister-Store View	10-4
Advanced Parameter View	10-4
Measures: Advanced Parameter View	10-5
Required Steps for Forecasting Using Each of the Like SKU/Sister-Store Methods	10-7
Procedures.....	10-7
Seasonality/SKU.....	10-7
Seasonality/STR.....	10-8
Seasonality/SKU_STR.....	10-8
Lifecycle/SKU	10-8
Lifecycle/STR.....	10-9
Lifecycle/SKU_STR.....	10-9
11 Forecast Maintenance Task	
Forecast Maintenance Workbook	11-1
Forecast Maintenance Workbook: Basic Settings	11-3
Final and Source Level Views - Basic Settings.....	11-3
Measures: Final and Source Level Views - Basic Settings.....	11-4
Forecast Maintenance Workbook: Advanced Settings	11-6
Advanced Final Parameter View	11-6
Measures: Advanced Final Parameter View	11-7
12 Run Batch Forecast Task	
Running a Batch Forecast Manually	12-1
Before You Begin	12-1
Procedure	12-1
Run Batch Forecast Workbook	12-2
13 Promote (Promotional Forecasting)	
About Promote	13-1
What Is Promotional Forecasting?.....	13-1
Comparison Between Promotional and Statistical Forecasting	13-2
Developing Promotional Forecast Methods.....	13-2
Promotional Forecasting Approach	13-3
Promotional Forecasting Terminology and Workflow.....	13-3
Examples of Promotion Events.....	13-4
Promote Tasks and Wizards.....	13-4
Promotion Effectiveness Task	13-5
Promotion Effectiveness Workbook.....	13-6
Promotion Model Type View	13-9
Measures: Promotion Model Type View.....	13-9
Promotions View	13-9
Measures: Promotions View.....	13-10
Aggregated Promotion Lift Override View	13-10
Measures: Aggregated Promotion Lift Override View	13-10
Promotional Forecast View.....	13-11
Measures: Promotional Forecast View	13-11

Promotion Parameters View.....	13-12
Measures: Promotion Parameters View	13-13
Promotion Maintenance Task	13-14
Promotion Maintenance Workbook	13-15
Promotion Parameters View.....	13-17
Measures: Promotion Parameters View	13-18
Promotion Management Task	13-19
Promotion Management Workbook.....	13-19
Promo Variable Model Type View	13-21
Measure: Promo Variable Model Type View.....	13-21
Promotion Enable View	13-21
Measures: Promotion Enable View	13-21
Accept Negative Lift View.....	13-22
Measures: Accept Negative Lift View.....	13-22
Promotion Planner Task	13-22
Promotion Planner Workbook	13-23
Promotions View	13-25
Measures: Promotions View.....	13-25
Procedures in Promotional Forecasting	13-26
Setting Up the System to Run a Promotional Forecast.....	13-26
Viewing a Forecast That Includes Promotion Effects	13-27
Viewing and Editing Promotion System-Calculated Effects	13-27
Promotion Simulation (What-if?) and Analysis	13-28

14 Cross Promotional Effects Module (CPEM)

About Promotional Effects	14-1
CPEM Workbooks and Wizards	14-2
Effect Estimation Administration Task	14-2
Effect Estimation Administration Workbook	14-2
Default Parameters View	14-4
Measures: Default Parameters View	14-4
Cannibalization Parameter Override View	14-5
Measures: Cannibalization Parameter Override View	14-5
Halo Parameter Override View	14-6
Measures: Halo Parameter Override View	14-6
Halo Products Mapping View.....	14-6
Measures: Halo Products Mapping View	14-7
Effect Estimation Review and Approval Task	14-7
Effect Estimation Review and Approval Workbook.....	14-7
Cannibalization Effects View	14-10
Measures: Cannibalization Effects View	14-10
Cannibalization Effects Drive View	14-10
Measures: Cannibalization Effects Drive View	14-11
Cannibalization History Data View	14-11
Measures: Cannibalization History Data View	14-11
Halo Effects View	14-12
Measures: Halo Effects View.....	14-12

Halo Effects Drive View	14-12
Measures: Halo Effects Drive View	14-13
Halo History Data View	14-13
Measures: Halo History Data View	14-14
Cannibalization Statistic Report View	14-15
Measures: Cannibalization Statistic Report View	14-15
Halo Statistic Report View	14-15
Measures: Halo Statistic Report View	14-16
15 Curve	
Dynamic Profiles	15-1
Curve Tasks and Wizards	15-2
Profile Administration Task	15-2
Profile Administration Workbook	15-2
Profile Parameter View	15-3
Measures: Profile Parameter View	15-3
Profile and Source Level Intersection View	15-6
Measures: Profile and Source Level Intersection View	15-7
Profile Approval Task	15-7
Profile Approval Workbook	15-8
Final Profile View	15-10
Measures: Final Profile View	15-10
Source Profile View	15-11
Measures: Source Profile View	15-11
Profile Approval View	15-11
Measures: Profile Approval View	15-12
Profile Maintenance Task	15-12
Profile Maintenance Workbook	15-12
Profile Maintenance: Final Approval and Sourcing View	15-15
Measures: Final Approval and Sourcing View	15-16
Profile Maintenance: Final Training Window View	15-16
Measures: Final Training Window	15-16
Run Batch Profile Task	15-17
Run Batch Profile Workbook	15-17
16 Grade	
Cluster Methods in Grade	16-1
Breakpoint	16-2
Index to Average	16-2
Batch Neural Gas Algorithm (BaNG)	16-2
BaNG versus Breakpoint	16-2
Grade in a Global Domain Environment	16-3
Grade Tasks and Wizards	16-3
Breakpoints Administration Task	16-3
Breakpoints Administration Workbook	16-3
Breakpoint Administration View	16-4

Delete Cluster Run Task	16-5
Delete Cluster Run Wizard	16-5
Cluster Review Task	16-5
Cluster Review Workbook.....	16-5
Cluster Input Summary View	16-6
Measures: Cluster Input Summary View	16-7
Generate Breakpoint Grades Task	16-8
Generate Breakpoint Grades Wizard	16-8
Generate Clusters Task	16-12
Generate Clusters Wizard	16-12

A Appendix: Preprocessing

Preprocessing Methods	A-2
Standard Median	A-2
Retail Median.....	A-3
Standard Exponential Smoothing.....	A-3
Lost Sales Standard Exponential Smoothing	A-5
Override.....	A-6
Increment.....	A-6
Forecast Sigma	A-7
Forecast Sigma Event.....	A-7
Clear	A-8
No Filtering	A-8
Preprocessing for Stock-outs	A-8
Preprocessing for Promotional Forecasting	A-8

B Appendix: RDF Workbooks (Fusion Client)

Forecast Product/Location Cloning Administration Task	B-1
Product/Location Cloning Administration Workbook	B-1
Forecast Administration Task	B-2
Forecast Administration Workbook.....	B-2
Forecast Approval Task	B-5
Forecast Approval Workbook	B-5
Forecast Delete Task	B-6
Wizard	B-6
Forecast Scorecard Task	B-6
Forecast Scorecard Workbook.....	B-6
Interactive Forecasting Task	B-7
Interactive Forecasting Workbook.....	B-7
Forecast Like-Item, Sister-Store Task	B-8
Forecast Like-Item, Sister-Store Workbook.....	B-8
Forecast Maintenance Task	B-8
Forecast Maintenance Workbook	B-8
Run Batch Forecast Task	B-9
Wizard	B-9
Promote (Promotional Forecasting)	B-9
Promotion Effectiveness Task	B-9

Promotion Effectiveness Workbook.....	B-10
Promotion Maintenance Task.....	B-10
Promotion Maintenance Workbook.....	B-10
Promotion Management Task.....	B-11
Promotion Management Workbook	B-11
Promotion Planner Task.....	B-11
Promotion Planner Workbook.....	B-11
Cross Promotional Effects Module (CPEM) Task	B-12
Effect Estimation Administration Workbook	B-12
Effect Estimation Review and Approval Workbook	B-12

C Appendix: Curve Workbooks (Fusion Client)

Profile Administration Task.....	C-1
Profile Administration Workbook.....	C-1
Profile Approval Task.....	C-2
Profile Approval Workbook.....	C-2
Profile Maintenance Task	C-2
Profile Maintenance Workbook	C-2
Run Batch Profile Task.....	C-2

D Appendix: Grade Workbooks (Fusion Client)

Breakpoints Administration Task.....	D-1
Delete Cluster Run Task	D-1
Cluster Review Task	D-1
Breakpoint Grades Task.....	D-2
Generate Clusters Task.....	D-2

Glossary

List of Tables

2-1	Description of Icons in the Taskflow Pane	2-5
2-2	Workbook Wizard Window User Interface Components.....	2-10
5-1	Fallback Methods	5-20
5-2	Winters Forecasting Approaches.....	5-22
5-3	Comparison of Winters Forecasting Models	5-22
6-1	Views for the Forecast Approval Workbook	6-1
10-1	Substitute Methods for Forecasting	10-7
13-1	Promote Tasks and Wizards.....	13-5
13-2	Promo Effect Types.....	13-18
B-1	Product/Location Cloning Administration Workbook	B-2
B-2	Forecast Administration Workbook.....	B-3
B-3	Forecast Approval Workbook.....	B-5
B-4	Forecast Scorecard Workbook.....	B-7
B-5	Interactive Forecasting Workbook	B-7
B-6	Forecast Like-Item, Sister-Store Workbook.....	B-8
B-7	Forecast Maintenance Workbook	B-9
B-8	Promotion Effectiveness Workbook.....	B-10
B-9	Promotion Maintenance Workbook.....	B-11
B-10	Promotion Management Workbook	B-11
B-11	Promotion Planner Workbook	B-11
B-12	Effect Estimation Administration Workbook	B-12
B-13	Effect Estimation Review and Approval Workbook	B-13
C-1	Profile Administration Workbook.....	C-1
C-2	Profile Approval Workbook.....	C-2
C-3	Profile Maintenance Workbook.....	C-2
D-1	Cluster Review Workbook.....	D-1

List of Figures

1-1	Final Level versus Source Level Forecasting	1-3
1-2	Global Domain versus Simple Domain Environment.....	1-5
1-3	Forecast Taskflow Pane.....	1-6
1-4	Promote Taskflow Pane	1-7
1-5	CPEM Taskflow Pane	1-8
1-6	Curve Taskflow Pane	1-8
1-7	Grade Taskflow Pane	1-9
1-8	RDF and the Oracle Retail Enterprise	1-10
2-1	RPAS Fusion Client Login Page.....	2-2
2-2	RPAS Fusion Client Login Page, Accessed Through SSO	2-2
2-3	RPAS Fusion Client Homepage	2-3
2-4	Illustration of Activity, Task, and Step	2-4
2-5	Taskflow Pane	2-4
2-6	Icons in the Taskflow Pane	2-5
2-7	Show List of Workbooks Icon	2-7
2-8	Open Workbook Window	2-7
2-9	Delete Warning Message	2-8
2-10	Renaming a Workbook.....	2-8
2-11	Rename Workbook Dialog Box.....	2-9
2-12	Renamed Workbook	2-9
2-13	Workbook Wizard	2-10
2-14	Online Help Link	2-11
2-15	Online Help Window, Book Field	2-12
2-16	Help Topic Within Online Help.....	2-12
2-17	Online Help: Figure List	2-13
2-18	Online Help: Index.....	2-13
2-19	Online Help: Search.....	2-14
2-20	Logout Link.....	2-14
3-1	AutoES Flowchart	3-6
3-2	Forecast Level Selection Process	3-7
3-3	Simple Exponential Smoothing.....	3-10
3-4	Croston's Method	3-10
3-5	Holt Exponential Smoothing	3-11
3-6	Multiplicative Winters Exponential Smoothing	3-14
3-7	Causal Forecasting Process	3-22
4-1	Product/Location Cloning Administration Task.....	4-2
4-2	Workbook Wizard: Select Item	4-2
4-3	Workbook Wizard: Select Stores.....	4-3
4-4	Adjustment Parameters View	4-3
4-5	Adjustment Ratio Formula	4-4
4-6	Cloned History Formula	4-4
4-7	Location Cloning View.....	4-4

4-8	Product Cloning View	4-6
4-9	Clone Mask View	4-7
5-1	Forecast Administration Task	5-5
5-2	Workbook Wizard: Select Forecast Level	5-6
5-3	Final Level Parameters View	5-7
5-4	Final and Source Level Parameters View	5-8
5-5	Causal Parameters View	5-10
5-6	Final Level Parameters View	5-14
5-7	Final and Source Level Parameters View	5-19
5-8	Regular Price Parameters View	5-23
5-9	Decaying Factor Trend	5-24
5-10	Promo Affinity Parameters View	5-25
6-1	Forecast Approval Task	6-2
6-2	Workbook Wizard: Select Final Forecast Level	6-2
6-3	Workbook Wizard: Select Forecast Levels	6-3
6-4	Workbook Wizard: Select Forecast Generation Date.....	6-4
6-5	Workbook Wizard: Select Stores.....	6-5
6-6	Workbook Wizard: Select Item	6-6
6-7	Workbook Wizard: History First Date.....	6-7
6-8	Workbook Wizard: Forecast Horizon Last Date	6-8
6-9	Workbook Wizard: Calendar Last Date	6-9
6-10	Workbook Wizard: Additional Forecasting Measures.....	6-10
6-11	Advanced Final Level View	6-10
6-12	Final Level View.....	6-12
6-13	Approval View	6-15
6-14	Valid Forecast Run View	6-16
6-15	Source Level View.....	6-16
6-16	System Parameters View	6-17
7-1	Forecast Delete Task	7-1
7-2	Workbook Wizard: Select Forecasts to Delete	7-2
7-3	Workbook Wizard: Verify Delete	7-2
8-1	Forecast Scorecard Task	8-1
8-2	Workbook Wizard: Forecast Level	8-2
8-3	Workbook Wizard: Generation Date	8-3
8-4	Workbook Wizard: Select Error Statistics	8-3
8-5	Workbook Wizard: Calculate Error Statistics	8-4
8-6	Workbook Wizard: Select Item	8-5
8-7	Workbook Wizard: Select Stores.....	8-6
8-8	Workbook Wizard: Select First Date of History	8-6
8-9	Workbook Wizard: Select Last Date of the Horizon	8-7
8-10	Workbook Wizard: Select Last Date in the Calendar	8-8
8-11	Error Measure View	8-9
8-12	Actuals versus Forecast View	8-10
9-1	Interactive Forecasting Task.....	9-2
9-2	Workbook Wizard: Forecast Level	9-2
9-3	Workbook Wizard: Forecast Start Date	9-3
9-4	Workbook Wizard: History Start Date	9-3
9-5	Workbook Wizard: Forecast End Date	9-4
9-6	Workbook Wizard: Forecast Horizon End Date.....	9-5
9-7	Workbook Wizard: Select Stores.....	9-5
9-8	Workbook Wizard: Select Items.....	9-6
9-9	Workbook Wizard: Forecast Results.....	9-6
9-10	Forecast Parameter View	9-7
9-11	Interactive Forecasting View	9-8
10-1	Forecast Like-Item, Sister-Store Task	10-1

10-2	Workbook Wizard: Forecast Level	10-2
10-3	Workbook Wizard: Select Stores.....	10-2
10-4	Workbook Wizard: Select Item	10-3
10-5	Final Like-Item View	10-3
10-6	Final Sister-Store View	10-4
10-7	Advanced Parameter View.....	10-5
11-1	Forecast Maintenance Task.....	11-1
11-2	Workbook Wizard: Causal Level	11-2
11-3	Workbook Wizard: Select Locations	11-2
11-4	Workbook Wizard: Select Item	11-3
11-5	Final Level View - Basic Settings	11-4
11-6	Source Level View - Basic Settings	11-4
11-7	Advanced Final Parameter View.....	11-7
12-1	Run Batch Forecast Task	12-2
12-2	Workbook Wizard: Select Final Forecast Level	12-3
12-3	Workbook Wizard: Verify	12-3
13-1	Promotion Effectiveness Task	13-6
13-2	Workbook Wizard: Causal Level.....	13-6
13-3	Workbook Wizard: Select Promotions.....	13-7
13-4	Workbook Wizard: Select Products	13-7
13-5	Workbook Wizard: Locations	13-8
13-6	Workbook Wizard: Select Calendar Positions	13-8
13-7	Promotion Model Type View.....	13-9
13-8	Promotions View.....	13-9
13-9	Aggregated Promotion Lift Override View	13-10
13-10	Promotional Forecast View	13-11
13-11	Promotion Parameters View	13-13
13-12	Promotion Maintenance Task	13-15
13-13	Workbook Wizard: Causal Level	13-15
13-14	Workbook Wizard: Select Promotions.....	13-16
13-15	Workbook Wizard: Select Products	13-16
13-16	Workbook Wizard: Select Locations	13-17
13-17	Promotion Parameters View	13-17
13-18	Promotion Management Task.....	13-19
13-19	Workbook Wizard: Select Causal Level.....	13-20
13-20	Workbook Wizard: Select Promotions	13-20
13-21	Promo Variable Model Type View	13-21
13-22	Promotion Enable View	13-21
13-23	Accept Negative Lift View.....	13-22
13-24	Promotion Planner Task	13-23
13-25	Workbook Wizard: Select Promotions.....	13-23
13-26	Workbook Wizard: Select Products	13-24
13-27	Workbook Wizard: Select Location.....	13-24
13-28	Workbook Wizard: Select Days	13-25
13-29	Promotions View.....	13-25
14-1	Effect Estimation Administration Task.....	14-2
14-2	Workbook Wizard: Select Products	14-3
14-3	Workbook Wizard: Locations	14-3
14-4	Workbook Wizard: Select RHS Products	14-4
14-5	Default Parameters View	14-4
14-6	Cannibalization Parameter Override View.....	14-5
14-7	Halo Parameter Override View	14-6
14-8	Halo Products Mapping View	14-6
14-9	Effect Estimation Review and Approval Task.....	14-7
14-10	Workbook Wizard: Select Products	14-8

14-11	Workbook Wizard: Locations	14-8
14-12	Workbook Wizard: Select RHS Products	14-9
14-13	Workbook Wizard: Select Dates	14-9
14-14	Cannibalization Effects View	14-10
14-15	Cannibalization Effects Drive View	14-10
14-16	Cannibalization History Data View	14-11
14-17	Halo Effects View	14-12
14-18	Halo Effects Drive View	14-13
14-19	Halo History Data View	14-14
14-20	Cannibalization Statistic Report View	14-15
14-21	Halo Statistic Report View	14-15
15-1	Profile Administration Task	15-2
15-2	Workbook Wizard: Curve Level	15-3
15-3	Profile Parameter View	15-3
15-4	Profile and Source Level Intersection View	15-7
15-5	Profile Approval Task	15-8
15-6	Workbook Wizard: Curve Level	15-8
15-7	Workbook Wizard: Select Item	15-9
15-8	Workbook Wizard: Select Stores	15-9
15-9	Workbook Wizard: Select Day	15-10
15-10	Final Profile View	15-10
15-11	Source Profile View	15-11
15-12	Profile Approval View	15-12
15-13	Profile Maintenance Task	15-13
15-14	Workbook Wizard: Curve Level	15-13
15-15	Workbook Wizard: Select Items	15-14
15-16	Workbook Wizard: Select Stores	15-14
15-17	Workbook Wizard: Select Day	15-15
15-18	Final Approval and Sourcing View	15-15
15-19	Final Training Window View	15-16
15-20	Run Batch Profile Task	15-17
15-21	Workbook Wizard: Curve Level	15-17
15-22	Workbook Wizard: Verify Curve level	15-18
16-1	Breakpoints Administration Task	16-4
16-2	Breakpoint Administration View	16-4
16-3	Delete Cluster Run Task	16-5
16-4	Cluster Review Task	16-6
16-5	Workbook Wizard: Select Grade Birth	16-6
16-6	Cluster Input Summary View	16-7
16-7	Generate Breakpoint Grades Task	16-8
16-8	Workbook Wizard: Source Measure for Cluster Analysis	16-8
16-9	Workbook Wizard: Hierarchy and Dimension to Cluster On	16-9
16-10	Workbook Wizard: Select Hierarchy and Dimension to Group By	16-10
16-11	Workbook Wizard: Select Items	16-10
16-12	Workbook Wizard: Select Stores	16-11
16-13	Workbook Wizard: Select Day	16-11
16-14	Workbook Wizard: Select Configuration Name	16-12
16-15	Generate Clusters Task	16-13
16-16	Workbook Wizard: Select Source Measure for Cluster Analysis	16-13
16-17	Workbook Wizard: Hierarchy and Dimension	16-14
16-18	Workbook Wizard: Select Hierarchy and Dimension to Group By	16-15
16-19	Workbook Wizard: Select Items	16-15
16-20	Workbook Wizard: Select Stores	16-16
16-21	Workbook Wizard: Select Day	16-16
16-22	Workbook Wizard: Set Name and Number of Clusters	16-17

16-23	Workbook Wizard: Cluster Run Succeeded	16-18
A-1	Preprocessing for Stock-outs	A-1
A-2	Standard Median Formula.....	A-2
A-3	Standard Median Example	A-2
A-4	Retail Median Formula.....	A-3
A-5	Retail Median Example	A-3
A-6	Standard Exponential Smoothing Formula.....	A-4
A-7	Standard Exponential Smoothing Example	A-4
A-8	Lost Sales Standard Exponential Smoothing Formula	A-5
A-9	Lost Sales Standard Exponential Smoothing Example.....	A-5
A-10	Override Formula	A-6
A-11	Override Example	A-6
A-12	Increment Formula	A-7
A-13	Increment Example.....	A-7

Send Us Your Comments

Oracle Retail Demand Forecasting User Guide for the RPAS Fusion Client, Release 13.4.1

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

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

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

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

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

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

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

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

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

Preface

The Oracle Retail Demand Forecasting User Guide for the RPAS Fusion Client describes the application's user interface and how to navigate through it.

Audience

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

Documentation Accessibility

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

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

Access to Oracle Support

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

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

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

Related Documents

For more information, see the following documents in the Oracle Retail Demand Forecasting Release 13.4.1 documentation set:

- *Oracle Retail Demand Forecasting Configuration Guide*
- *Oracle Retail Demand Forecasting Installation Guide*
- *Oracle Retail Demand Forecasting Release Notes*
- *Oracle Retail Demand Forecasting User Guide for the RPAS Classic Client*
- Oracle Retail Predictive Application Server documentation

Supplemental Documentation

The following document is available through My Oracle Support at the following URL:

<https://support.oracle.com>

Oracle Retail Demand Forecasting 13.4.1 Cumulative Fixed Issues (Note ID 1547653.1)

This document details the fixed issues and defects for all RDF, Curve, and Grade patch releases prior to and including the current release.

Customer Support

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

<https://support.oracle.com>

When contacting Customer Support, please provide the following:

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

Review Patch Documentation

When you install the application for the first time, you install either a base release (for example, 13.4) or a later patch release (for example, 13.4.1). If you are installing the base release, additional patch, and bundled hot fix releases, read the documentation for all releases that have occurred since the base release before you begin installation. Documentation for patch and bundled hot fix releases can contain critical information related to the base release, as well as information about code changes since the base release.

Oracle Retail Documentation on the Oracle Technology Network

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

http://www.oracle.com/technology/documentation/oracle_retail.html

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

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

Conventions

The following text conventions are used in this document:

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

Introduction

Oracle Retail Demand Forecasting (RDF) is a statistical and promotional forecasting solution. It uses state-of-the-art modeling techniques to produce high quality forecasts with minimal human intervention. Forecasts produced by the RDF system enhance the retailer's supply-chain planning, allocation, and replenishment processes, enabling a profitable and customer-oriented approach to predicting and meeting product demand.

Today's progressive retail organizations know that store-level demand drives the supply chain. The ability to forecast consumer demand productively and accurately is vital to a retailer's success. The business requirements for consumer responsiveness mandate a forecasting system that more accurately forecasts at the point of sale, handles difficult demand patterns, forecasts promotions and other causal events, processes large numbers of forecasts, and minimizes the cost of human and computer resources.

Forecasting drives the business tasks of planning, replenishment, purchasing, and allocation. As forecasts become more accurate, businesses run more efficiently by buying the right inventory at the right time. This ultimately lowers inventory levels, improves safety stock requirements, improves customer service, and increases the company's profitability.

The competitive nature of business requires that retailers find ways to cut costs and improve profit margins. The accurate forecasting methodologies provided with RDF can provide tremendous benefits to businesses.

A connection from RDF to Oracle Retail's Planning and Optimization suite is built directly into the business process by way of the automatic approvals of forecasts, which may then be fed directly to any Planning and Optimization suite application. This process allows you to accept all or part of a generated sales forecast. Once that decision is made, the remaining business measures may be planned within an Planning and Optimization suite application such as Merchandise Financial Planning (MFP).

Forecasting Challenges and RDF Solutions

A number of challenges affect the ability of organizations to forecast product demand accurately. These challenges include selecting the best forecasting method to account for:

- Level
- Trending
- Seasonal demand

- Spiky demand
- Generating forecasts for items with limited demand histories
- Forecasting demand for new products and locations
- Incorporating the effects of promotions and other event-based challenges on demand
- Accommodating the need of operational systems to have sales predictions at more detailed levels than planning programs provide

Selecting the Best Forecasting Method

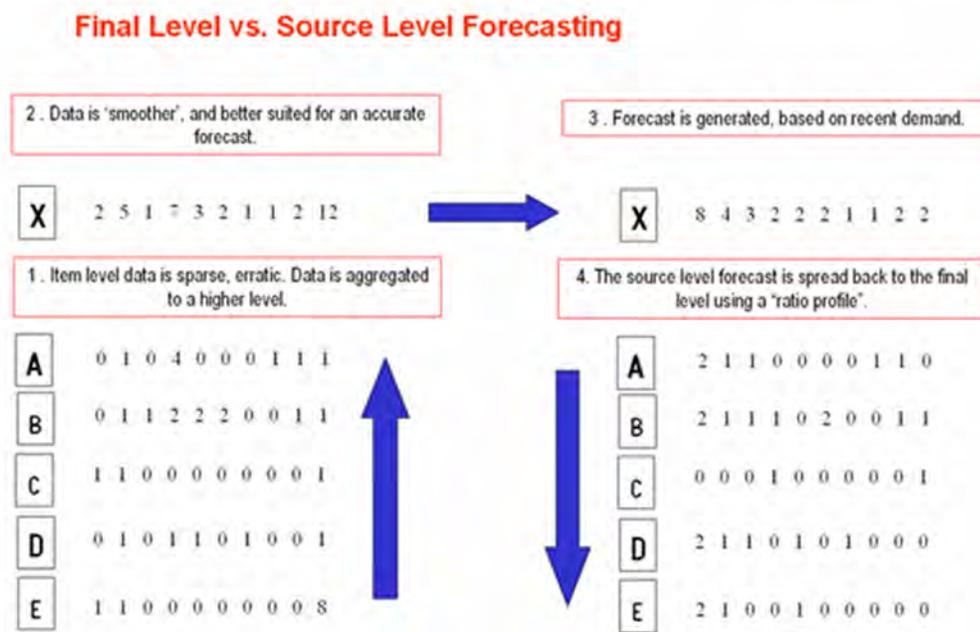
One challenge to accurate forecasting is the selection of the best model to account for level, trending, seasonal, and spiky demand. Oracle Retail's AutoES (Automatic Exponential Smoothing) forecasting method eliminates this complexity.

The AutoES method evaluates multiple forecast models, such as Simple Exponential Smoothing, Holt Exponential Smoothing, Additive and Multiplicative Winters Exponential Smoothing, Croston's Intermittent Demand Model, and Seasonal Regression forecasting to determine the optimal forecast method to use for a given set of data. The accuracy of each forecast and the complexity of the forecast model are evaluated in order to determine the most accurate forecast method. You simply select the AutoES forecast generation method and the system finds the best model.

Overcoming Data Sparsity Through Source Level Forecasting

It is a common misconception in forecasting that forecasts must be directly generated at the lowest levels (final levels) of execution. Problems can arise when historic sales data for these items is too sparse and noisy to identify clear selling patterns. In such cases, generating a reliable forecast requires aggregating the sales data from a final level up to a higher level (source-level) in the hierarchy in which demand patterns can be seen, and then generate a forecast at this source-level. After a forecast is generated at the source-level, the resulting data can be allocated (spread) back down to the lower level, based on the lower level's (final level) relationship to the total. This relationship can then be determined through generating an additional forecast (interim forecast) at the final level. Curve is then used to dynamically generate a profile based on the interim forecasts. Also, a non-dynamic profile can be generated and approved to be used as this profile. It is this profile that determines how the source-level forecast is spread down to the final level. For more information on Curve, refer to [Chapter 15, "Curve"](#).

Figure 1–1 Final Level versus Source Level Forecasting



Some high-volume items may possess sufficient sales data for robust forecast calculations directly at the final forecast level. In these cases, forecast data generated at an aggregate level and then spread down to lower levels can be compared to the interim forecasts run directly at the final level. Comparing the two forecasts, each generated at a different hierarchy level, can be an invaluable forecast performance evaluation tool.

Your RDF system may include multiple final forecast levels. Forecast data must appear at a final level for the data to be approved and exported to another system for execution.

Forecasting Demand for New Products and Locations

RDF also forecasts demand for new products and locations for which no sales history exists. You can model a new product's demand behavior based on that of an existing similar product for which you do have a history. Forecasts can be generated for the new product based on the history and demand behavior of the existing one. Likewise, the sales histories of existing store locations can be used as the forecast foundation for new locations in the chain. For more details, refer to the [Forecast Like-Item, Sister-Store Workbook](#) section.

Managing Forecasting Results Through Automated Exception Reporting

The RDF end user may be responsible for managing the forecast results for thousands of items, at hundreds of stores, across many weeks at a time. The Oracle Retail Predictive Application Server (RPAS) provides users with an automated exception reporting process (called Alert Management) that indicates to you where a forecast value may lie above or below an established threshold, thereby reducing the level of interaction needed from you.

Alert management is a feature that provides user-defined and user-maintained exception reporting. Through the process of alert management, you define measures

that are checked daily to see if any values fall outside of an acceptable range or do not match a given value. When this happens, an alert is generated to let you know that a measure may need to be examined and possibly amended in a workbook.

The Alert Manager is a dialog box that is displayed automatically when you log on to the system. This dialog provides a list of all identified instances in which a given measure's values fall outside of the defined limits. You may pick an alert from this list and have the system automatically build a workbook containing that alert's measure. In the workbook, you can examine the actual measure values that triggered the alert and make decisions about what needs to be done next.

For more information on the Alert Manager, refer to the *Oracle Retail Predictive Application Server User Guide*.

Incorporating the Effects of Promotions and Other Event-Based Challenges on Demand

Promotions, non-regular holidays, and other causal events create another significant challenge to accurate forecasting. Promotions such as advertised sales and free gifts with purchase might have a significant impact on a product's sales history, as can fluctuating holidays such as Easter.

Using Promotional Forecasting (an optional, add-on module to RDF), promotional models of forecasting can be developed to take these and other factors into account when forecasts are generated. Promotional Forecasting attempts to identify the causes of deviations from the established seasonal profile, quantify these effects, and use the results to predict future sales when conditions in the selling environment is similar. This type of advanced forecasting identifies the behavioral relationship of the variable you want to forecast (sales) to both its own past and explanatory variables such as promotion and advertising.

Suppose that your company has a large promotional event during the Easter season each year. The exact date of the Easter holiday varies from year to year, as a result, the standard time-series forecasting model often has difficulty representing this effect in the seasonal profile. The Promotional Forecasting module allows you to identify the Easter season in all years of your sales history, and then define the upcoming Easter date. By doing so, you can causally forecast the Easter-related demand pattern shift.

Oracle Retail Demand Forecasting Architecture

This section includes these topics:

- [The Oracle Retail Predictive Application Server and RDF](#)
- [Global Domain versus Simple Domain Environment](#)

The Oracle Retail Predictive Application Server and RDF

The RDF application is a member of the Planning and Optimization suite, including other solutions such as Merchandise Financial Planning, Item Planning, Category Management, and Advance Inventory Planning. The Planning and Optimization suite applications share a common platform called the Oracle Retail Predictive Application Server (RPAS). RDF leverages the versatility, power, and speed of the RPAS engine and user-interface.

These features characterize RPAS:

- Multidimensional databases and database components (dimensions, positions, hierarchies)

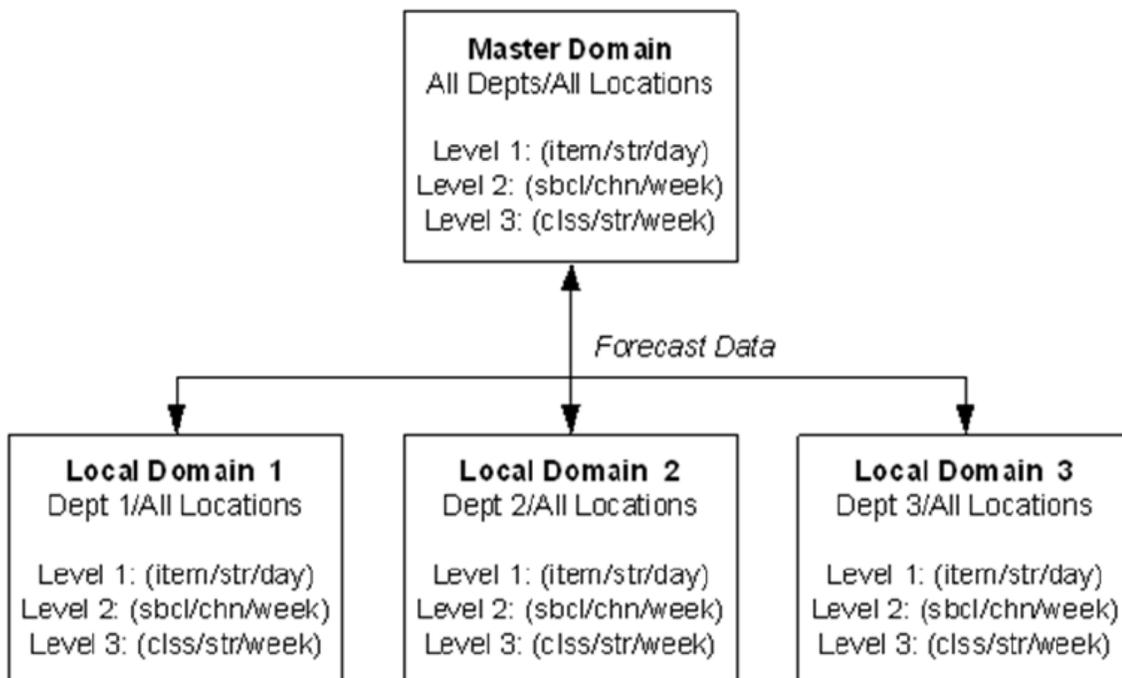
- Product, location, and calendar hierarchies
- Aggregation and spreading of sales data
- Client-server architecture and master database
- Workbooks and views for displaying and manipulating forecast data
- Wizards for creating and formatting workbooks and views
- Menus, quick menus, and toolbars for working with sales and forecast data
- An automated alert system that provides user-defined and user-maintained exception reporting
- Charting and graphing capabilities

More details about the use of these features can be found in the *Oracle Retail Predictive Application Server User Guide* and online help provided within your RDF solution.

Global Domain versus Simple Domain Environment

A simple domain environment supports isolated partitions of data. This type of environment does not allow for data to be aggregated across partitions into a single view. Whereas a global domain environment allows for data partitions to exist, however, certain data may be edited and viewed across partitions. Within this structure we refer to data within a partition as the Local domain (or sub-domain) and the view to data across multiple local domains as the Master domain. [Figure 1–2](#) represents a global domain environment:

Figure 1–2 Global Domain versus Simple Domain Environment



Within this structure, batch forecast results across all domains may be viewed within the master domain. This is achieved by passing measures with the same forecast birth date (date/time stamp) to each Local Domain when the batch forecast is generated. It is important to note that this cannot be achieved through the use of the [Run Batch](#)

[Forecast Workbook](#) wizard. The *Oracle Retail Demand Forecasting Implementation Guide* provides more information on execution of batch forecast processes to support a global domain environment.

The RDF solution in a global domain environment also supports centralized administration and maintenance of forecast parameters in the Master domain. Additional details on the availability and limitations of all of the workbooks in the Master domain environment are provided in the following sections.

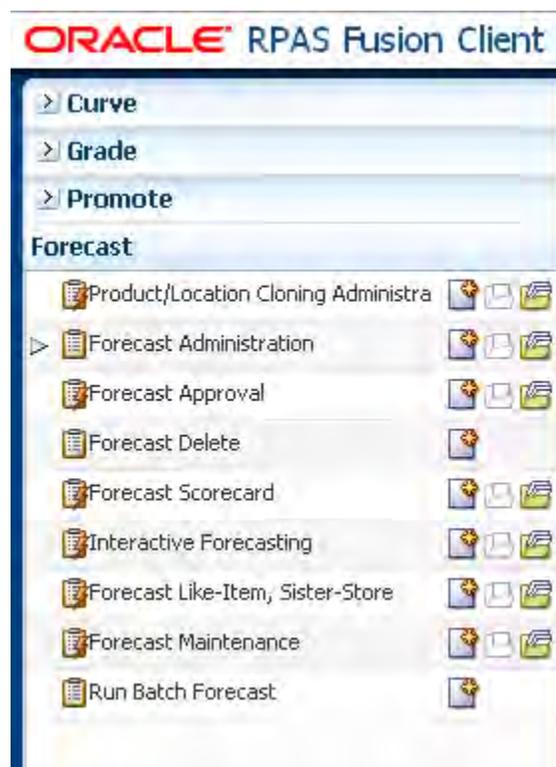
Oracle Retail Demand Forecasting Workbook Groups

In addition to the standard RPAS Administration and Analysis workbook groups, there are several workbook groups that are associated with the RDF solution which may include: Forecast, Promote, Curve, or any Planning and Optimization suite application (available modules are based upon licensing agreement).

Forecast

The Forecast module refers to the primary RDF functionality and consists of the workbooks, measures, and forecasting algorithms that are needed to perform time-series forecasting.

Figure 1–3 Forecast Taskflow Pane



This module includes the:

- [Product/Location Cloning Administration Workbook](#)
- [Forecast Administration Workbook](#)

- [Forecast Approval Workbook](#)
- [Forecast Delete Wizard](#)
- [Forecast Scorecard Workbook](#)
- [Interactive Forecasting Workbook](#)
- [Forecast Like-Item, Sister-Store Workbook](#)
- [Forecast Maintenance Workbook](#)
- [Run Batch Forecast Workbook](#)

The Forecast module also includes the batch forecasting routine and all of its component algorithms.

Promote (Promotional Forecasting)

The Promote module consists of the workbooks and algorithms required to perform promotional forecasting, which uses both past sales data and promotional information (for example, advertisements, holidays) to forecast future demand.

Figure 1–4 Promote Taskflow Pane



This module includes the:

- [Promotion Effectiveness Workbook](#)
- [Promotion Maintenance Workbook](#)
- [Promotion Management Workbook](#)
- [Promotion Planner Workbook](#)

For more information on the Promote workbooks and views, refer to [Chapter 13, "Promote \(Promotional Forecasting\)."](#)

Cross Promo Effect Module (CPEM)

The Cross Promo Effect Module consists of the workbooks and algorithms that perform data mining which determines promotional Cannibalization, or Halo relationships, or both between items or groups of items.

Figure 1–5 CPEM Taskflow Pane

This module includes the:

- [Effect Estimation Administration Task](#)
- [Effect Estimation Review and Approval Task](#)

For more information on the CPEM workbooks and views, refer to [Chapter 14, "Cross Promotional Effects Module \(CPEM\)."](#)

Curve

The Curve module consists of the workbooks and batch algorithms that are necessary for the creation, approval, and application of profiles that may be used to spread source-level forecasts down to final levels as well to generate profiles, which may be used in any RPAS solution. The types of profiles typically used to support forecasting are: store Contribution, Product, and Daily profiles. These profiles may also be used to support Profile-Based Forecasting; however, Curve may be used to generate profiles that are used by other Planning and Optimization suite applications for reasons other than forecasting. Profiles Types include Daily Seasonal, Lifecycle, Size, Hourly, and User-Defined profiles.

Figure 1–6 Curve Taskflow Pane

This module includes the:

- [Profile Administration Workbook](#)

- [Profile Approval Workbook](#)
- [Profile Maintenance Workbook](#)
- [Run Batch Profile Workbook](#)

For more information on the Curve workbooks and views, refer to [Chapter 15, "Curve."](#)

Grade

The Grade module consists of the workbooks and wizards that are necessary for clustering. Typically, a retailer may cluster stores over item sales to create logical groupings of stores based upon sales of particular products. This provides increased visibility to where products are selling, and it allows the retailer to make more accurate decisions in merchandising. Beyond this traditional use of clusters, Grade is flexible enough to cluster any business measure based on products, locations, time, promotions, customers, or any hierarchy configured in the solution.

Figure 1–7 *Grade Taskflow Pane*



This module includes the:

- [Breakpoints Administration Workbook](#)
- [Delete Cluster Run Wizard](#)
- [Cluster Review Workbook](#)
- [Generate Breakpoint Grades Wizard](#)
- [Generate Clusters Wizard](#)

For more information on the Grade workbooks and views refer to [Chapter 16, "Grade."](#)

RDF Solution and Business Process Overview

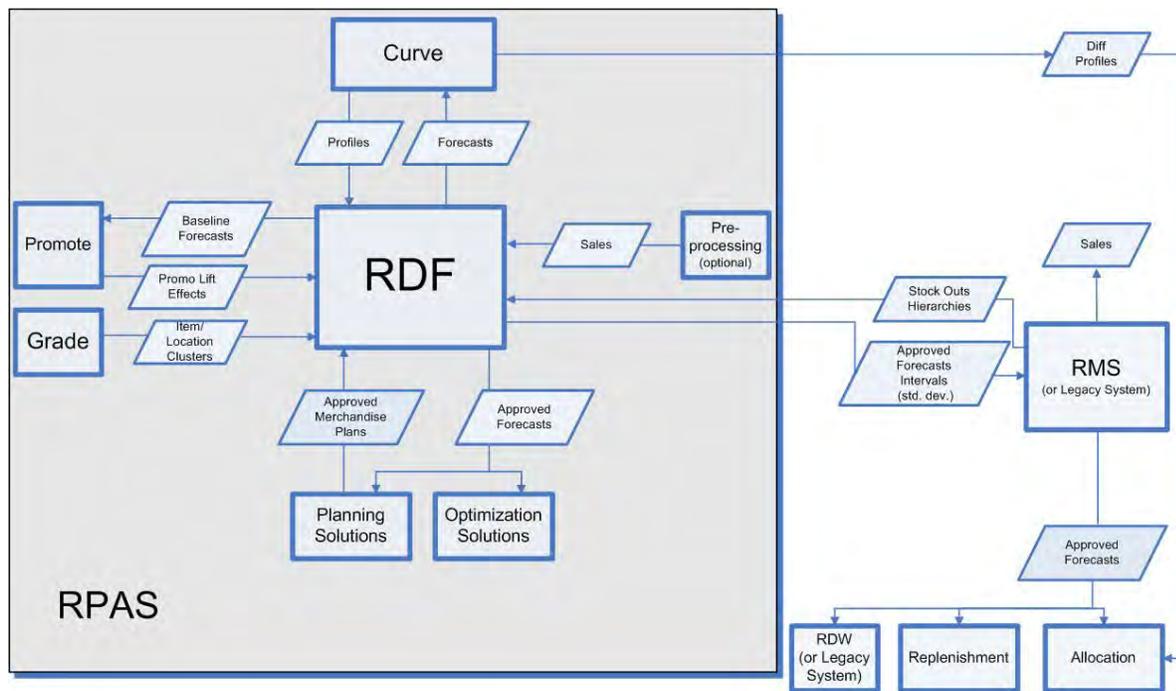
Oracle Retail has designed a forecasting solution separate from replenishment, allocation or planning. In order to provide a single version of the truth, it is crucial to free up your time and supply the tools to focus on the analysis of forecast exceptions, historical data, and different modeling techniques. This empowers you to make better decisions, thus improving overall accuracy and confidence in the forecast downstream.

RDF and the Oracle Retail Enterprise

Within the Oracle Retail Enterprise, Oracle Retail Merchandising System (RMS) supplies RDF with Point-of-Sale (POS) and hierarchy data that is used to create a forecast. Once the forecast is approved, it is exported to RMS in order to calculate a recommended order quantity. Forecasts can also be utilized (no export process required) in any RPAS solution to support merchandise, financial, collaborative, and price planning processes.

Figure 1–8 illustrates an overview of RDF and the Oracle Retail Enterprise.

Figure 1–8 RDF and the Oracle Retail Enterprise



RDF Primary Workflow

There are a number of core super-user/end-user forecasting steps in the RDF workflow that are essential for producing accurate forecasts for the millions of item and location combinations that exist in a domain.

Navigation

This chapter provides a brief description of how to get started using RDF on the RPAS Fusion Client. It includes the following sections:

- [Before you Begin](#)
- [Logging into RDF](#)
- [Taskflow Pane](#)
- [Opening a Workbook](#)
- [Deleting a Workbook](#)
- [Renaming a Workbook](#)
- [Understanding the Workbook Wizard Window](#)
- [Accessing Online Help](#)
- [Logging Out of RDF](#)

Before you Begin

Note: Before you log on to the RPAS Fusion Client, ensure that your system meets the recommended configuration. For more information, refer to the *Oracle Retail Predictive Application Server Installation Guide*.

Once you check the configuration, obtain the following information:

- Uniform Resource Locator, URL – Enter the URL or the Web address of the application in the Web browser to access the application. For example:
`http://<fullyqualifieddomainname>:<port>/aip`
- User name, Password, and Profile Name – Based on the tasks you want to perform, obtain a user account (that includes user name and password) to log on to the application.

Logging into RDF

To log on to the RPAS Fusion Client, follow these steps:

1. Start a supported browser such as IE, Firefox, or Chrome.
2. In the **Address** bar, enter the RPAS Fusion Client URL and press **Enter**. The RPAS Fusion Client Login page opens, as shown in [Figure 2-1](#).

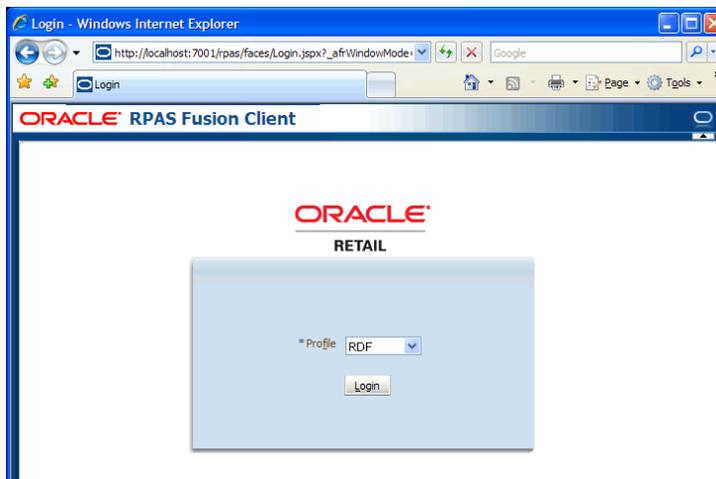
Figure 2–1 RPAS Fusion Client Login Page



3. On the **Login** page, enter your Username and Password, and select your domain from the **Profile** drop-down list.

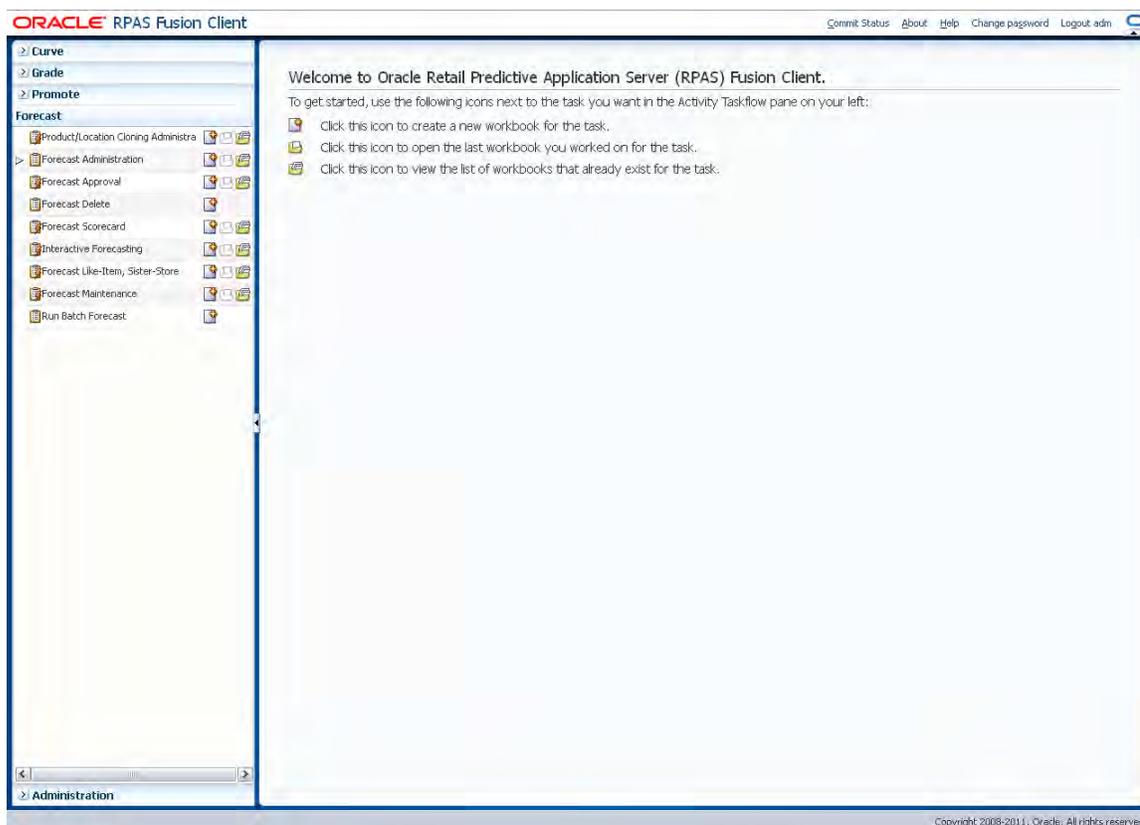
Note: If you have accessed the RPAS Fusion Client through a single sign-on environment such as the Oracle Retail Workspace portal, then only the profile field appears on the RPAS Fusion Client login page. Select your domain from the drop-down list, as shown in [Figure 2–2](#).

Figure 2–2 RPAS Fusion Client Login Page, Accessed Through SSO



4. Click **Login**. The window refreshes and the homepage of the RPAS Fusion Client is shown.

Figure 2–3 RPAS Fusion Client Homepage



Taskflow Pane

The taskflow pane opens on the left of your window and enables you to easily navigate through the activities in the application. It provides a pre-configured business workflow organized into expandable and collapsible activities, tasks, and steps.

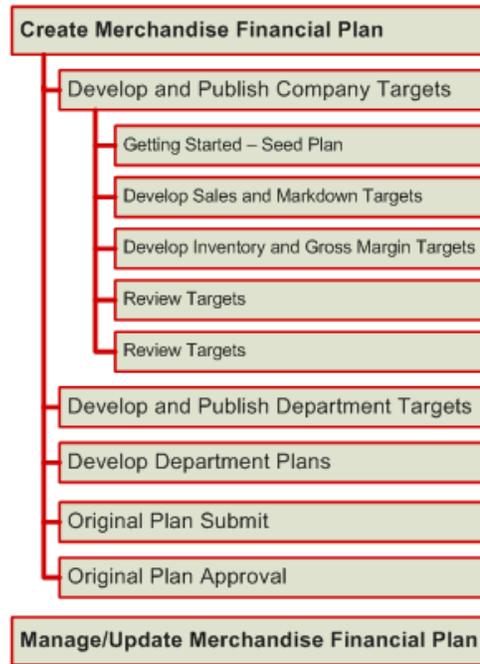
Note: The taskflow pane also opens when you log on to the application.

Each activity consists of more than one task, and each task may consist of one or more steps. In the RPAS Fusion Client, each solution (spanning across multiple workbooks) is represented as a set of activities, tasks, and steps.

Note: The taskflow is configured for your domain during implementation. For more information on this configuration, refer to the *Oracle Retail Predictive Application Server Configuration Tools User Guide*.

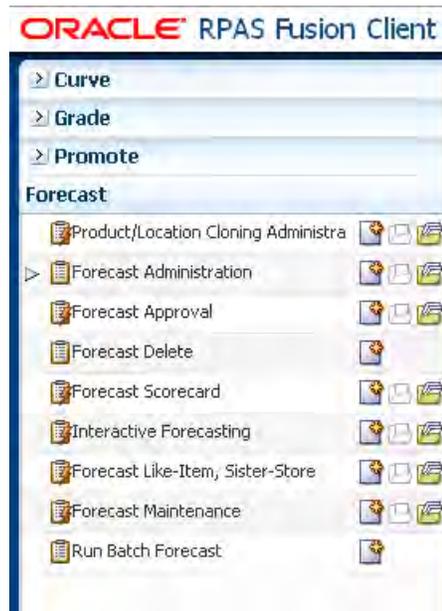
Figure 2–4 shows an example of an RPAS solution configured with activities, tasks, and steps:

Figure 2–4 Illustration of Activity, Task, and Step



In the RPAS Fusion Client, the workflow shown in [Figure 2–5](#) opens in the following manner:

Figure 2–5 Taskflow Pane



When you are working with a specific step, a left arrow icon appears in the taskflow pane, indicating the current step and your position in the workflow.

The Create New Workbook, Open Latest Workbook, and Show List of Workbooks icons that appear next to each task or step enable you to create new workbooks or

open existing workbooks. You can find more information on creating new workbooks and opening existing workbooks in the section [Opening a Workbook](#).

Figure 2–6 illustrates the icons that appear in an RPAS taskflow.

Figure 2–6 Icons in the Taskflow Pane

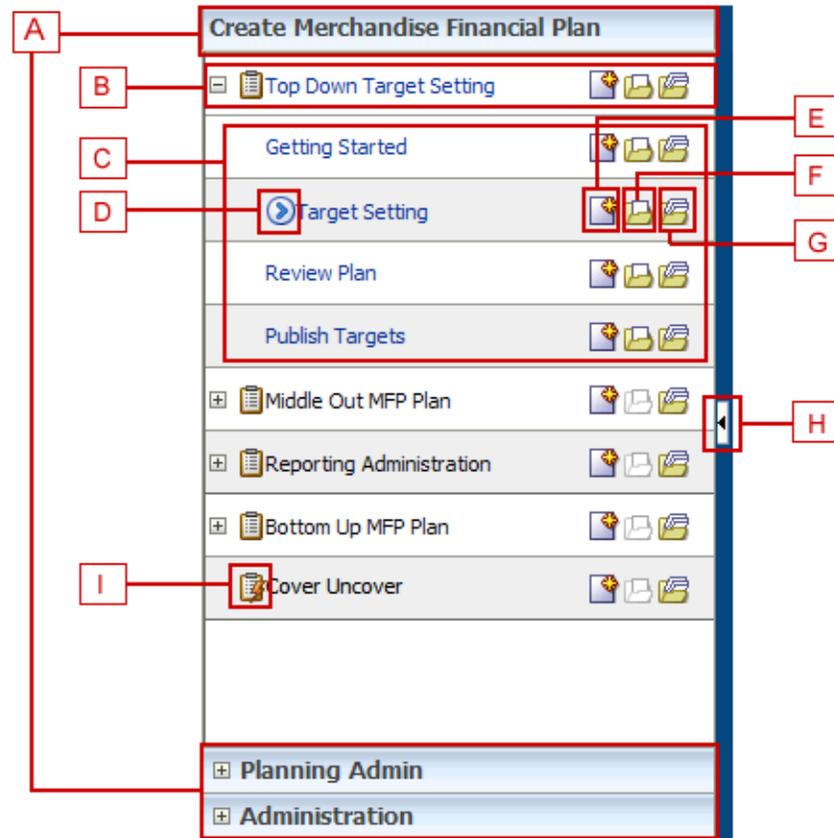


Table 2–1 describes the icons that appear with all the entries in the taskflow.

Table 2–1 Description of Icons in the Taskflow Pane

Legend	Icon Name	Description
A	Activities	These tabs represent the predefined activities of the application.
B	Tasks	These are individual tasks within an activity. Tasks that have a workbook associated with them display a clipboard icon.
C	Steps	One or more steps make up each task.
D	Current Position Icon	Indicates the current step and your position in the workflow.
E	Create a New Workbook	Click this icon to create a new workbook. The new workbook wizard associated with the relevant task opens.
F	Open the Latest Workbook	Click this icon to view the latest workbook (associated with the relevant task/step) you worked on.

Table 2–1 (Cont.) Description of Icons in the Taskflow Pane

Legend	Icon Name	Description
G	Show List of Workbooks	Click this icon to view a list of all workbooks (associated with the relevant task/step) accessible to you. The Open Workbook window opens with a list of accessible workbooks. If you do not own any workbooks, a message appears allowing you to view all workbooks. In case there are no workbooks associated with the task/step, created by you or by someone else, a message appears that allows you to create a new workbook.
H	Collapse/Restore Icon	Click this icon to collapse or restore the taskflow pane. This allows you to view more content in the Contents area.
I	Dynamic Task	A dynamic task is a task that has been defined in the taskflow file but does not have a workbook associated with it. Dynamic tasks display a clipboard with a lightning bolt icon.

Access-based Visibility

The activity and tasks that appear in the taskflow pane are access-based. Depending upon the security settings, you may not have access to some tasks or activities. Access to a task is defined by whether you have access to the workbook that the task is assigned to.

The access to the workbook is maintained in the Security Administration step. Refer to the *Oracle Retail Predictive Application Server Administration Guide for the Fusion Client* for more information about the Administration activity. If you do not have access to a workbook, then you are not able to build the workbooks for the associated tasks and steps.

Hiding specific tasks or activities based on user access is configurable. In the configuration, a task can be set to either hidden or disabled. If it is hidden and you do not have access to it, then the task is not displayed. If it is disabled and you do not have access, then the task is displayed but the links to access the workbooks are disabled. For more information, refer to the *Oracle Retail Predictive Application Server Configuration Tools User Guide*.

Switching Between Multiple Tasks

When working within multiple steps or tasks, all changes you make in a specific step are maintained when you move to a step in the same task or a different task associated with the same workbook. In this case, you do not need to save your work when you switch between tasks within the same workbook.

If you switch to a task that is associated with a different workbook, you are prompted to save or discard the changes before opening or building a new workbook. To save your changes, use the Save As dialog box.

Opening a Workbook

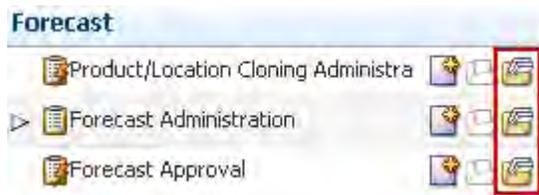
After you log on to the application, a taskflow pane opens that enables you to navigate through the activities and tasks associated with your user account.

Note: This procedure outlines the basic steps for creating a new workbook. For detailed information about building specific RDF workbooks, refer to the specific chapter for that workbook in this guide.

To open a workbook:

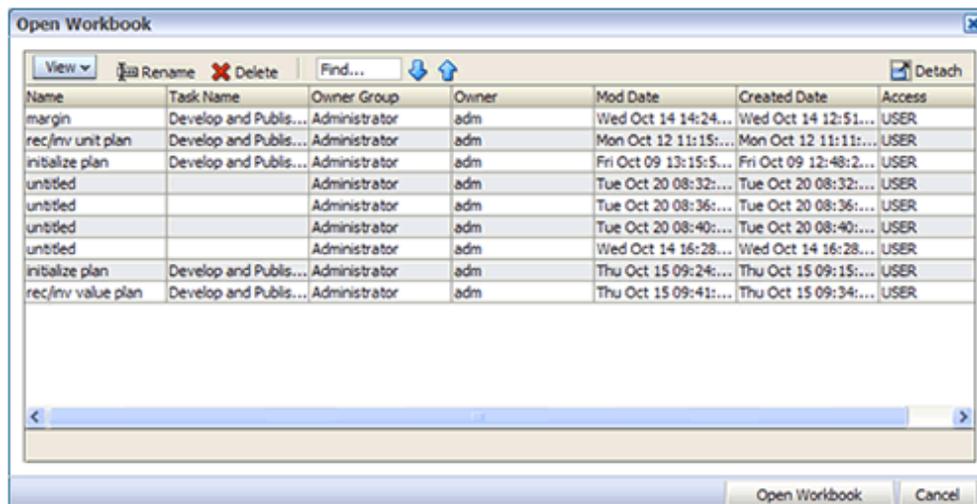
1. On the taskflow pane, click the **Show List of Workbooks** icon next to the task or step you want. For more information on the taskflow pane, refer to the section [Taskflow Pane](#).

Figure 2–7 Show List of Workbooks Icon



The **Open Workbook** window opens.

Figure 2–8 Open Workbook Window



2. Select the workbook you want and click **Open Workbook**.

Deleting a Workbook

From the list of workbooks, you can delete any workbook you have write access to.

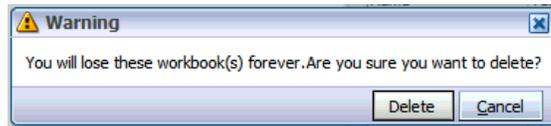
To delete a workbook:

1. On the Open Workbook window, select the workbook you want.

To delete more than one workbook, hold down the **Ctrl** key and select the workbooks you want to delete. Or, you can hold down **Shift** and select a group of workbooks.

- Click **Delete**. A warning message appears.

Figure 2–9 Delete Warning Message



- Review the warning message and then click **Delete** to delete the workbooks or click **Cancel** to cancel this operation.

Renaming a Workbook

The new workbook wizard also enables you to rename existing workbooks without affecting the data within the workbook or the other workbook information such as the created date, modified date, formatting information, and so on.

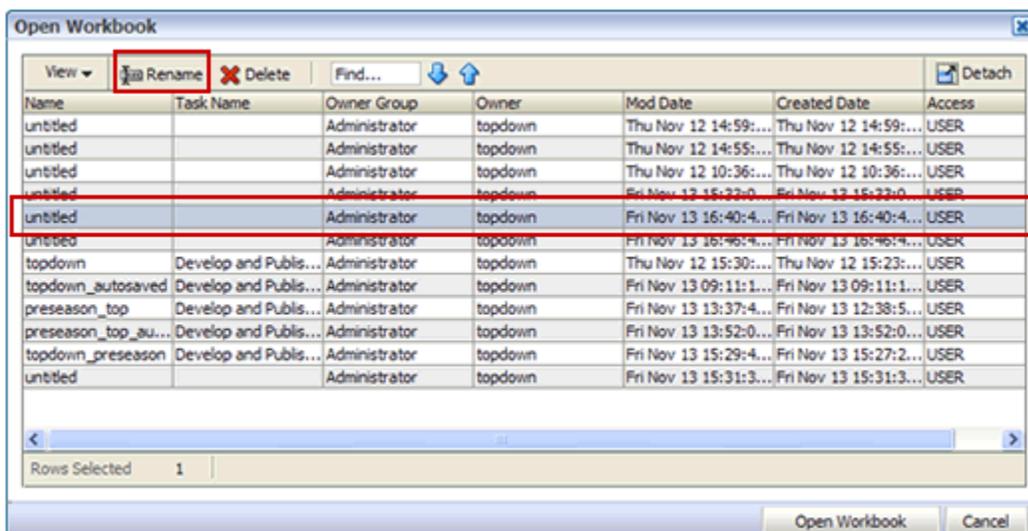
There are a few key points to know about renaming workbooks:

- You can only rename workbooks that you have write access to.
- Workbook names can be no more than 32 characters.
- Workbook names cannot contain double or single quotation marks.
- Workbooks cannot be named *Untitled*. This name is reserved.

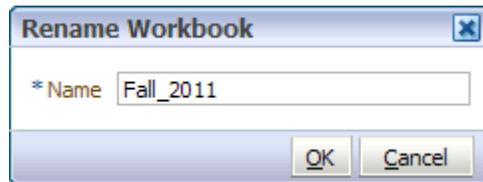
To rename a workbook in the new workbook wizard, perform the following steps:

- Select the workbook you want to rename from the list. When selected, it becomes shaded, as shown in [Figure 2–10](#).
- Click **Rename**.

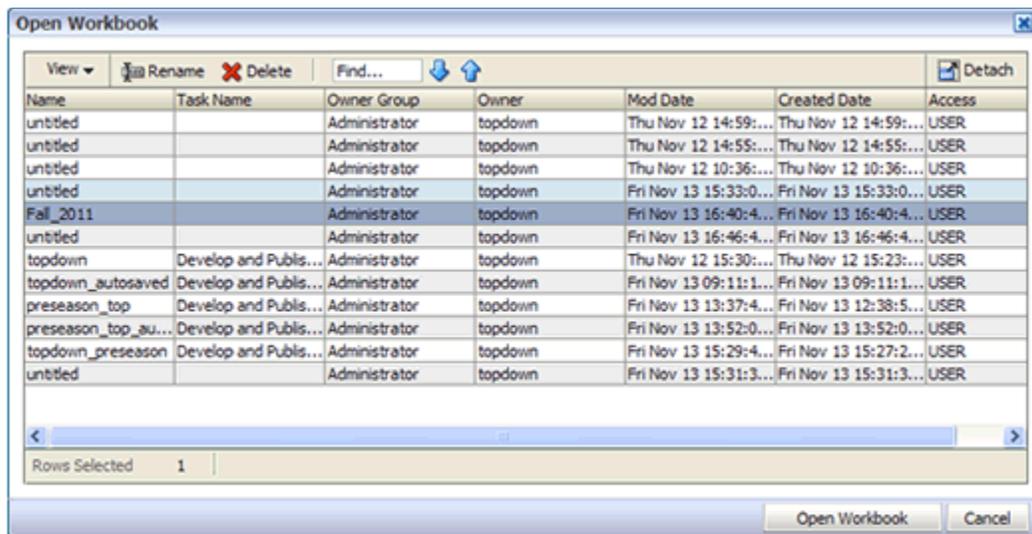
Figure 2–10 Renaming a Workbook



- The Rename Workbook dialog box appears. Enter the new name of the workbook and click **OK**.

Figure 2–11 Rename Workbook Dialog Box

4. The renamed workbook appears in the workbook list.

Figure 2–12 Renamed Workbook

You can also rename a workbook from the File menu.

Understanding the Workbook Wizard Window

Figure 2–13 highlights the various components of the workbook wizard window.

Figure 2–13 Workbook Wizard

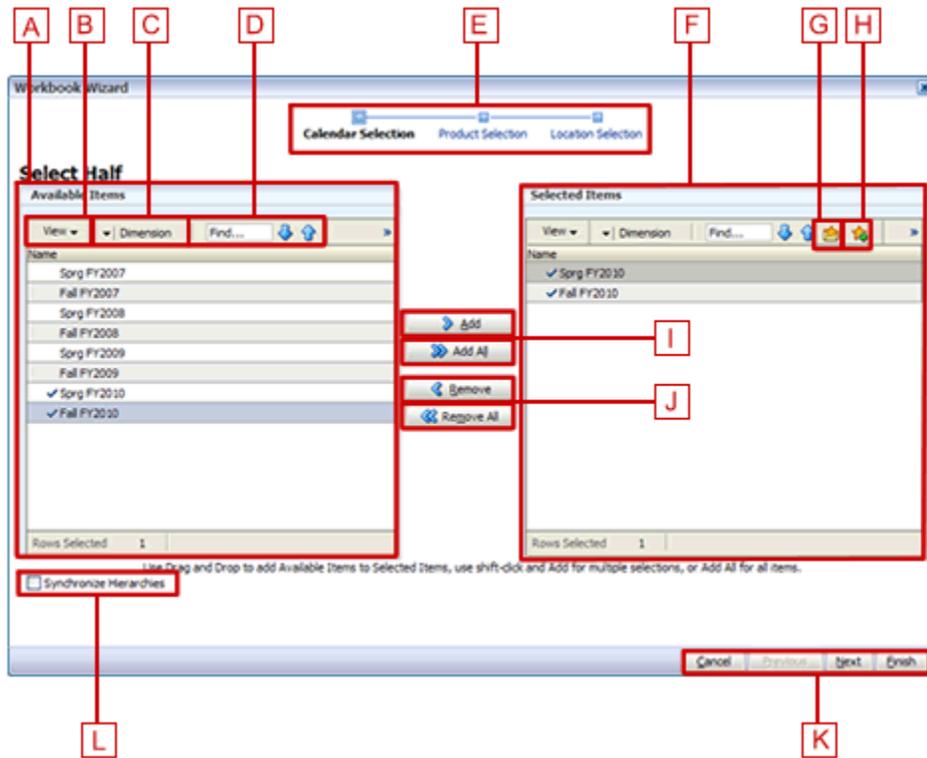


Table 2–2 describes the components of the workbook wizard window.

Table 2–2 Workbook Wizard Window User Interface Components

Legend	Window Area Name	Position	Description
A	Available Items area	Center Left	Displays the positions that are available for you to select.
B	View menu	Center Left, within the Available Items area	Provides options for viewing the available positions. You can adjust the column setting, detach the list to view it in a larger window, expand or collapse the positions, or scroll to the beginning or end of the list.
C	Dimension menu	Center Left, within the Available Items area	Opens the Dimension options window where you can select the levels of the dimension you want to view in the Available Items list. You can also select the attributes that you would like to see and sort by in the list.
D	Find field, Previous and Next arrows	Center Left, within the Available Items area	Enables you to identify the positions with the keywords you entered in the Find field. Use the Previous and Next arrows to navigate between the workbooks matching the search pattern. The search is not case-sensitive.

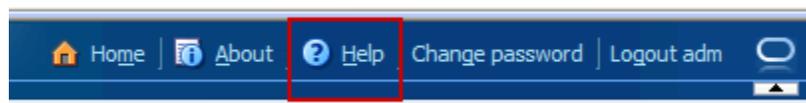
Table 2–2 (Cont.) Workbook Wizard Window User Interface Components

Legend	Window Area Name	Position	Description
E	Wizard Taskflow	Top Center	Displays the steps in the wizard process and shows you where you are within that process. The wizard taskflow is configured in the RPAS Configuration Tools. For more information, refer to the <i>Oracle Retail Predictive Application Server Configuration Tools User Guide</i> .
F	Selected Items area	Center Right	Displays the positions you selected. It also includes a toolbar that enables you to perform various functions.
G	Load Favorite icon	Center Right, within the Selected Items area	Enables you to select a previously saved group of positions to load into the workbook.
H	Save Favorite icon	Center Right, within the Selected Items area	Enables you to save the positions you have selected as group. The next time you build a new workbook, you can select and load that group rather than choosing the same positions individually again.
I	Add and Add All icons	Center	Enables you to add positions that are selected in the Available Items area. To add all positions in the Available Items area, click the Add All icon.
J	Remove and Remove All icons	Center	Enables you to remove positions in the Selected Items area. To remove all positions in the Selected Items area, click the Remove All icon.
K	Wizard Navigation icon	Bottom Right	Enables you to navigate from one wizard page to another. Once you have made the selections for the workbook and clicked Finish , the workbook builds with the selected positions.
L	Synchronize Hierarchies check box	Bottom Left	When selected, the shown levels within the Selected Items area match the ones in the Available Items area. This is selected by default.

Accessing Online Help

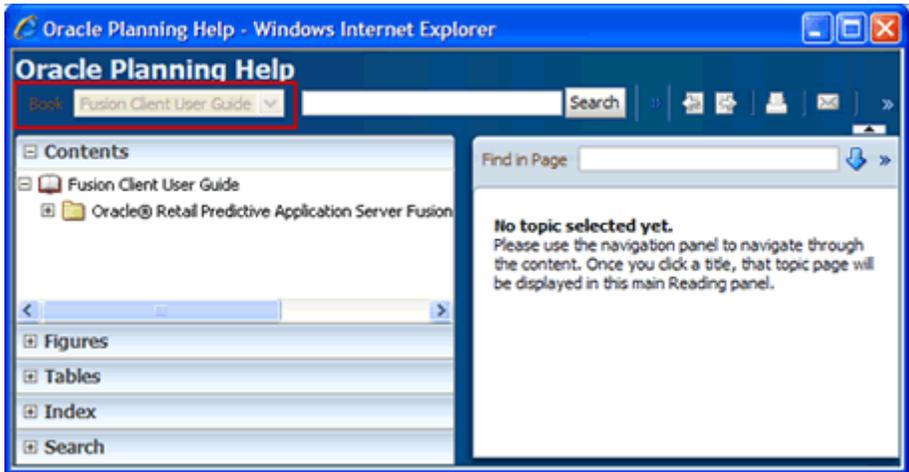
To access online help, click **Help** in the global header.

Figure 2–14 Online Help Link



The online help opens in a new browser window. If more than one help set is available, you can choose which one to view by selecting it from the Book drop-down field. If there is only one help set, the field is shaded.

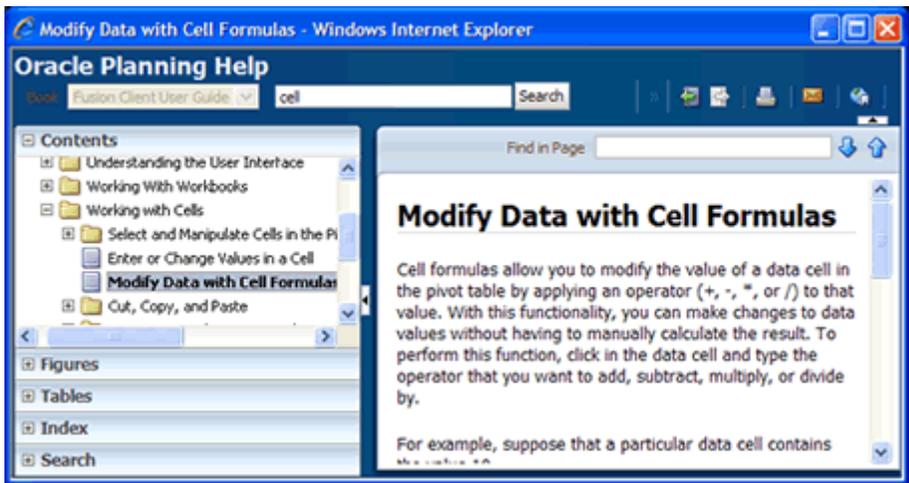
Figure 2–15 Online Help Window, Book Field



Contents

The contents of the online help set are shown in the Contents section of the navigation bar on the left. Click the **Expand** icon next to the folders to drill down to the help topics. Click the page that you want to view. It appears in the content area.

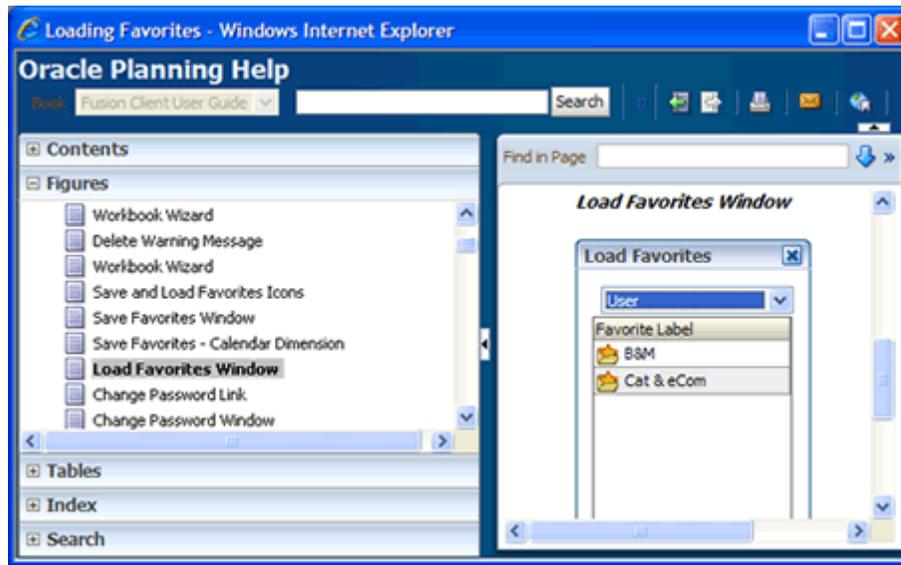
Figure 2–16 Help Topic Within Online Help



Figures and Tables

To view a list of figures or tables, click the **Figures** or **Tables** bar.

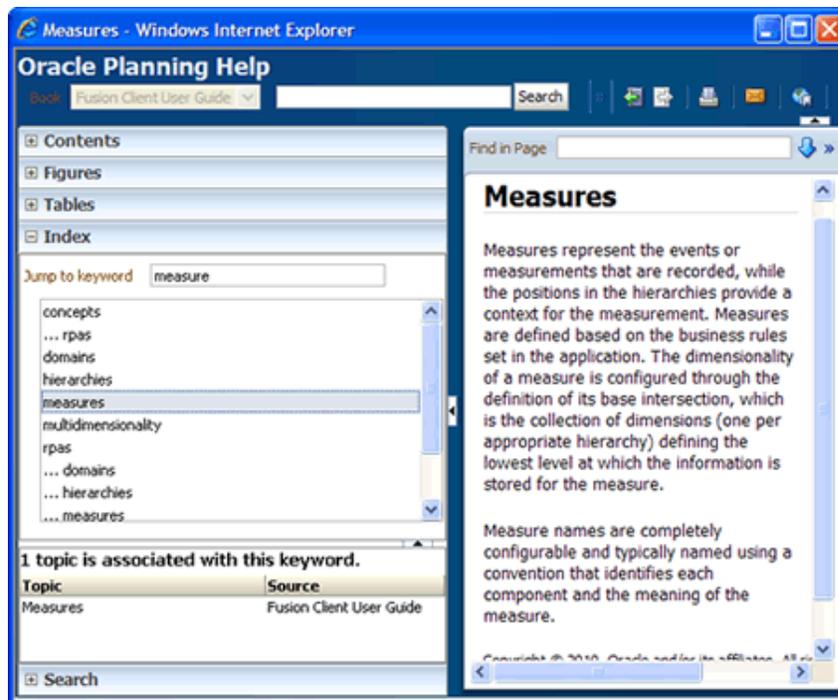
Figure 2–17 Online Help: Figure List



Index

In the Index, you can enter keywords to find topics that contain this keyword. Click the topic name to see the page.

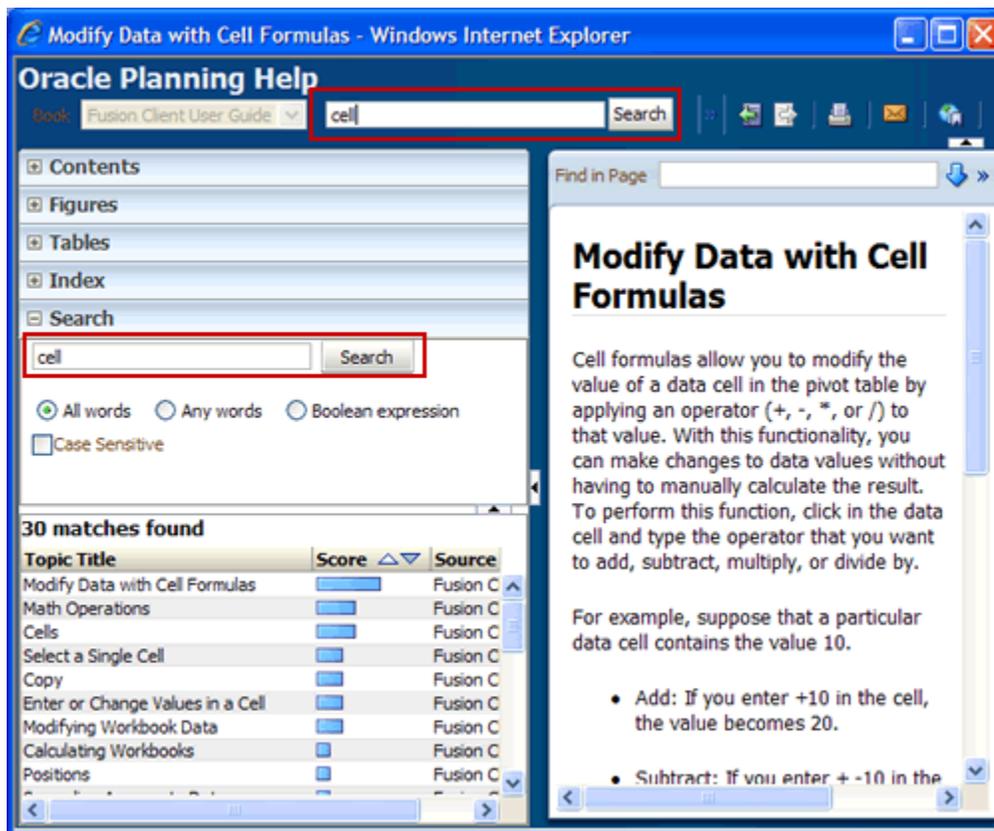
Figure 2–18 Online Help: Index



Search

Use the Search field at the top of the online help window or the Search section within the navigation bar to search for words or phrases.

Figure 2–19 Online Help: Search



Logging Out of RDF

To log out of the application, click **Logout** in the global header.

Figure 2–20 Logout Link



Note: You may be prompted to save and commit the workbook data before logging out of the application.

Oracle Retail Demand Forecasting Methods

This chapter discusses the forecasting methods used in Oracle Retail Demand Forecasting in detail.

Forecasting Techniques Used in RDF

RDF uses a variety of predictive techniques to generate forecasts of demand. The technical methods used are driven by the goal to provide the most accurate forecasts possible in an automatic and efficient manner. These methods have been analyzed, optimized, and refined over years of research on retail-specific data.

The primary techniques RDF uses include:

- [Exponential Smoothing](#)
- [Regression Analysis](#)
- [Bayesian Analysis](#)
- [Prediction Intervals](#)
- [Automatic Method Selection](#)
- [Source Level Forecasting](#)
- [Promotional Forecasting](#)

Exponential Smoothing

Exponential smoothing models fit basic features of the demand pattern such as level, trend, and seasonality, and project these into the future. These models provide computational benefits and have been chosen for their ability to handle different types of time series, including short and/or noisy series that are characteristic of retail sales. They are smoothing models because they use weighted averages on historic data. They are exponential smoothing models because the weighting uses decays at an exponential rate. That is, more recent data is weighted more heavily than the past.

Regression Analysis

Regression analysis is another standard technique used in prediction. Regression uses a least-squares estimator to fit a model of predictor variables to another set of target variables. Seasonal Regression is an Oracle Retail specific extension of this procedure for use in seasonal models with between one and two years of history. Causal Forecasting uses stepwise regression to determine which Promo Effects are significant.

Bayesian Analysis

Bayesian analysis considers *a priori* information as a starting point in development of a prediction. Bayesian forecasting, as developed by Oracle Retail, uses a sales plan as the starting point that is adjusted based on observed data. This method fills a gap in standard time series forecasting when new, short lifecycle or products with significant lifecycles are being forecast.

Prediction Intervals

Predictions from these various models gives the estimated mean outcome. By using standard statistical distributional assumptions, RDF develops measures of uncertainty associated with forecast point estimates from these models. While this is of key concern for various optimization solutions of the forecast, the technical details are beyond the scope of this document. For further details on prediction interval calculations, refer to *Char&Yatfield, International Journal of Forecasting, March 1992*.

Automatic Method Selection

Providing multiple forecasting methods is only valuable if the appropriate model can be selected in an accurate and efficient manner. In order to make this feasible in a retail environment, Oracle Retail has developed a number of different meta-methods that can automatically select the best method among a number of competing models. Automatic Exponential Smoothing (AutoES) is an example of one such method that clients can select. The final selection between the competing models is made according to a performance criterion that involves a trade-off between the model's fit over the historic data and its complexity. A description of the competing models used within AutoES is described in "[Exponential Smoothing \(ES\) Forecasting Methods](#)" on page 3-9. In academia, this discipline is known as Information Theory and is used in the combination and selection of various competing models.

Source Level Forecasting

Sometimes it is difficult to capture seasonality, trend, or causal effects on the final-level (item/store) due to scarcity of the data. Also, time series are often too noisy at that level. To overcome these issues, RDF utilizes source-level forecasting. In source-level forecasting, data is aggregated first to a higher level across the product or location hierarchy (or both). Then the forecast is generated and proportionally spread down to the final-level. We have experimentally proven that source-level forecasting technique often improves the accuracy on the final-level.

Promotional Forecasting

In some instances, especially in retail, pure time series techniques are inadequate for forecasting demand. Instead of using only historic demand patterns to forecast future demand, additional causal or promotional factors are used to better explain past performance. With the help of a promotional calendar, an indication of when promotions is run in the future, these promotional forecasting techniques can better predict demand in the future.

Time Series (Statistical) Forecasting Methods

This section describes those techniques within RDF that generate forecasts directly from only a single time series. Generally the time series provided is past sales history for a given item/store that is used to predict what future demand might be. In actual

practice these algorithms have been and can be used to forecast a myriad of different data streams at any product/location level (shipment data at item/warehouse, financial data at dept./chain, and so on).

The following topics present fundamentals of the RDF statistical forecasting processes. Included is a discussion of the importance of confidence intervals and confidence limits, the time series methods used to generate forecasts, and how the best forecasting method is selected from a list of candidate models.

A wide variety of statistical forecasting techniques are available, ranging from very simple to very sophisticated. All of these methods attempt to best capture the statistical probability distribution previously discussed, and they do this by fitting quantitative models to statistical patterns from historical data. Put simply, the better the history of the variable being forecast, the stronger these statistical patterns are. Increased forecast accuracy depends on the strength of these patterns in relation to background irregularities.

RDF is able to use several time series methods to produce forecasts. Time series methods extrapolate features from the past (in this case, past sales data) to the future. The time series methods that the system offers include:

- [Average](#) (Simple Moving Average)
- [Bayesian Information Criterion \(BIC\)](#)
- [Exponential Smoothing \(ES\) Forecasting Methods](#) - Automatic Exponential Smoothing (AutoES)
 - [Simple Exponential Smoothing](#)
 - [Croston's Method](#)
 - [Simple/Intermittent Exponential Smoothing](#)
 - [Holt Exponential Smoothing](#)
 - [Multiplicative Winters Exponential Smoothing](#)
 - [Oracle Winters](#)
 - [Oracle Winters Decomposition](#)
 - [Winters Standard](#)
 - [Winters Responsive](#)
 - [Additive Winters Exponential Smoothing](#)
 - [Seasonal Exponential Smoothing \(SeasonalES\)](#)
 - [Seasonal Regression](#)

Why Use Statistical Forecasting?

The purpose of statistical forecasting is to make the process of predicting future events both objective and quantitative. Statistical forecasting utilizes information from the past (such as sales data) to predict what will happen in the future. Forecast accuracy depends on the degree to which a mathematical model can detect and extract statistical patterns from historic data. The most common statistical methodologies used are univariate. This means that they are based solely on the history of one variable, such as sales. Each forecast observation reflects a future value of the sole input variable. Statistical forecasting processes are relatively easy to implement, and the better the historical data, the better the resulting forecasts.

Businesses benefit greatly from the use of systematic statistical forecasting techniques that aim to accurately predict product demand, enabling these businesses to maintain sufficient product inventory levels. When inventory levels are optimized, lost sales due to product stock-outs are greatly reduced, as are the costs incurred by overstocking.

Average

A simple moving average forecast involves taking the average of the past n time periods and using that average as the forecast for all future time periods (where n is the length of fitting period). Simple moving average forecasts are frequently used in the system because they:

- Make few assumptions about the historical time series.
- Can be generated with little historical data.
- Are very fast to generate.

Typically, moving average forecasts are generated at the final forecast level (for example, item/store) and their results used to spread more sophisticated higher-level forecasts (for example, those generated with exponential smoothing).

A Simple Moving Average model assumes that historical data is too short or noisy to consider seasonal effects or local trend and is based on the level of the series. Since this model does not use a smoothing parameter to place added weight on more recent historic values, a Simple Moving Average model is not actually in the exponential smoothing family. However, it is an adequate model to use when low-level (final forecast) ratios are needed for RDF's spreading of high-level (aggregate) forecasts. That is, when aggregate forecasts can be calculated for long and less noisy aggregate time series, Simple Moving Average models provide an adequate (and computationally quick) forecast to determine the ratios needed for RDF spreading. User input in overriding the automatic training horizon further enhances the simple robustness of this model for base-level data.

Bayesian Information Criterion (BIC)

Within AutoES, the model that minimizes the Bayesian Information Criterion (BIC) is selected as the final model. The BIC criterion attempts to balance model complexity with goodness-of-fit over the historical data period (between history start date and forecast start date). The BIC criterion rewards a model for goodness-of-fit and penalizes a model for its complexity. The complexity penalty is necessary to avoid over fitting.

There are various equivalent versions of the Bayesian Information Criterion, but RDF minimizes the following:

$$BIC = s \cdot n^{k/2n}$$

where n is the number of periods in the available data history, k is the number of parameters to be estimated in the model (a measure of model complexity), and s is the root mean squared error computed with one-step-ahead forecast errors resulting from the fitted model (a measure of goodness-of-fit). Note that since each member of the model candidate list is actually a family of models, an optimization routine to select

optimal smoothing parameters is required to minimize s for each model form (that is, to select the best model).

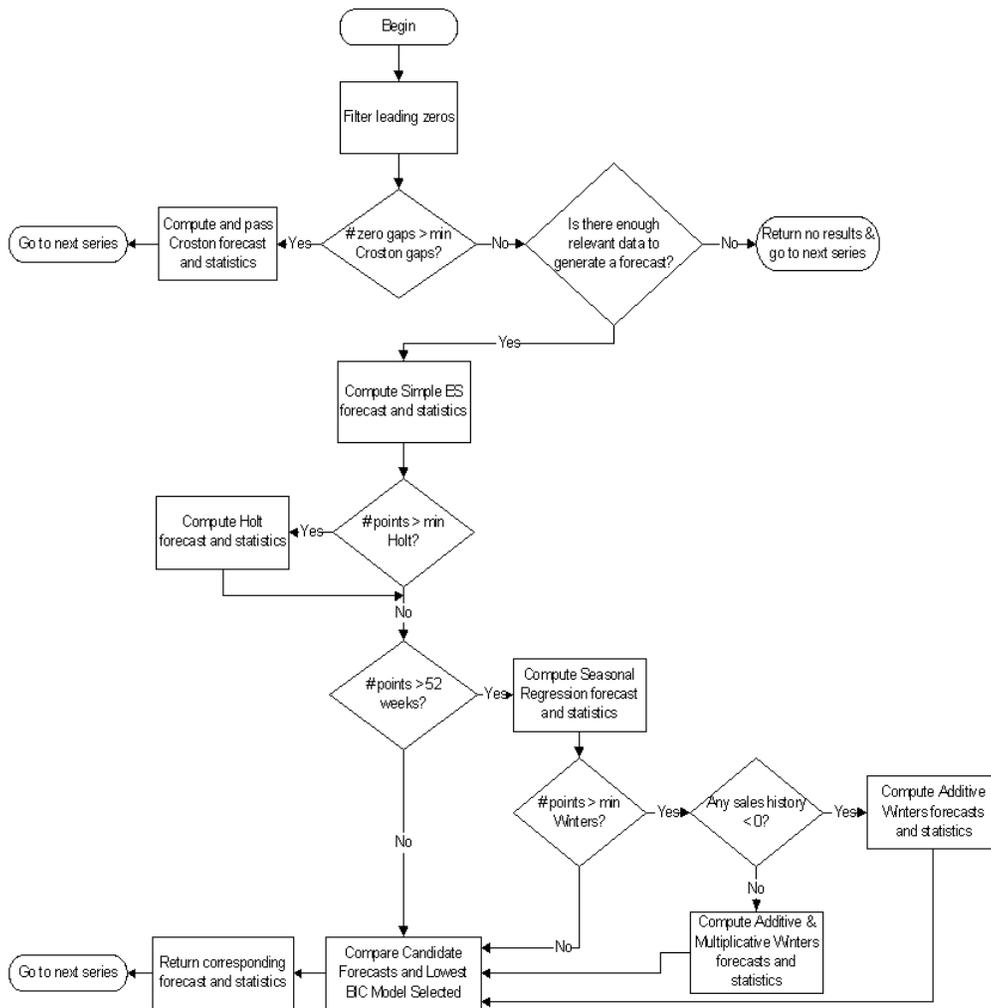
Within RDF, a few modifications to the standard selection criteria have been made. These include reducing the number of parameters the Winter's model is penalized by discounting seasonal indices that have little impact on the forecast (multiplicative indices close to one (1), additive indices close to zero (0)). These changes tend to favor the seasonal models to a slightly higher degree that improves the forecasts on retail data, especially for longer forecast horizons.

AutoES Flowchart

The following procedure outlines the processing routine steps that the system runs through to evaluate each time series set to forecast using the AutoES method. Refer to [Figure 3–1, "AutoES Flowchart"](#).

1. Filter all leading zeros in the input data that is within the training window.
2. Does the time series contain the minimum data points to qualify to forecast using the Croston's method? If yes, generate the forecast and statistics using the Croston's method and move on to the next time series. If no, move on to Step 3.
3. Does the time series contain enough relevant data to generate a forecast? If yes, generate a forecast and statistics using the SimpleES method and move on to Step 4. If no, do not forecast and go to the next time series.
4. Does the time series contain the minimum data points to qualify to forecast using the Holt method? If yes, generate a forecast and statistics using the Holt method and move on to Step 5. If no, move on to Step 9.
5. Does the time series contain more than 52 weeks of input data? If yes, generate a forecast and statistics using the Seasonal Regression method and move on to Step 6. If no, move on to Step 9.
6. Does the time series contain the minimum data points to qualify to forecast using Winters methods? If yes, move on to Step 7. If no, move on to Step 9.
7. Does the time series contain any data point with sales equal qualify to forecast using Additive Winters method? If yes, generate the forecast and statistics using the Additive Winters method and move on to Step 9. If no, move on to Step 8.
8. Does the time series qualify to forecast using the Multiplicative Winters method? If yes, generate the forecast and statistics using both the Additive Winters and Multiplicative Winters methods and move on to Step 9.
9. Compare all candidate forecasts using BIC Criterion.
10. Return the corresponding forecast and statistics for the system-selected forecast method and move on to the next time series.

Figure 3–1 AutoES Flowchart



Automatic Forecast Level Selection (AutoSource)

This section describes how the automatic forecast level selection (AutoSource) could help improve the accuracy of your forecasts.

In the system, one of the key elements to producing accurate forecasts is using the system's ability to aggregate and spread sales data and forecasts across the product and location hierarchies. Low selling or relatively new products can use aggregated data from similar products/locations at a higher level in the hierarchy, generate forecasts using this data, and then spread these higher level forecasts back down to provide more accurate forecasts. The difficulty comes in deciding which products/locations will benefit from this technique and from what level in the hierarchy these source-level forecasts should be spread.

The Automatic Forecast Level Selection feature of the system automates the selection of best aggregation level (forecast source-level) for each product/location combination. While providing invaluable information regarding the best aggregate level for source forecasts, the Automatic Forecast Level Selection process may be very CPU intensive. To solve this problem, the task of selecting best aggregation levels for product/location combinations is decomposed and processed piecemeal during times when the computer would normally be idle.

Identifying the best aggregation levels for sets of products and locations can be divided into a number of sub-problems:

- Forecasting
- Determining the best source-level forecast
- Status and scheduling

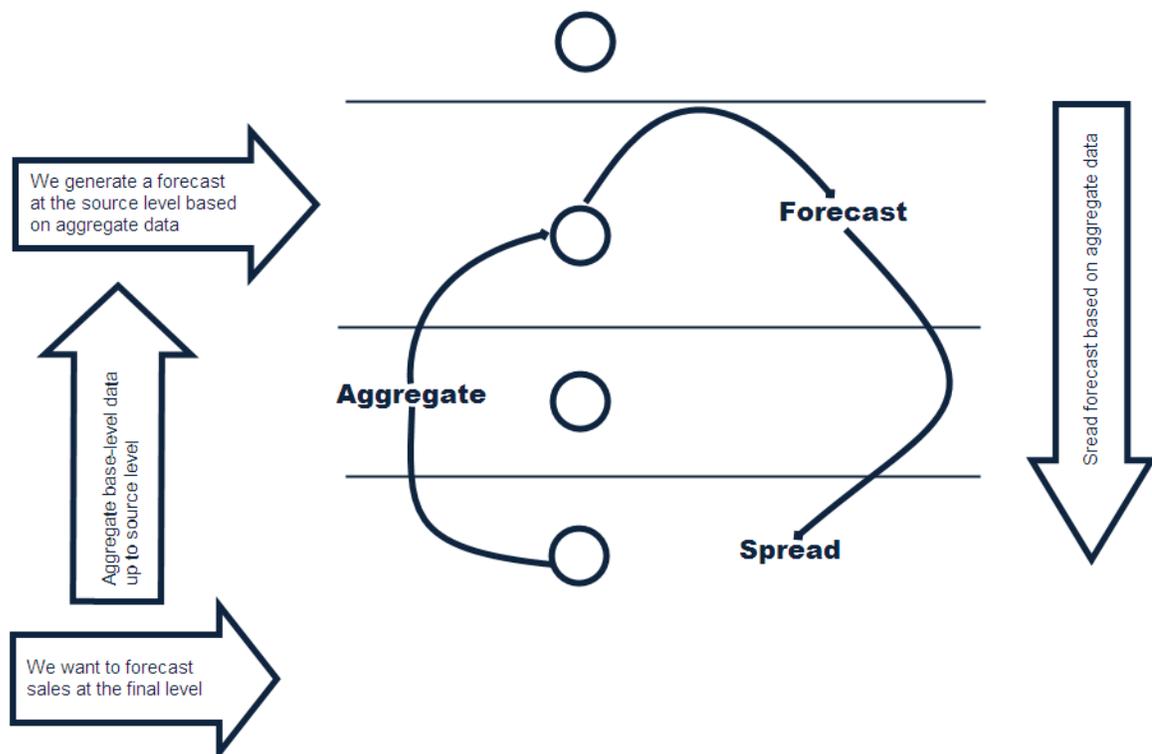
The Forecast Level Selection Process

The automatic source generation level selection subsystem selects the best source generation level for each product/location in a given final forecast level. In order to determine the best level, a final forecast is generated for each product/location using each candidate source generation level. As illustrated in [Figure 3–2](#), a final forecast is generated by:

1. Aggregating up from the base level to the source-level
2. Generating a source-level forecast
3. Spreading the source-level forecast down to the final-level

For example, assume base-level sales data is at the item/store level, the final forecast level is at the item/store level, and the candidate source generation level is at the style/store level. In this case, base-level sales data is aggregated from the item/store level up to the style/store level. A style/store forecast is generated, and the forecast data is spread back down to the item/store level. This forecast represents the final forecast.

Figure 3–2 Forecast Level Selection Process



Determining the Best Source Level Forecast

The selection of the best level is based on a train-test approach. In this process, historical data is used to generate a forecast for a test period for which actual sales data already exists. The forecast, generated over the train period, can be compared to the actual sales figures in the test period to calculate the Percent Absolute Error (PAE) between the two.

A final-level forecast is generated for each product/location combination using each potential source generation level. Each time a source-level forecast is generated, a PAE is calculated for that level. If that PAE is better than the current best PAE (corresponding to the current best source generation level), the source generation level that generated that better PAE becomes the new best level.

Status and Scheduling

Identifying the best aggregation level for a given set of products and locations may take a significant amount of time (that is, an amount of time that is greater than the duration of the computer's shortest idle period). This task, however, can be partitioned meaning that the problem of selecting the best aggregation levels can be decomposed into smaller sub-problems to be solved piecemeal during times when the computer would normally be idle.

For each product/location combination at the final forecast level, the problem consists of:

- Generating forecasts at each unique aggregation level
- Using the train-test approach to evaluate the percent absolute error statistics for each

One or more of these subtasks is performed during each period that the computer is idle. The best aggregation status keeps track of which sub-problems have been performed and which sub-problems remain. In this way, when the best aggregation procedure is run, the procedure knows what the next sub-problem is.

Best aggregation level procedures are run during idle computer periods. The scheduling of the Automatic Forecast Level Selection process (AutoSource) must be integrated with the schedules of other machine processes. In general, you should select a schedule so that source generation-level selection does not conflict with other activities. The following is an example of a typical schedule for the Automatic Forecast Level Selection process: Monday through Thursday, the selection process starts at midnight and runs for eight hours. On Friday and Saturday, the process is allowed to run for 20 hours. Sunday is reserved for generating forecasts.

Using the System-Selected Forecast Level

You have the option of accepting the system-generated source-level selection or manually selecting a different source-level to be used. The value for the source forecast level can be manipulated in the Final Level view of the Forecast Maintenance workbook. For each product/location combination, the best source forecast level identified by RDF appears in the Optimal source-level measure on this view. You can enable the use of this level by placing a check mark in the Pick Optimal Level measure for that product/location. The absence of a check mark in this measure causes the system to default to the Default Source level or the Source Level Override value if this has been set by you.

Exponential Smoothing (ES) Forecasting Methods

The primary process by which RDF automatically fits an exponential smoothing model to a time series is called Automatic Exponential Smoothing (AutoES). When AutoES forecasting is chosen in RDF, a collection of candidate models is initially considered. The models in the candidate list include:

- [Simple Exponential Smoothing](#)
- [Croston's Method](#)
- [Simple/Intermittent Exponential Smoothing](#)
- [Holt Exponential Smoothing](#)
- [Multiplicative Winters Exponential Smoothing](#)
 - [Oracle Winters](#)
 - [Oracle Winters Decomposition](#)
 - [Winters Standard](#)
 - [Winters Responsive](#)
- [Additive Winters Exponential Smoothing](#)
- [Seasonal Exponential Smoothing \(SeasonalES\)](#)
- [Seasonal Regression](#)

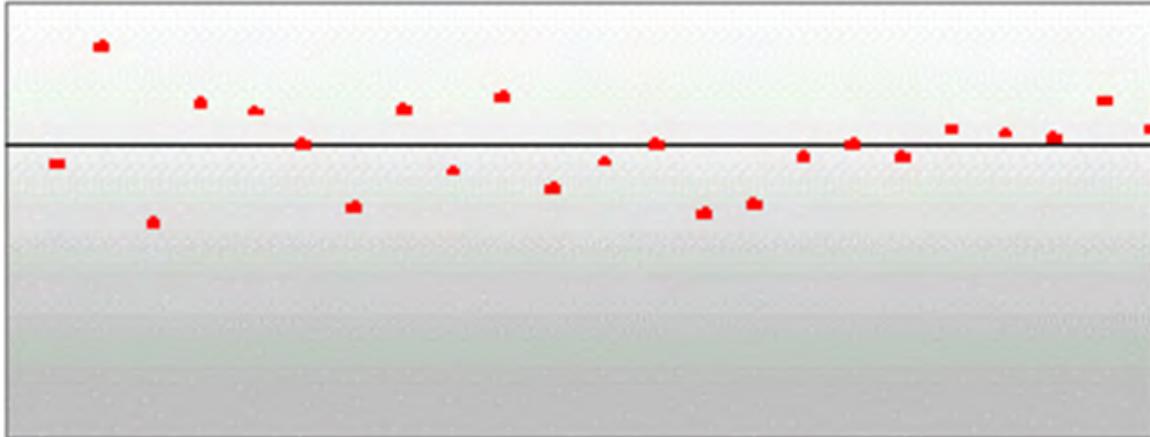
These models include level information, level and trend information, and level, trend and seasonality information, respectively. The optimal smoothing parameters for each model form are determined automatically (that is, greater smoothing is applied to noisier data). The final selection between the resulting models is made according to a performance criterion that involves a trade-off between the model's fit over the historic data and its complexity.

The amount of available historic information can affect the complexity of the model that can be fit. For example, fitting a seasonal model would not be appropriate without a complete year of historic data. In fact, one prefers to see each season occur multiple times. For a particular series, even if the amount of available history allows one to fit a complex model (that is, one with seasonal components), the resulting model is not necessarily superior to a simpler model. If a simpler model (for example, a model with only a level component or level and trend components) fits as well as a seasonal model, the AutoES forecasting process finds the simpler model to be preferable. In such a case, the simpler model captures the basic features supported by the data without over fitting and therefore generally projects better forecasts.

Simple Exponential Smoothing

Simple Exponential Smoothing does not consider seasonality or trend features in the demand data (if they exist). It is the simplest model of the exponential smoothing family, yet still adequate for many types of RDF demand data. Forecasts for short horizons can be estimated with Simple Exponential Smoothing when less than a year of historic demand data is available and acts-like associations are not assigned in RDF.

[Figure 3–3](#) is an example of a forecast in which data seems to be un-trended and un-seasonal; note the flat appearance of the forecast.

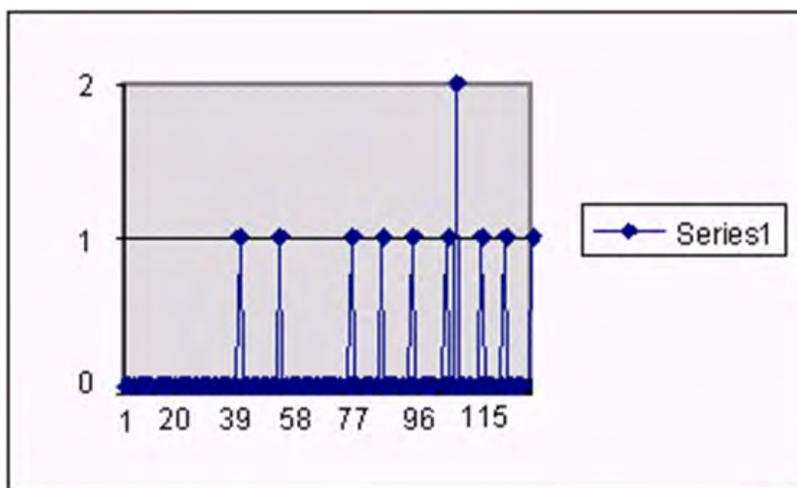
Figure 3–3 Simple Exponential Smoothing**Croston's Method**

Croston's Method is used when the input series contains a large number of zero data points (that is, intermittent demand data). The method involves splitting the original time series into two new series:

- Magnitude series
- Frequency series

The magnitude series contains all the non-zero data points, while the frequency series consists of the time intervals between consecutive non-zero data points. A Simple Exponential Smoothing model is then applied to each of these newly created series to forecast a magnitude level as well as a frequency level. The ratio of the magnitude estimate over the frequency estimate is the forecast level reported for the original series.

Figure 3–4 shows a sales history of data where the demand for a given period is often zero.

Figure 3–4 Croston's Method

Simple/Intermittent Exponential Smoothing

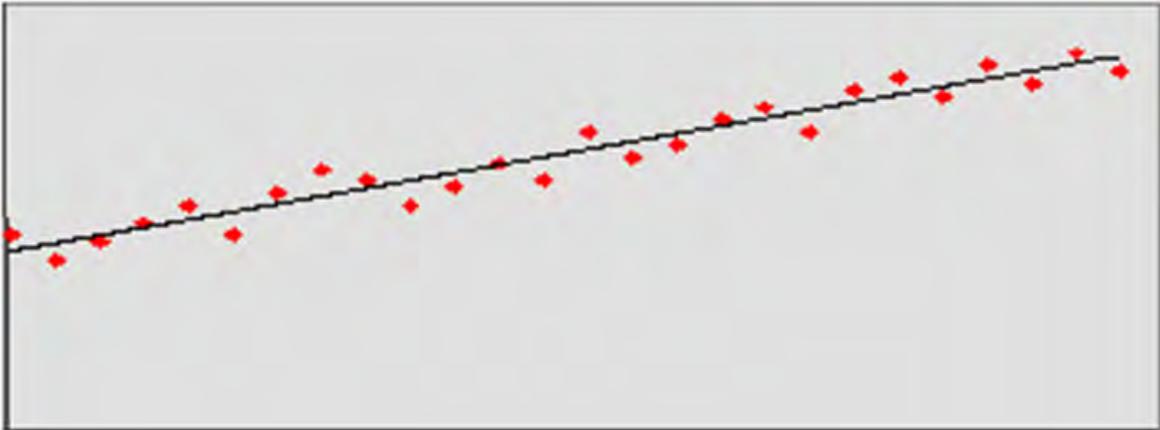
This method is a combination of the SimpleES and Croston's (IntermittentES) methods. The SimpleES model is applied to the time series unless a large number of zero data points are present. In this case, the Croston's model is applied.

Holt Exponential Smoothing

Holt exponential smoothing treats data as linearly trended but non-seasonal. The Holt model provides forecast point estimates by combining an estimated trend (for the forecast horizon - h) and the smoothed level at the end of the series. RDF uses a damped Holt model that decays the trend component so that it disappears over the first few weeks. This improves forecasts created using Holt over longer forecast horizons.

When this forecasting method is selected, the forecasts are seen as trending either up or down, as shown in [Figure 3-5](#).

Figure 3-5 Holt Exponential Smoothing



Multiplicative Winters Exponential Smoothing

RDF now offers a choice of four different forecasting approaches within the Multiplicative Seasonal Forecasting Method.

The Winter's forecasting approaches, including the current seasonal approach, are:

- Oracle Winters (current or default approach)
- Oracle Winters Decomposition
- Winters Standard
- Winters Responsive

All four approaches use the same calculation for arriving at forecasts.

Overall, a forecast point estimate is evaluated as:

$$\hat{Y}_t(h) = \left(L_t + \left[\sum_{i=1}^h \varphi^i \right] T_t \right) \hat{S}_t(h),$$

a function of level, trend, seasonality, and trend dampening factor.

Winters Standard and Winters Responsive both use a basic Winter's method:

The Level at the end of the series (time t) is:

$$L_t = \alpha \frac{Y_t}{S_{t-p}} + (1 - \alpha)(L_{t-1} + T_{t-1})$$

The Trend at the end of the series (time t) is:

$$T_t = \gamma(L_t - L_{t-1}) + (1 - \gamma)T_{t-1}$$

The Seasonal Index for the time series (applied to the forecast horizon) is:

$$S_t = \delta \frac{Y_t}{L_t} + (1 - \delta)S_{t-p}$$

The difference between the two approaches is how these values are optimized within the optimization routine.

Oracle Winters and Oracle Winters Decomposition use a Winters-based decomposition approach to update the level, trend, and Seasonal Indexes. Each of these methods calculates initial seasonal indices from a baseline Holt forecast. Seasonal indices, level, and trend are then updated in separate stages, using Winter's model as a basis for the updates. The difference between the two approaches is which parameters are optimized and in which stages seasonal indices are updated.

Oracle Winters

Oracle Winters is the current seasonal forecasting approach, which uses a combination of Winters approach and decomposition. Decomposition allows level and trend to be optimized independently while maintaining a seasonal curve.

From sufficient data, RDF extracts seasonal indexes that are assumed to have multiplicative effects on the de-seasonalized series.

Note: The component describing the de-seasonalized values, which is multiplied by the seasonal index:

$$\hat{S}_t(h)$$

is the Holt model previously described.

In this case, three parameters are used to control smoothing of the components of level, trend, and seasonality.

Oracle Winters Decomposition

Like the Oracle Winters seasonal forecasting approach, Oracle Winters Decomposition uses a combination of Winters approach and decomposition.

From sufficient data, RDF extracts seasonal indexes that are assumed to have multiplicative effects on the de-seasonalized series.

Note: The component describing the de-seasonalized values, which is multiplied by the seasonal index:

$$\hat{S}_t(h)$$

is the Holt model previously described.

The key difference between Oracle Winters and Oracle Decomposition is that the Seasonal smoothing parameter is also optimized, and the calculation of Seasonal Indices and de-seasonalizing of data is done with in the optimization routine for Oracle Winters Decomposition.

Also, the optimization routine with in Oracle Winters Decomposition, while minimizing forecast errors, tends to weigh more recent Forecast errors heavier.

As with Oracle Winters, three parameters are used to control smoothing of level, trend, and seasonality.

Key differences between Oracle Winters and Oracle Winters Decomposition include:

- Seasonal smoothing factor is optimized.
- The optimization places more importance on recent sales than historical sales.
- Seasonal indices are recalculated based on optimization.

When this forecasting approach is selected, the forecasts tend to be more responsive to recent changes in levels while maintaining a substantial seasonal component.

Winters Standard

The Winters Standard uses a standard Winters model, with no additional calculations. All three smoothing parameters are optimized here.

Initial Seasonal Indices are set to one. Initial level and trend are calculated from average selling levels over initial periods. The optimization routine chooses optimal level, trend, and smoothing parameters that minimize Forecast Errors in the fit region. All forecast errors within the fit region are weighed equally.

Seasonal Indices, level, and trend are calculated based on the optimal values of level, trend, and Seasonal Smoothing parameters:

$$\alpha, \gamma, \delta$$

When this forecasting approach is selected, the forecasts tend to be responsive to changes in levels. In cases where there is no obvious correlation in seasonality from period to period or where there are significant changes in level, these forecasts may generate smoother, less seasonal forecasts.

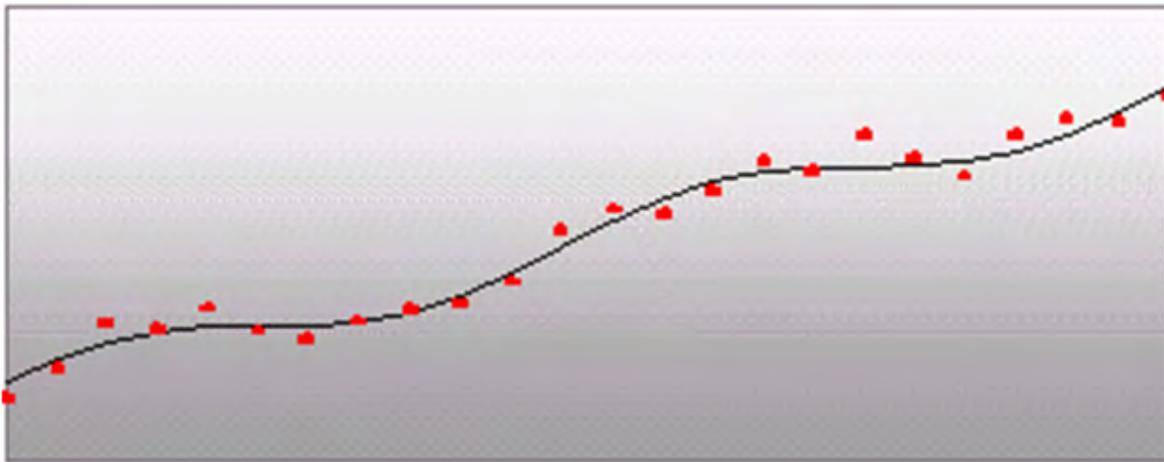
Winters Responsive

Winters Responsive is similar to Winters Standard except the Optimization routine, while minimizing forecast errors, weighs more recent Forecast errors heavier. Winters Responsive, like Standard Winters, optimizes all three smoothing parameters.

When this forecasting approach is selected, the forecasts tend to be more responsive to recent changes in levels. In cases where there are significant shift in levels, resulting in high responsiveness to levels, these forecasts tend to be smoother, with less seasonality.

When the Multiplicative Seasonal forecasting method is selected, the forecasts tend to look squiggly, as shown in [Figure 3–6](#).

Figure 3–6 *Multiplicative Winters Exponential Smoothing*



Additive Winters Exponential Smoothing

RDF now offers a choice of four different Forecasting approaches within the Additive Seasonal Forecasting Method.

The Winters forecasting approaches, including the current seasonal approach, are:

- Oracle Winters (current or default approach)
- Oracle Winters Decomposition
- Winters Standard
- Winters Responsive

All four approaches use the same calculation for arriving at forecasts.

Overall, a forecast point estimate is evaluated as:

$$\hat{Y}_t(h) = L_t + \left[\sum_{i=1}^h \varphi^i \right] T_t + \hat{S}_t(h)$$

a function of level, trend, seasonality and trend dampening factor.

Winters Standard and Winters Responsive both use a basic Winter's method.

The Level at the end of the series (time t) is:

$$T_t = \gamma(L_t - L_{t-1}) + (1 - \gamma)T_{t-1}$$

and the Trend at the end of the series (time t) is:

$$T_t = \gamma(L_t - L_{t-1}) + (1 - \gamma)T_{t-1}$$

and the Seasonal Index for the time series (applied to the forecast horizon) is:

$$S_t = \delta(Y_t - L_t) + (1 - \delta)S_{t-p}$$

The difference between the two approaches is how these values are optimized within the optimization routine.

Oracle Winters and Oracle Winters Decomposition use a Winters-based decomposition approach to update the level, trend, and seasonal indexes.

Refer to Multiplicative Winters Exponential Smoothing in this document for a description of each of the forecasting approaches.

Seasonal Exponential Smoothing (SeasonalES)

A combination of several seasonal methods. This method is generally used for known seasonal items or forecasting for long horizons. This method lets the Multiplicative Seasonal and Additive Seasonal models compete and picks the one with the better fit. If less than two years of data is available, a Seasonal Regression model is used. If there is too little data to create a seasonal forecast (in general, less than 52 weeks), the system selects from the Simple ES, Trend ES, and Intermittent ES methods.

Seasonal Regression

A common benchmark in seasonal forecasting methods is sales last year. A sales last year forecast is based entirely on sales from the same time period of last year. Forecasting using only sales last year involves simple calculations and often outperforms other more sophisticated seasonal forecasting models. This method performs best when dealing with highly seasonal sales data with a relatively short sales history.

The seasonal models used in earlier releases of RDF (Additive and Multiplicative Winters) were designed to determine seasonality. However, they were not designed to work with sales histories of shorter than two years. Because sales histories of longer than two years are often difficult to obtain, many retail environments need a seasonal forecast that can accommodate sales data histories of between one and two years. In addition, the Additive and Multiplicative Winters models search for short-term trends and have difficulties with trends occurring inside the seasonal indices themselves.

The current RDF Seasonal Regression forecasting model is designed to address these needs.

The Seasonal Regression Model uses simple linear regression with last year's data as the predictor variable and this year's sales as the target variable. The system determines the multiplicative and additive weights that best fit the data on hand. When optimizing the Seasonal Regression Model, the sales last year forecast is inherently considered, and it will automatically be used if it is the model that best fits the data. If there have been significant shifts in the level of sales from one year to the next, the model learns that shift and appropriately weigh last year's data (keeping the same shape from last year, but adjusting its scale).

As with other seasonal models, you can forecast demand for products with insufficient sales histories using this method if:

- You paste in fake history as needed, providing a seasonal profile for the previous year.
- You also forecast for a source-level (with the same seasonality profile as the forecast item and with more than one year of history) using seasonal regression and spread these forecast results down to the member products.

The Seasonal Regression Model is included in the AutoES family of forecasting models and is thus a candidate model that is selected if it best fits the data.

This method captures the trend of a series through the slope of the regression line while the series shifted by a cycle provides its seasonal profile. Using this method, the resulting forecast for the original series is calculated. The regression method provides a much better forecast of the series than was possible with the other exponential smoothing algorithms.

Based on the assumptions of the model that this method is trying to describe, versus the noisy data it is likely to receive, several exceptions to this regression technique are caught and corrected. First, since it is logically impossible to receive a negative value for the slope (such a value suggesting an inverse seasonality), whenever a negative slope is detected, the regression is rerun with the intercept fixed to zero. This guarantees that a positive slope is calculated and thus a more logical forecast is given.

The second noise-driven concession is to check the slope to determine if it is either too slight or too great. If this is the case, the method rejects itself out of hand and allows one of the other competing methods to provide the forecast.

Profile-Based Forecasting

The Profile-based forecasting method generates a forecast based on a seasonal profile. The profile may be loaded, manually entered, or generated by Curve. It can also be copied from another profile and adjusted.

Forecast Method

The Profile-based forecasting method proceeds as follows:

1. The historical data and the profile are loaded.
2. The data is de-seasonalized using the profile and then fed to Simple method.
3. The alpha is capped by 0.5.
4. The Simple forecast is re-seasonalized using the profiles.

Profile-Based Method and New Items

The Profile-based forecasting method can be successfully used to forecast new items. In order to do that, we need to have a profile (which can be copied from an item that shares the same seasonality) and a number that specifies the de-seasonalized demand (DD value). The forecast is calculated using the DD value multiplied by the profile. The confidence interval is set to 1/3 of the DD value.

If the DD value is used to forecast, the history (if it exists) of the product is ignored. Once we have enough history (number of data points exceed a global parameter), the forecast stops using the DD value, and it defaults to the normal Profile Based method.

Bayesian Forecasting

The Bayesian Forecasting method is based on combining historic sales data with sales plan data. It is especially effective for new products with little or no historic sales data.

Your sales plan can incorporate expert knowledge in two ways — shape and scale.

- Shape is the selling profile or lifecycle that can be derived from a sales plan. For example, the shape for certain fashion items might show sales ramping up quickly for the first four weeks and then trailing off to nothing over the next eight weeks.
- Scale, or magnitude, of a sales plan is the total quantity expected to be sold over the plan's duration.

Bayesian Forecasting assumes that the shape that sales takes is known, but the scale is uncertain. In Bayesian Forecasting, when no sales history is available, the sales forecast figures are equal to the sales plan figures. At this point, there is no reason to mistrust the sales plan. As point-of-sale data becomes available, the forecast is adjusted and the scale becomes a weighted average between the initial plan's scale and the scale reflected by known sales history. Confidence in the sales plan is controlled by the amount of sales data on hand and a Bayesian sensitivity constant (Bayesian Alpha), which you can set between zero and infinity.

Unlike standard time series forecasting, which requires only sales history to produce a forecast, Bayesian Forecasting requires a sales plan and sales history (if available). Because of this difference, Bayesian Forecasting is not included in AutoES. You must select it manually as a forecasting method in Forecast Administration or Forecast Maintenance.

Obtaining accurate short life-cycle product forecasts is very difficult, and standard statistical time series forecasting models frequently do not offer an adequate solution for many retailers. The following are major problems in automatically developing these forecasts:

- The lack of substantial sales history for a product (which especially makes obtaining seasonal forecasts very difficult).
- The difficulty of automatically matching a new product to a previous product or profile.
- The inability to include planners' intuition into a forecasting model. For example, the overall sales level of the product, how quickly the product takes off, how the product's sales is affected by planned promotions.

Using a Bayesian approach, a short life-cycle forecasting algorithm was developed that begins with a product's seasonal sales plan (that is developed externally to the system by the planner).

As sales information arrives during the first few days or weeks of the season, the model generates a forecast by merging the information contained in the sales plan with the information contained in the initial sales data. These forecast updates can be critical to a company's success and can be used to increase or cancel vendor orders.

As forecasting consultants and software providers, Oracle Retail assists clients in obtaining good forecasts for future demands for their products based upon historical sales data and available causal information. Depending on the information available, Oracle Retail's software supports various forms of exponential smoothing and regression-based forecasting. Frequently, clients already have some expectations of future demands in the form of sales plans. Depending on the quality of their plans, they can provide very useful information to a forecasting algorithm (especially when only limited historical sales data is available). A forecasting algorithm was developed that merges a customer's sales plans with any available historical sales in a Bayesian fashion (that is, it uses new information to update or revise an existing set of probabilities).

Sales Plans versus Historic Data

In most retail situations, clients are interested in obtaining good product forecasts automatically with little or no human intervention. For stable products with years of historic sales data, our time series approaches (Simple, Holt, Winters, Regression based Causal, and so on) produce adequate results. The problem arises when attempting to forecast products with little or no history. In such instances, expert knowledge (generally in the form of sales plans) is required. Given that both sales plans and time series forecasts are available, an obvious question exists: When should the transition from sales plan to time series forecasting occur? In answering that question (in a particular scenario), suppose that we have determined that 13 weeks of history is the transition point. Does that mean that at 12 weeks the time series results are irrelevant and that at 14 weeks the sales plan has no value? Our intuition tells us that instead of a hard-edge boundary existing, there is actually a steady continuum where the benefits from the sales plan decrease as we gather more historic sales data. This was the motivation for developing an approach that would combine the two forecasts in a reasonable manner.

Bayesian Algorithm

Bayesian forecasting is primarily designed for product/location positions for which a plan exists. These product/locations can be, but are not limited to: items new in the assortment, fashion items, and so on. The following guidelines should be followed:

1. No more than one plan should exist for a given product/location position. If multiple plans are to be set up for different time periods, the domain should be set up with different forecasting levels for each time period of interest.
2. Any time period with non-zero Actuals for a given product/location position should have a corresponding plan component. Otherwise, the system assumes a plan exists and equals zero and acts accordingly.
3. The time period of interest for the Bayesian algorithm starts with the first non-zero value of the plan or the history start date (whichever is more recent), and ends at the end of the forecast horizon.

LoadPlan Forecasting Method

LoadPlan Forecasting Method copies the measure that was specified as Data Plan in the [Forecast Administration Workbook](#) into the Forecast measure. This method does

not generate confidence and cumulative intervals when it is the final level method and no source level is specified. This method only generates confidence and cumulative intervals when a source level is specified and the source and final levels are the same.

Copy Forecasting Method

Copy Forecasting Method copies the measure that was specified as Forecast Data Source in the [Forecast Administration Workbook](#) into the Forecast measure. This method does not generate confidence and cumulative intervals.

Causal (Promotional) Forecasting Method

Causal, or promotional, forecasting requires three input streams:

- Time Series Data
- Historical Promotional Calendar
- Future Promotional Calendar

Promote decomposes the problem of promotional forecasting into two subtasks:

- Estimating the effect that promotions have on demand
- Forecasting baseline (that is no promotions) demand

To accomplish the first task, a stepwise regression sub-routine is used. This routine takes a time series and a collection of promotional variables and determines which variables are most relevant and what effect those relevant variables have on the series. Thus, the output from the algorithm is a selection of promotional variables and the effects of those variables on the series. In the second step, the time series is de-causalized using the promo effects. Then, AutoES is used to calculate the baseline demand. Once we know the effects and we have the baseline demand, we can generate a promotional forecast by applying the effects wherever the promotion is active in the future.

It should be noted that just because promotional forecasting is selected, it does not necessarily imply that a promotional forecast results. In some instances, no promotional variables are found to be statistically significant. In these cases, the forecast ends up equivalent to standard time series forecasts.

If you want to force certain promotional variables into the model, this can be managed through forecasting maintenance parameters.

The Oracle Retail experience in promotional forecasting has led us to believe that there are a few requirements that are necessary to successfully forecast retail promotions:

- Baseline forecasts need to consider seasonality; otherwise normal seasonal demand is attributed to promotional effects.
- Promotional Effects need to be able to be analyzed at higher levels in the retail product and location hierarchies. This produces cleaner signals and alleviates issues involved in forecasting new items and new stores and issues involving data sparsity.
- Users need to be aware that the forecasting models cannot tell the difference between causal effects and correlated effects. What this means is that users should be wary of promotional effects attributed to an event that occurs at the same time every year. The system cannot distinguish between the promotional effect and the normal seasonality of the product. The same can be said for any two events that

always occur at the same time. The combined effect is most likely attributed to one or the other event.

The Causal Forecasting Algorithm

For purposes of understanding the algorithm, a promotional variable is defined as a causal event that is only active for certain series at certain points in time. It is assumed that these events are entered into the system independently of the forecasting algorithm and that a vehicle exists for determining which promotional variables are relevant to a particular series.

The core of the causal forecasting algorithm uses a stepwise regression sub-routine. This sub-routine takes a time series and a collection of promotional variables, and it determines which variables are most relevant and what effect those relevant variables have on the series. Thus, the output from the algorithm is a selection of promotional variables and the effects of those variables on the series.

Promo Effect types define how Promo Effects are treated in the causal model-fitting process (which includes a call to the lower level regression engine) and the forecast generation process where the model is used to extend the forecast over the forecast horizon. Refer to [Chapter 13, "Promote \(Promotional Forecasting\),"](#) for information on each of the Promo Effect types:

- Automatic
- Force In The Model
- Disabled
- Override All
- Override Future
- Override Higher Level

Causal Forecasting Process

The causal forecasting algorithm itself lies in the AutoES binary code, and it executes in the following manner.

1. The binary reads the history of the time series.
2. The binary reads the type of each promotional variable into the system.
3. The binary reads in all the promotional variables that apply to the series.
4. The binary creates an internal promotional variable to allow the modeling of trend.
5. Promotional variables, internal promotional variables, promotional variable types, and the series itself are passed to the stepwise regression routine, with the historic data serving as the dependent variables.
6. If the regression finds no significant promotional variables, the casual method is considered to have failed to fit. In this case, the standard time series methods are used to generate a forecast and skips to step 14.
7. The fit at time t , $fit(t)$, is defined in terms of b_0 , the intercept of the regression, b_1 , the effect corresponding to promotional variable i , and $p_i(t)$, is the value of promo variable I , in time t as:

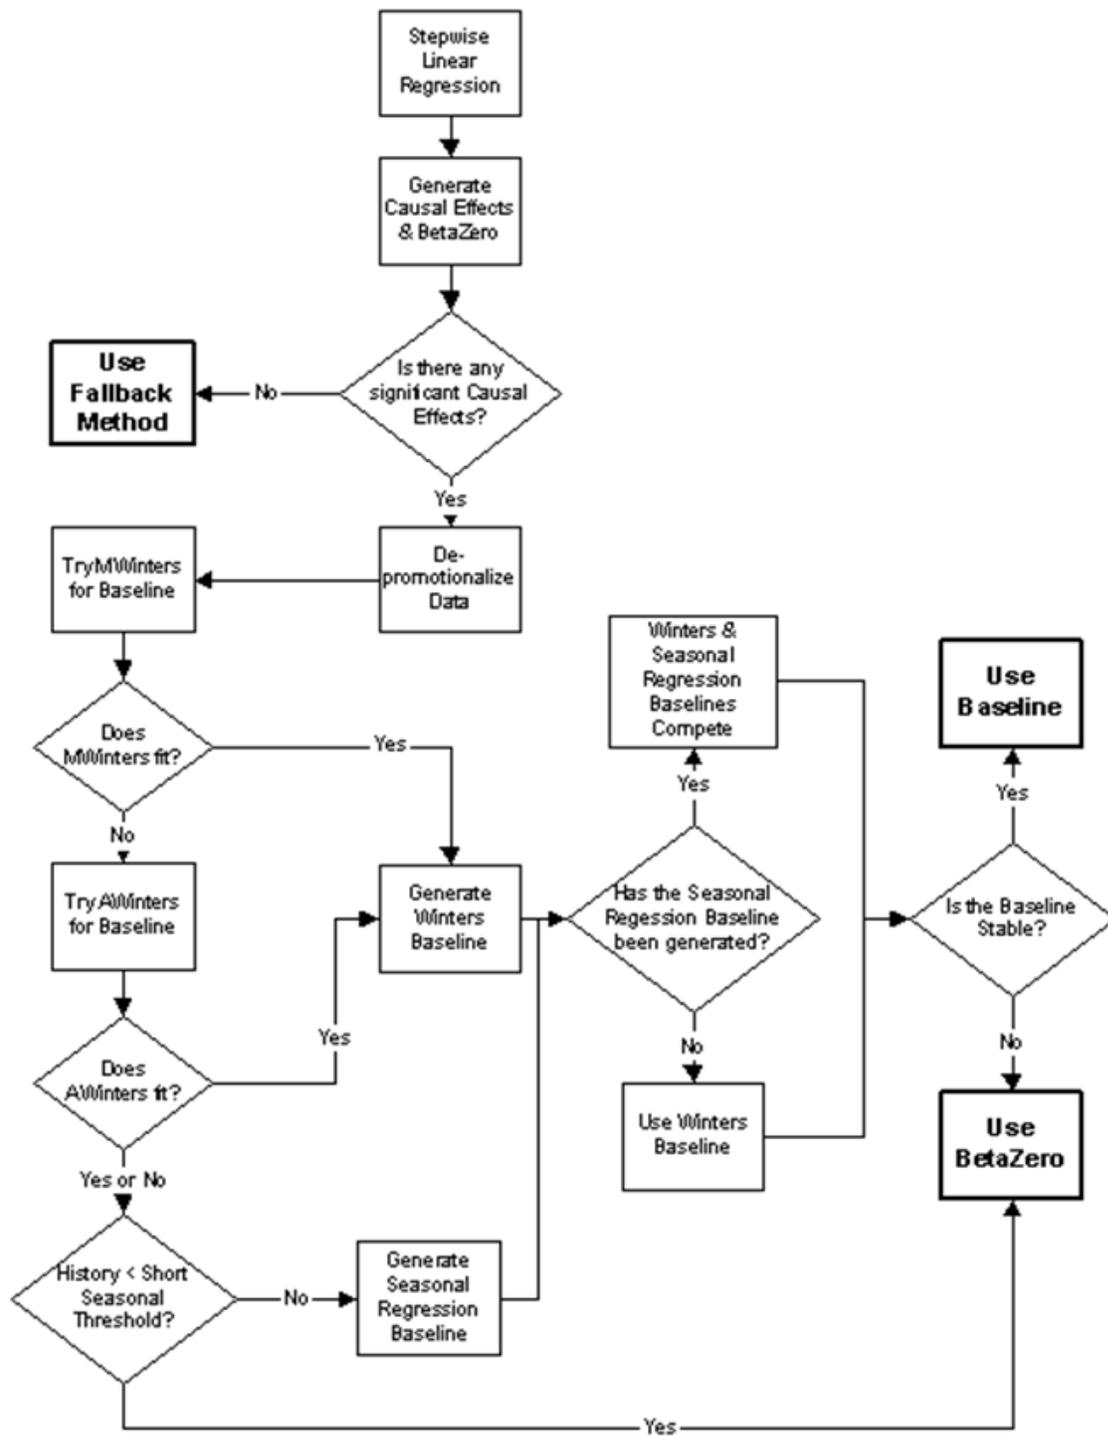
$$fit(t) = b_0 + \sum b_i * p_i(t)$$

8. The Time Series is de-causalized in the fit region of the history by removing the causal effects. That is, subtracting the additive causal effects for the product/location/time positions where the corresponding promo variables are on.
9. A seasonal model is fitted to the de-causalized series. The seasonal model is the winner of a competition between SeasonalES, Additive Winters, and Multiplicative Winters. The RMSE is determined across the fit region on the de-causalized series, and the winning model is the one that produces the least RMSE.
10. The winning seasonal model is then used to determine the seasonal forecast by applying the $\text{fit}(t)$ function across the forecast horizon. This forecast is also exported as the baseline forecast (that is forecast without any causal influences).
11. The forecast is obtained by re-causalizing the seasonal forecast. This is done by adding back the additive causal effects for the product/location/time positions where the corresponding promo variables are on in the forecast region.
12. The binary writes the winning promotional variables effects back to the database.
13. The selected model is recorded in the database.
14. The binary records the forecast and the baseline in the database.

Flowchart

[Figure 3-7](#) summarizes the Causal Forecasting process.

Figure 3-7 Causal Forecasting Process.



Causal Forecasting with External Baseline

Rather than use a sales history that may not have sufficient or accurate data, users can load an external baseline into the RDF Causal Engine instead. External baselines are often generated using data that is rolled up to a higher dimension than SKU/store, providing a greater depth of data and hence a less-noisy sales history. The sales data used to generate the external baseline can be corrected for out-of-stock, promotional, and other short-term event information using Preprocessing. These external baselines are then spread back to the SKU/store level and then loaded in the RDF Causal Engine. The Engine then uses this external baseline (rather than the causal baseline) along with the historic promotional data and future planned promotions data to create the system forecast, which is the external baseline with the lifts, which were calculated from the promotional data, applied on top.

Causal Forecasting Using External Baseline Process

When an external baseline is used, the AutoES binary code executes in the following manner:

1. The binary reads the history of the time series.
2. The binary reads the type of each promotional variable into the system.
3. The binary reads in all the promotional variables that apply to the series.
4. The binary creates an internal promotional variable to allow the modeling of trend.
5. Promotional variables, internal promotional variables, promotional variable types, and the series itself are passed to the stepwise regression routine, with the historic data serving as the dependent variables.
6. If the regression finds no significant promotional variables, the casual method is considered to have failed to fit. In this case, the forecast equals the external baseline.
7. The fit at time t , $fit(t)$, is defined in terms of b_0 , the intercept of the regression, b_i , the effect corresponding to promotional variable i , and $p_i(t)$, is the value of promo variable I , in time t as:

$$fit(t) = b_0 + \sum b_i * p_i(t)$$
8. The forecast is obtained by re-causalizing the external baseline. This is done by adding back the additive causal effects for the product/location/time positions, where the corresponding promo variables are on in the forecast region.
9. The binary writes the winning promotional variables effects back to the database.
10. The selected model is recorded in the database.
11. The binary records the forecast and the baseline in the database.

Causal Forecasting for Short-lifecycle Items

To forecast short-lifecycle promotional items, Causal deprices, depromotes, and smoothes the forecasting data source to generate the short lifecycle forecast causal baseline. In essence, causal forecasting uses the external baseline, where the generation of the baseline and the promotional lifts is modified for items with short lifecycle.

First, the external baseline is generated. Then the forecast data source is depriced, depromoted and smoothed. The resulting feed is aggregated and then spread down using rate of sales as a profile to create a lifecycle curve. The lifecycle curve is shifted

and stretched or shrunk to fit the new season length. This curve represents the pre-season baseline forecast. In season, the pre-season forecast serves as a forecast plan to the Bayesian forecasting method. The Bayesian forecast is the causal external baseline for short lifecycle items. The next step is to calculate the promotional lifts.

To do this, the promotional lifts are filtered from the historical sales and applied on top of the item's rate of sale. The resulting measure is regressed against the promotional variables to determine the promotion effects. Finally, the promotion effects are applied on top of the short-lifecycle causal baseline to generate the final forecast.

Causal Forecasting at the Daily Level

The causal forecasting at the daily level is calculated by spreading the weekly causal forecast down to day. The spreading utilizes causal daily profiles, thus obtaining a causal forecast at the day granularity.

The daily casual forecast process executes in the following manner:

1. Preprocess the day-level promotional variables by multiplication with daily profiles. Aggregate the preprocessed continuous day level promotional variables to the week level.
2. Calculate the causal forecast at the weekly level. Set promotional effects if desired. Use the RDF causal engine to generate the forecast.
3. Calculate the multiplicative promotional effects at the item/store level for every promo variable. The effects can be either:
 - Manually preset (See Step 1).
 - Calculated. When calculating the causal forecast, the calculated causal effects are written back to the database. If the effects are calculated at higher level than item/store, the effects are replicated down to item/store since the effects are multiplicative. If source-level forecasting is used and causal method is used both at the source-level and at the final-level, the effects from the final-level is used.
4. Daily profiles are calculated using the Curve module. Since as much history as possible is used and is averaged over seven days, it's assumed that these profiles are de-causalized. The de-causalized daily profiles capture the day-of-week effect and should be quite stable.
5. Causal effects are applied to the daily profiles. The profiles are multiplied by the causal effects and then the profiles have to renormalize.

Example

For every item/store combination, calculate a normal week-to-day profile based on historic data. Note that this profile is already computed for spreading the weekly forecasts to the day level. Suppose for a certain product, the profile is as follows:

Mon	Tues	Wed	Thu	Fri	Sat	Sun	Week
10%	10%	10%	10%	20%	30%	10%	100%

Suppose that in the past, the promotion was held on Wednesday, Thursday, and Friday of week w6:

Mon	Tues	Wed	Thu	Fri	Sat	Sun
		P	P	P		

Then the continuous weekly indicator for this promotion in w6 should be set to 0.4, which is the sum of the weights of Wednesday, Thursday, and Friday.

Now assume that the same promotion is held in a future week (w36), but only on Thursday:

Mon	Tues	Wed	Thu	Fri	Sat	Sun
			P			

Then the continuous weekly indicator for w36 should be set to 0.1, which is the weight of Thursday only.

The approach to use the continuous promotion indicators to generate an accurate causal forecast at the day level is as follows:

- Calculate the weekly multiplicative effect for the promotion using the standard causal forecasting system with continuous indicators.
- Calculate the forecast for w36 using the standard causal forecasting system with continuous indicators.
- Update the week-to-day profile of w36 so that the weight of Thursday is doubled (the multiplicative factor is 2):

Mon	Tues	Wed	Thu	Fri	Sat	Sun	Week
10%	10%	10%	20%	20%	30%	10%	110%

- Normalize the profile for w36:

Mon	Tues	Wed	Thu	Fri	Sat	Sun	Week
9%	9%	9%	18%	18%	27%	9%	100%

Finally, spread the forecast of w36 using the normalized profile.

Final Considerations About Causal Forecasting

In the Oracle Retail approach to causal forecasting, the causal effects are obtained by fitting a stepwise linear regression model that determines which variables are most relevant and what effect those relevant variables have on the series. The data used to fit the regression is the fit history of each time series, so basically a model is fit per time series. A problem arises due to potential lack of significant data (that is, when a promotional variable is not represented in the history, but it is present in the forecast region). In that case, the effect for that variable would not be computed at all, thus affecting the accuracy of the forecast. There are a few solutions that make use of the effects from other similar time series. One solution would be to do source-level causal forecasting and then spread down to the final-level.

This would be equivalent to using the effects at the source-level for time series that have no Promo Effect instances in the history. This has a serious conceptual drawback. By aggregating the promotional variables at the source-level, we would force the

effects on the other time series in the same aggregation class that would otherwise not have the Promo Effects on at the same time. An alternate solution is whenever a causal effect cannot be computed because of lack of significant data. An averaged effect from another time-series in the same aggregation class is going to be used instead. Refer to [Table 13-2, "Promo Effect Types"](#) for information on the Override Higher Level type.

Forecast Product/Location Cloning Administration Task

The Product/Location Cloning Administration task allows you to specify parameters to clone or copy history for new items or locations from existing items and locations. Since RDF needs history for forecasting, this functionality allows you to specify clone items and clone stores for these new Product/Locations.

Cloning is different from Like-item/Sister-Store functionality. Like-item/Sister-Store functionality copies forecasts or history of the Like-item or Sister-Store as the forecasts for the new SKU/store. However, in cases where the two items have different promotion calendars, it is necessary to forecast based on history data and promotion calendar. Item/Location Cloning functionality facilitates this process.

Also, by copying history from the parent product, you can start forecasting almost immediately after the launch of the new product.

Product/Location Cloning Administration can be performed from the Master or Local domain. The actual cloning or copying of history can be performed as part of batch run.

Note: Clone is available as a special expression that can be configured to be called as a batch run to perform the actual copying of history.

When the Cloning administration task is built in the Master domain, if clone assignments are made such that a product is assigned a clone that belongs to a different local domain, then at the time of batch run RDF does not clone that specific product and continue processing after logging a message.

Note: This task is only valid at final-levels. Therefore, this task may include hierarchy dimensions that are at higher positions than item or store (for example, subclass or region).

Product/Location Cloning Administration Workbook

To build the Product/Location Cloning Administration workbook, perform these steps:

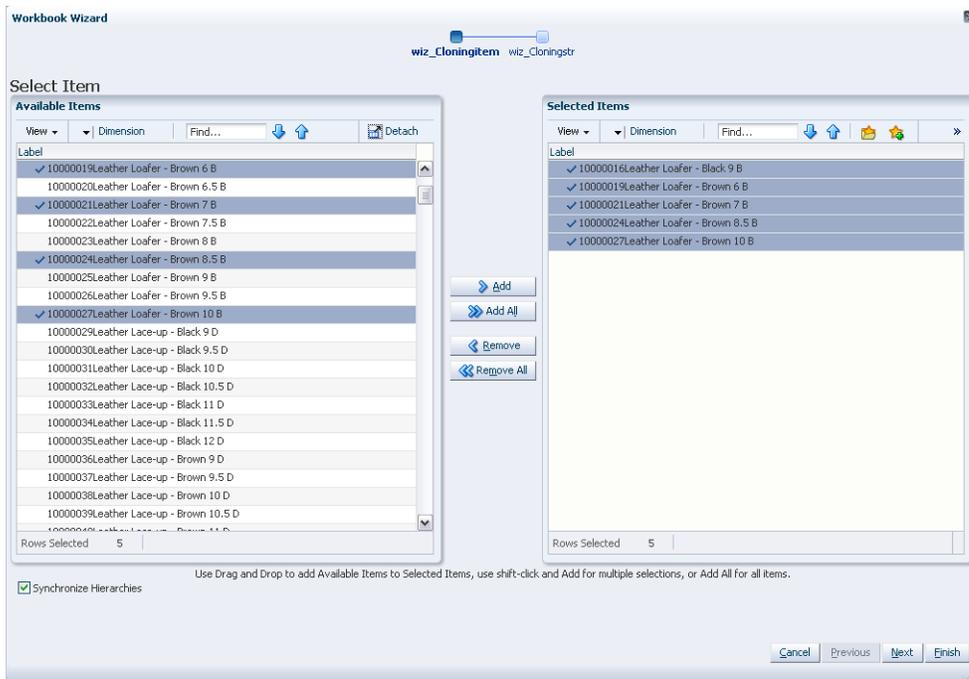
1. Click the New Workbook icon in the Product/Location Cloning Administration task.

Figure 4–1 Product/Location Cloning Administration Task



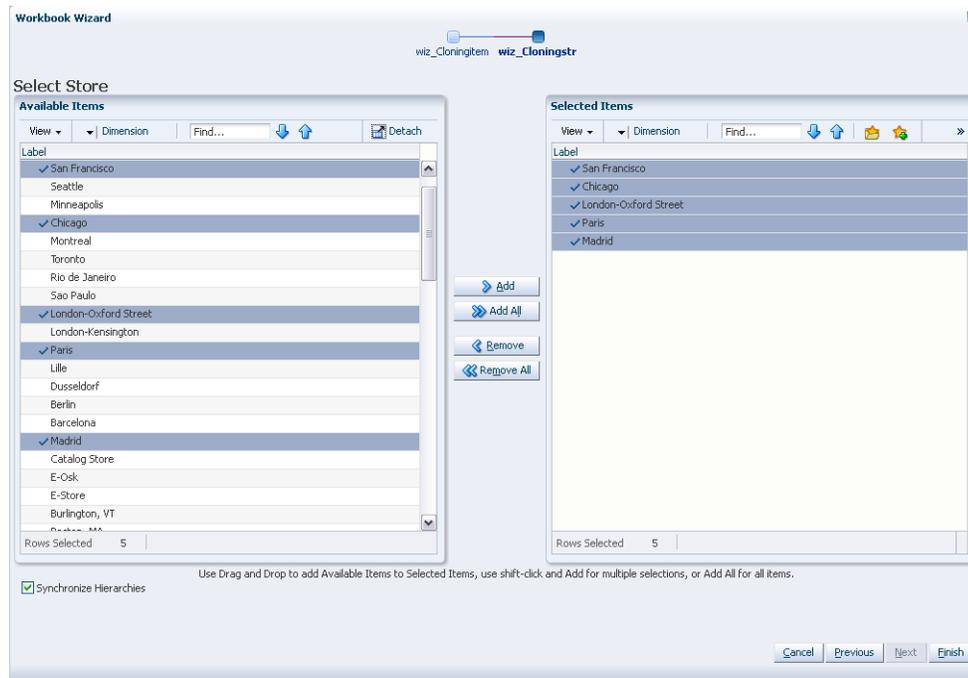
2. The Workbook wizard opens. Select the items you want to work with and click **Next**.

Figure 4–2 Workbook Wizard: Select Item



3. Select the stores you want to work with and click **Finish**.

Figure 4–3 Workbook Wizard: Select Stores



The Product/Location Cloning Administration workbook is built.

Adjustment Parameters View

The Adjustment Parameters view allows you to specify parameters that determine how much the cloned history is adjusted to match the level of the actual sales.

Figure 4–4 Adjustment Parameters View

Adjustment Parameters	
Cloned History Adjustment - Alpha	0.00
Cloned History Adjustment - Recent Sales Threshold	0
Clone History Adjustment - Calculation Window Length	0

Measures: Adjustment Parameters View

The Adjustment Parameters view contains the following measures:

Cloned History Adjustment - Alpha (Range 0-1)

This parameter is used to calculate the ratio that is applied to the cloned history. If this parameter is close to zero (0), then the cloned sales will be scaled so that they are aligned with the sales level. If this parameter is close to one (1), then little scaling of the cloned history will occur.

Cloned History Adjustment - Recent Sales Threshold

This parameter represents a specific number of periods. If it is less than the number of periods from the first sale date to the current date, then the cloned history is adjusted. If it is greater than the number of periods from the first sale date to the current date, then no adjustment is made to the cloned history.

Cloned History Adjustment - Calculation Window Length

This parameter is the number of periods that are considered when adjusting the cloned history.

Figure 4–5 Adjustment Ratio Formula

$$\text{adjustment ratio} = \left(\frac{\text{sum of first } M \text{ periods of sales}}{\text{sum of last } M \text{ periods of cloned history}} \right)^{(1-\alpha)}$$

Figure 4–6 Cloned History Formula

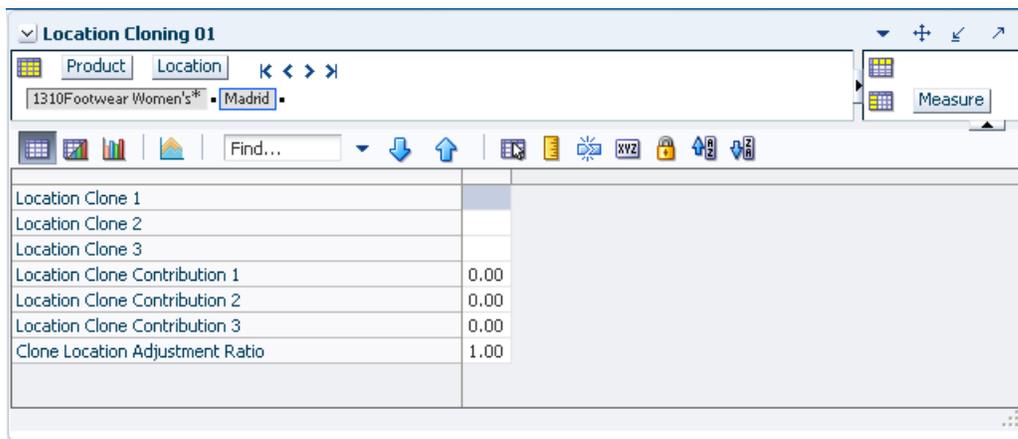
$$\text{cloned history} = \text{adjustment ratio} * \text{unadjusted cloned history}$$

Location Cloning View

Location Cloning view allows you to make clone Location assignments for new Locations. Location Cloning view allows you to specify up to three Clone Locations for a given Location that is being cloned. Also, it allows you to specify different parent Locations for different products. For example, the Beverly Hills store sells like the Miami store for swimwear and sells like the New York store for electronics. The Location Cloning view thus is at Location and Product intersection.

You have the ability to specify associated contributions for each of the Parent Locations and have the ability to specify an overall Adjustment Ratio.

Figure 4–7 Location Cloning View



Measures: Location Cloning View

The Location Cloning view contains the following measures:

Location Clones

Location Clone1, Location Clone 2, Location Clone 3

Parent Locations for the new Location being cloned. It is optional to specify more than one parent clone. The selection or assignment of Location Clone is made by clicking this measure, and a Single Hier Select wizard opens. Assignments are made by selecting a Location and clicking **OK**. Refer to the *Oracle Retail Predictive Application Server User Guide* for more details on the Single Hier Select wizard.

Location Clone Contributions

Location Clone Contribution1, Location Clone Contribution2, Location Clone Contribution3

Relative contributions of all Clone Locations. If user does not specify contributions, but makes selections for Location Clones and commits the changes, at the time of commit the contributions are calculated to have equal weights for all specified Location Clones. If user specifies contributions, at the time of commit, Location clone contributions get normalized so that they add up to one (1). Location Clone contributions with no associated Location Clone values get set to zero at the time of commit.

Clone Location Adjustment Ratio

You may enter an Adjustment Ratio to apply to the cloned history of the new location. This is a real number between [zero (0), infinity). The default (NA) value is 1.00 (in other words 100%), which translates to no adjustment.

Example 4–1 Clone Location Adjustment Ratio

If demand for a new store is expected to be 30% greater than its clone store, the Adjustment Ratio would be set to 1.30. If demand for a new store is expected to be 30% less than its clone store, the Adjustment percent would set to 0.70.

Note: Adjustment Ratios and Clone Contributions specified in the Product Cloning and Location Cloning views are used together while evaluating the result for a given item/store.

Example 4–2 Adjustment Ratios and Clone Contributions

SKU1	
SKU2	SKU3
20%	80%
Adjustment Ratio = 1	

STR1:	
STR2	STR3
50%	50%

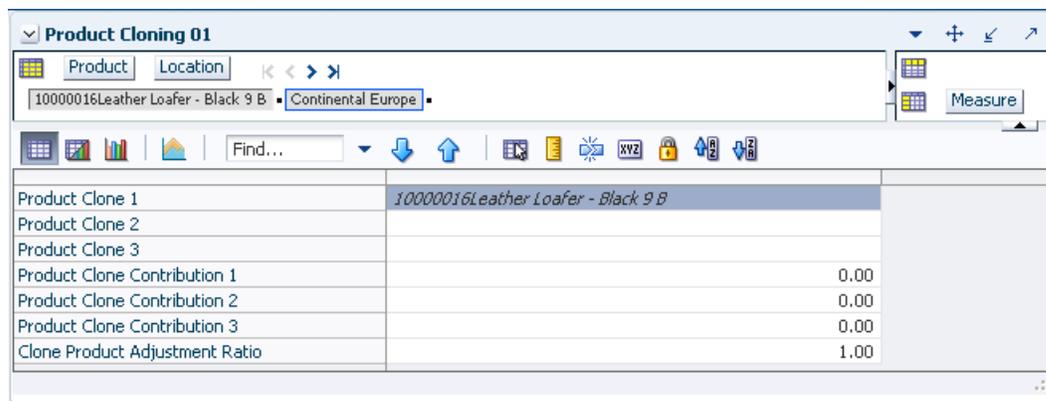
STR1:	
STR2	STR3
Adjustment Ratio = 0.5	
The contributions are calculated as:	
SKU2/STR2 at 5% (=20% x 0.5x50%), SKU2/STR3 at 5%,SKU3/STR2 at 20%(=80%x0.5x50%),SKU3/STR3 at 20%	

Product Cloning View

The Product Cloning view allows you to make Clone Product assignments for new Products. From this view, you can specify up to three parent Products for a given Product that is being cloned. Also, it allows you to specify different parent Products for different Locations. For example, CS Spring Water 1 sells like Diet Soda single in North America, but sells like Regular Soda in Asia. The Product Cloning view thus is at Product and Location intersection.

You have the ability to specify associated contributions for each of the Clone Products and have the ability to specify an overall Adjustment Ratio.

Figure 4–8 Product Cloning View



Measures: Product Cloning View

The Product Cloning view contains the following measures:

Product Clones

Product Clone1, Product Clone 2, Product Clone 3

Parent Products for the new Product being cloned. It is optional to specify more than one product clone. The selection or assignment of product clone is made by clicking this measure, and a Single Hier Select wizard opens. Assignments are made by selecting a product and clicking OK. Refer to the *Oracle Retail Predictive Application Server User Guide* for more details on the Single Hier Select wizard.

Product Clone Contributions

Product Clone Contribution1, Product Clone Contribution2, Product Clone Contribution3

Relative contributions of all clone products. If user does not specify contributions, but makes selections for product clones and commits the changes, at the time of commit

the contributions are calculated to have equal weights for all specified product clones. If user specifies Product clone contributions, they get normalized such that they add up to one at the time of commit. Product clone contributions with no associated product clone values get set to zero at the time of commit.

Clone Product Adjustment Ratio

You may enter an Adjustment Ratio to apply to the cloned history of the new product. This is a real number between [zero (0), infinity]. The default (NA) value is 1.00 (in other words 100%), which translates to no adjustment.

Example 4–3 Clone Product Adjustment Ratio

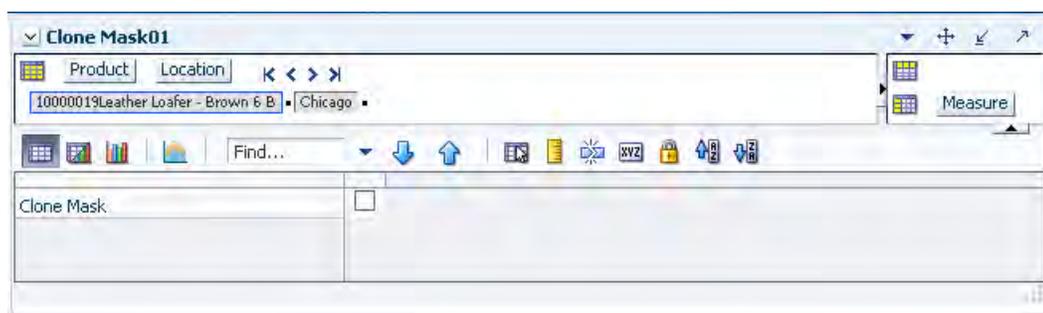
If demand for a new item is expected to be 30% greater than its clone item, the Adjustment Ratio would be set to 1.30. If demand for a new item is expected to be 30% less than its clone item, the Adjustment percent would set to 0.70.

Clone Mask View

This view allows you to enable cloning for product/locations by setting the Clone Mask measure to **true**.

Note: Clone Mask measure can also be brought in using the Show/Hide option in the Product Cloning and Location views. The clone mask can also be set based on business rules by configuring RPAS rules.

Figure 4–9 Clone Mask View



Measures: Clone Mask View

The Clone Mask view contains the following measures:

Clone Mask

This measure specifies if cloning should be performed for the item/store combination.

Forecast Administration Task

The Forecast Administration task allows you to set up RDF to generate forecasts.

Forecast Administration Task Components

This section describes the components of the Forecast Administration.

Basic versus Advanced

Forecast Administration is the first task used in setting up RDF to generate forecasts. It provides access to forecast settings and parameters that govern the whole domain (database). These settings and parameters are divided into two areas, Basic and Advanced, and are available beneath the Forecast Administration task.

The Basic Settings are used to establish a final-level forecast horizon, the commencement and frequency of forecast generation, and the specification of aggregation levels (Source Levels) and spreading (Profile) methods used to yield the final-level forecast results. For information on the views in Basic Settings, refer to [Forecast Administration Workbook: Basic Settings](#).

The Advanced Settings are used to set default values for parameters affecting the algorithm and other forecasting techniques used to yield final-level and source-level forecasts, thus eliminating the need to define these parameters individually for each product and location. If certain products or locations require parameter values other than the defaults, these fields can be amended on a case-by-case basis in the [Forecast Maintenance Workbook](#). For information on the views in Advanced Settings, refer to [Forecast Administration Workbook: Advanced Settings](#).

Final versus Source Level Forecasts

Often, forecast information is required for items at a very low level in the hierarchy. However, problems can arise in that data is often too sparse and noisy to identify clear patterns at these lower levels. For this reason, it sometimes becomes necessary to aggregate sales data from a low level to a higher level in the hierarchy in order to generate a reasonable forecast. Once this forecast is created at the higher or source-level, the results can be allocated to the lower-level or final-level dimension based on the lower level's relationship to the total.

In order to spread this forecast information back down to the lower level, it is necessary to have some idea about the relationship between the final-level and the source-level dimensions. Often, an additional interim forecast is run at the low level in order to determine this relationship. Forecast data at this low level might be sufficient to generate reliable percentage-to-whole information, but the actual forecast numbers are more robust when generated at the aggregate level.

The Final Level view represents forecast parameters for the lower (final) level, the level to which source forecast values are ultimately spread. Forecast data must appear at some final-level in order for the data to be approved or exported to other systems. The Source Level view represents the default values for forecast parameters at the more robust aggregate (source) level.

Forecasting Methods Available in RDF

A forecasting system's main goal is to produce accurate predictions of future demand. The RDF solution utilizes the most advanced forecasting algorithms to address many different data requirements across all retail verticals. Furthermore, the system can be configured to automatically select the best algorithm and forecasting level to yield the most accurate results.

Following is a summary of the various forecasting methods employed in RDF. Some of these methods may not be visible in your solution based on configuration options set in the RPAS Configuration Tools. More detailed information on these forecasting algorithms is provided in [Chapter 3, "Oracle Retail Demand Forecasting Methods."](#)

Average

RDF uses a simple average model to generate forecasts.

Moving Average

RDF uses a simple moving average model to generate forecasts. You can specify a Moving Average Window length.

IntermittentES

RDF fits the data to the Croston's model of exponential smoothing. This method should be used when the input series contains a large number of zero data points (that is, intermittent demand data). The original time series is split into a Magnitude and Frequency series, and then the SimpleES model is applied to determine level of both series. The ratio of the magnitude estimate over the frequency estimate is the forecast level reported for the original series.

TrendES

RDF fits the data to the Holt model of exponential smoothing. The Holt model is useful when data exhibits a definite trend. This method separates base demand from trend and then provides forecast point estimates by combining an estimated trend and the smoothed level at the end of the series. For instance, where the forecast engine cannot produce a forecast using the TrendES method, the Simple/IntermittentES method is used to evaluate the time series.

Causal

Causal is used for promotional forecasting and can only be selected if Promote is implemented. Causal uses a Stepwise Regression sub-routine to determine the promotional variables that are relevant to the time series and their lift effect on the series. AutoES utilizes the time series data and the future promotional calendar to generate future baseline forecasts. By combining the future baseline forecast and each promotion's effect on sales (lift), a final promotional forecast is computed. For instances where the forecasting engine cannot produce a forecast using the Causal method, the system evaluates the time series using the SeasonalES method.

To forecast short lifecycle, promotional items, Causal deprices, depromotes, and smoothes the forecasting data source to generate the short-lifecycle forecast causal baseline. The promotion effects are calculated using the same Stepwise Regression

code as Causal. However, the promo lifts are first filtered from the sales and applied on top of the item's rate of sale. The resulting measure is regressed against the promotional variables to determine the promotion effects. Finally, the promotion effects are applied on top of the short lifecycle baseline to generate the final forecast.

No Forecast

No forecast is generated for the product/location combination.

Bayesian

Useful for short lifecycle forecasting and for new products with little or no historic sales data, the Bayesian method requires a product's known sales plan (created externally to RDF) and considers a plan's shape (the selling profile or lifecycle) and scale (magnitude of sales based on Actuals). The initial forecast is equal to the sales plan, but as sales information comes in, the model generates a forecast by merging the sales plan with the sales data. The forecast is adjusted so that the sales magnitude is a weighted average between the original plan's scale and the scale reflected by known history. A Data Plan must be specified when using the Bayesian method. For instances where the Data Plan equals zero (0), the system evaluates the time series using the SeasonalES method.

Profile-based

RDF generates a forecast based on a seasonal profile that can be created in RPAS or legacy system. Profiles can also be copied from another profile and adjusted. Using historic data and the profile, the data is de-seasonalized and then fed to the SimpleES method. The Simple forecast is then re-seasonalized using the profiles. A Seasonal Profile must be specified when using the Profile-Based method. For instances where the Seasonal Profile equals zero, the system evaluates the time series using the SeasonalES method.

Components

The Components forecast method multiplies pre-calculated baseline, promotional lifts, and regular price lifts to generate the forecast. These components are specified in Forecast Administration task, Advanced Settings activity, Final Level Parameters view. The following table lists the components and their referenced measures.

Component	Measure
Baseline	Components - Baseline
Regular Price Lifts	Components - Regular Price Lifts
Promotional Lifts	Components - Promotional Lifts
Interval	Components - Baseline Interval

Intervals are not generated and need to be provided. This method is useful when the forecast components are available since the forecast can be quickly generated.

Automatic Exponential Smoothing (AutoES)

RDF fits the sales data to a variety of exponential smoothing models of forecasting, and the best model is chosen for the final forecast. The candidate methods considered by AutoES are:

- [SimpleES](#)
- [Croston's Method](#)

- [Simple/IntermittentES](#)
- [HoltES](#)
- [Multiplicative Seasonal](#)
- [Additive Seasonal](#)
- [SeasonalES](#)
- [Seasonal Regression](#)

The final selection between the models is made according to a performance criterion (Bayesian Information Criterion) that involves a trade-off between the model's fit over the historic data and its complexity.

SimpleES

RDF uses a simple exponential smoothing model to generate forecasts. SimpleES ignores seasonality and trend features in the demand data and is the simplest model of the exponential smoothing family. This method can be used when less than one year of historic demand data is available. For more information, see [Simple Exponential Smoothing](#).

Croston's Method

Croston's Method is used when the input series contains a large number of zero data points (that is, intermittent demand data). For more information, see [Croston's Method](#).

Simple/IntermittentES

Simple/IntermittentES is a combination of the SimpleES and IntermittentES methods. This method applies the SimpleES model unless a large number of zero data points are present, in which case the Croston's model is applied. For more information, see [Simple/Intermittent Exponential Smoothing](#).

HoltES

Holt exponential smoothing treats data as linearly trended but non-seasonal. The Holt model provides forecast point estimates by combining an estimated trend (for the forecast horizon - h) and the smoothed level at the end of the series. RDF uses a damped Holt model that decays the trend component so that it disappears over the first few weeks. This improves forecasts created using Holt over longer forecast horizons. For more information, see [Holt Exponential Smoothing](#).

Multiplicative Seasonal

Also referred to as Multiplicative Winters Model, this model extracts seasonal indices that are assumed to have multiplicative effects on the un-seasonalized series. For more information, see [Multiplicative Winters Exponential Smoothing](#) which includes:

- Oracle Winters
- Oracle Winters Decomposition
- Winters Standard
- Winters Responsive

Additive Seasonal

Also referred to as Additive Winters Model, this model is similar to the Multiplicative Winters model, but is used when zeros are present in the data. This model adjusts the

un-seasonalized values by adding the seasonal index (for the forecast horizon). For more information, see [Additive Winters Exponential Smoothing](#).

SeasonalES

This method, a combination of several Seasonal methods, is generally used for known seasonal items or forecasting for long horizons. This method applies the Multiplicative Seasonal model unless too many zeros are present in the data, in which case the Additive Winters model of exponential smoothing is used. If less than two years of data is available, a Seasonal Regression model is used. If there is too little data to create a seasonal forecast (in general, less than 52 weeks), the system selects from the SimpleES, TrendES, and IntermittentES methods. For more information, see [Seasonal Exponential Smoothing \(SeasonalES\)](#).

Seasonal Regression

Seasonal Regression cannot be selected as a forecasting method, but is a candidate model that is used only when the SeasonalES method is selected. This model requires a minimum of 52 weeks of history to determine seasonality. Simple Linear Regression is used to estimate the future values of the series based on a past series.

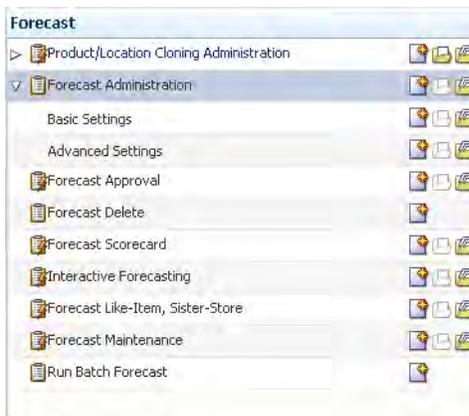
The independent variable is the series history one-year or one cycle length prior to the desired forecast period, and the dependent variable is the forecast. This model assumes that the future is a linear combination of itself one period before plus a scalar constant. For more information, see [Seasonal Regression](#).

Forecast Administration Workbook

To build the Forecast Administration workbook, perform these steps:

1. Click the New Workbook icon in the Forecast Administration task.

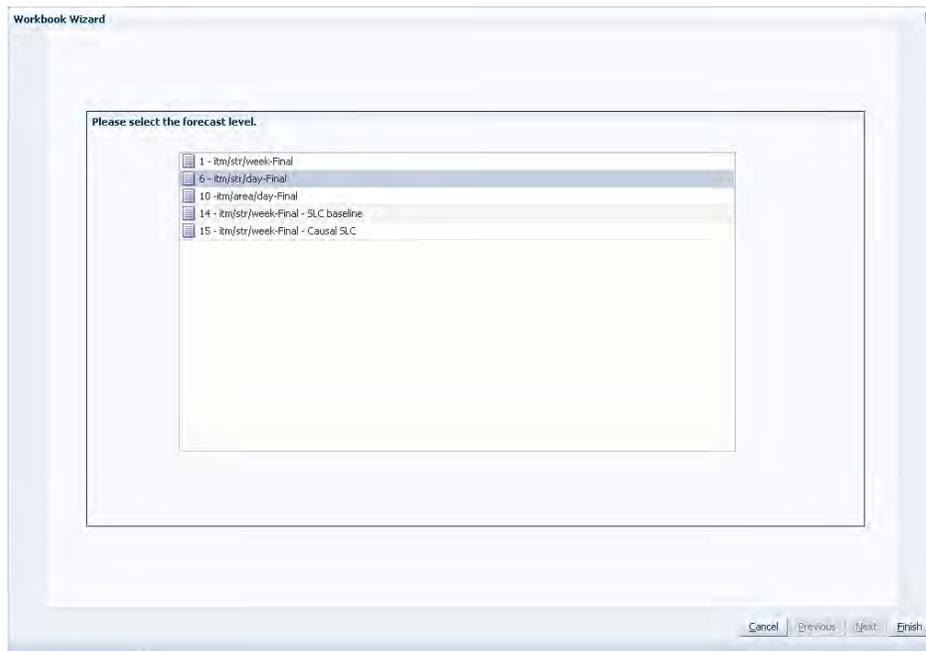
Figure 5–1 Forecast Administration Task



2. The Workbook wizard opens. Select the forecast level you want to work with and click **Finish**.

Note: The final forecast level is a level at which approvals and data exports can be performed. Depending on your organization's setup, you may be offered a choice of several final forecast levels. Make the appropriate selection.

Figure 5–2 Workbook Wizard: Select Forecast Level



The Forecast Administration workbook is built.

Forecast Administration Workbook: Basic Settings

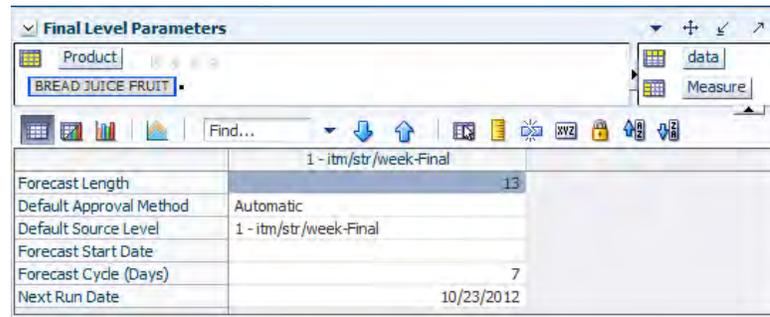
The Forecast Administration Basic Settings workbooks contains two views:

- [Final Level View - Basic Settings](#)
- [Final and Source Level Parameters View - Basic Settings](#)

Final Level View - Basic Settings

The Final Level view allows you to set the forecast horizon information, frequency of review, and all default parameters for the lower or final-level forecast (the level to which aggregate forecast data is ultimately spread). Forecast approvals and data exports can only be performed on forecasts at a final-level. [Figure 5–3](#) provides an example of a view of the Final Level Parameters view in a master domain with three partitions/local domains, partitioned based on group.

Figure 5–3 Final Level Parameters View



Measures: Final Level View - Basic Settings

The Final Level View - Basic Settings contains the following measures:

Default Approval Method

This field is a list from which you select the default automatic approval policy for forecast items. Valid values are:

Field	Description
Manual	The system-generated forecast is not automatically approved. Forecast values must be manually approved by accessing and amending the Forecast Approval Workbook .
Automatic	The system-generated quantity is automatically approved as is.
By Alert	This list of values may also include any Forecast Approval alerts that have been configured for use in the forecast approval process. Alerts are configured during the implementation. Refer to the <i>Oracle Retail Predictive Application Server Configuration Tools User Guide</i> for more information on the Alert Manager.

Default Source Level

The pick list of values displayed in this field allows you to change the forecast level that is used as the primary level to generate the source forecast. The source-levels are set up in the RPAS Configuration Tools. A value from the pick list is required in this field at the time of forecast generation.

Forecast Cycle

The Forecast Cycle is the amount of time (measured in days) that the system waits between each forecast generation. Once a scheduled forecast has been generated, this field is used to automatically update the Next Run Date field. A non-zero value is required in this field at the time of forecast generation.

Forecast Start Date

This is the starting date of the forecast. If no value is specified at the time of forecast generation, the system uses the data/time at which the batch is executed as the default value. If a value is specified in this field and it is used to successfully generate the batch forecast, this value is cleared.

Forecast Length

The Forecast Length is used with the Forecast Start Date to determine forecast horizon. The forecast length is based on the calendar dimension of the final-level. For example,

if the forecast length is to be 10 weeks, the setting for a final-level at day is 70 (10 x 7 days).

Next Run Date

The Next Run Date is the date on which the next batch forecast generation process automatically runs. RDF automatically triggers a set of batch processes to be run at a pre-determined time period. When a scheduled batch is run successfully, the Next Run Date automatically updates based on the Start Date value and the Forecast Cycle. No value is required in this field when the [Run Batch Forecast Workbook](#) wizard is used to generate the forecast or if the batch forecast is run from the back-end of the domains using the override **true** option. Refer to the *Oracle Retail Demand Forecasting Implementation Guide* for more information on forecast generation.

Final and Source Level Parameters View - Basic Settings

The Final and Source Level view allows you to set the default parameters that are common to both the final and source-level forecasts.

Figure 5–4 Final and Source Level Parameters View

	1 - itm/str	2 -	3 -	4 -	5 -	9 -
Forecast Data Source	pos	pos	pos	pos	pos	pos
Default History Start Date						
Default Forecast Method	Simple	AutoES	AutoES	AutoES	AutoES	AutoES
Data Plan						
Seasonal Profile						
Spreading Profile			01	03	05	19

Measures: Final and Source Level Parameters View - Basic Settings

The Final and Source Level Parameters View - Basic Settings contains the following measures:

Forecast Data Source

This is a read-only value that displays the sales measure (the measure name) that is the data used for the generation of forecasts (for example, POS). The measure that is displayed here is determined at configuration time in the RPAS Configuration Tools. Different Data Sources can be specified for the Final and Source Levels.

Default History Start Date

This field indicates to the system the point in the historical sales data at which to use in the forecast generation process. If no date is indicated, the system defaults to the first date in your calendar. It is also important to note that the system ignores leading zeros that begin at the history start date. For example, if your history start date is January 1, 2011 and an item/location does not have sales history until February 1, 2011, the system considers the starting point in that item/location's history to be the first data point where there is a non-zero sales value. Different History Start Dates can be specified for the Final and Source Levels.

Data Plan

Used in conjunction with the Bayesian forecast method, Data Plan is used to input the measure name of a sales plan that should be associated with the final-level forecast. Sales plans, when available, provide details of the anticipated shape and scale of an item's selling pattern. If the Data Plan is required, this field should include the measure name associated with the Data Plan.

Default Forecast Method

The Default Forecast Method is a complete list of available forecast methods from which you can select the primary forecast method that is used to generate the forecast. A summary of methods is provided in [Forecasting Methods Available in RDF](#). The chapter, "[Oracle Retail Demand Forecasting Methods](#)", covers each method in greater detail.

It is important to note that Causal should not be selected unless the forecast level was set as a Causal level during the configuration. Refer to the *Oracle Retail Demand Forecasting Configuration Guide* for more information on configurations using the Causal forecast method.

Seasonal Profile

Used in conjunction with the Profile-Based forecasting method, this is the measure name of the seasonal profile that is used to generate the forecast at either the source or final-level. Seasonal profiles, when available, provide details of the anticipated seasonality (shape) of an item's selling pattern. The seasonal profile can be generated or loaded, depending on your configuration. The original value of this measure is set during the configuration of the RDF solution.

Spreading Profile

Used for source-level forecasting, the value of this measure indicates the profile level that is used to determine how the source-level forecast is spread down to the final-level. No value is needed to be entered at the final-level. For dynamically generated profiles, this value is the number associated with the final profile level (for example, 01). Note that profiles one (1) through 9 have a zero (0) preceding them in Curve—this is different than the forecasting level numbers.

For profiles that must be approved, this is the measure associated with the final profile level. This measure is defined as apvp +level (for example, apvp01 for the approved profile for level 01 in Curve).

Source Level Forecast Data Source

You can specify a measure to serve as a forecast data source at the source-level. If the measure is left blank for a certain source-level, the source measure from the corresponding final-level is used.

Source Level History Start Date

This is the starting date for historical sales data at the source-level. For example, if your system start date is January 1, 2003, but you only want to use historical sales data from the beginning of 2011, you need to set your History Start Date to January 1, 2011. Only history after this date is used for generating the forecast. The default is the system start date unless otherwise specified. If sales data is collected weekly or monthly, RDF generates forecasts only using data from sales periods after the one containing the history start date.

It is also important to note that the system ignores leading zeros that begin at the history start date. For example, if your history start date is January 1, 2011, and an

item/location does not have sales history until February 1, 2011, the system considers the starting point in that item/location's history to be the first data point where there is a non-zero sales value.

Note: History start dates for final- and source-levels can be different.

Forecast Administration Workbook: Advanced Settings

The Forecast Administration Advanced Settings workbooks contains these views:

- [Causal Parameters View - Advanced Settings](#)
- [Final Level Parameters View - Advanced Settings](#)
- [Final and Source Level Parameters View - Advanced Settings](#)
- [Regular Price Parameters View- Advanced Settings](#)

The Forecast Administration Advanced Settings workflow activity is used to set parameters related to either the data that is stored in the system or the forecasting methods that are used at the final or source-levels. The parameters on this workflow activity are not as likely to be changed on a regular basis as the ones on the Basic Settings workflow activity.

Causal Parameters View - Advanced Settings

The Causal Parameters view allows you to set the parameters that support promotional forecasting. This view only includes the causal forecast levels for the final-level selected during the wizard process. This view is not visible if Promote/Promotional Forecasting is not implemented in your RDF environment.

Figure 5–5 Causal Parameters View

	1 - itm/str/week-Final	2 - itm/str/week
Causal Aggregation Profile		
Causal Data Source		
Causal Calculation Intersection		
Causal Calculation Intersection Periodicity	0	0
Causal Higher Intersection		
Disable Causal	<input type="checkbox"/>	<input type="checkbox"/>
Causal Spread Profile		
Use Causal Capping	<input type="checkbox"/>	<input type="checkbox"/>
Causal Capping Number		
Causal Capping Ratio		
Minimum Causal Capping History (Periods)	0	0
Causal Beta Scale	50.00	50.00
Causal Short Season Threshold	52	52
Causal - Disable Seasonal Regression	<input type="checkbox"/>	<input type="checkbox"/>
Causal - Seasonal Baseline Selection Criteria	Avg Criteria	Avg Criteria
Causal Baseline Minimum Winters History Length	0	0
Causal External Baseline		
Causal - Lift Estimate	Hybrid	Hybrid
Causal - Smoothing Method	IndicatorBased	IndicatorBased
Causal SLC - ROS		
Causal - Effects Only	<input type="checkbox"/>	<input type="checkbox"/>
Aggregated Promotional Lift Override	Lift Override	Lift Override

Measures: Causal Parameters View - Advanced Settings

The Causal Parameters View - Advanced Settings contains the following measures:

Aggregated Promotional Lift Override

This parameter contains the name of the measure that stores the higher level (aggregated) promotional peak override. The name is implementation specific. If populated, during the causal forecasting generation the aggregated lift is spread down to the appropriate final level intersection.

Causal - Disable Seasonal Regression

This parameter gives you the option to remove Seasonal Regression from choice of Forecast Methods for baseline calculation within Causal Forecasting Method. It defaults to **false**, which means Seasonal regression is included in the choice of methods for baseline calculation within causal.

Causal - Effects Only

If Causal - Effects Only is set to **true**, the Causal forecasting method will only calculate promotional effects, and will not generate the forecast. If the measure is not set to **true**, then the Causal forecasting method generates promotional effects and forecasts.

Causal - Lift Estimate

The parameter gives you the option to use either the multiplicative or hybrid models to decausalize the sales history.

- Multiplicative—this model extracts promotional lifts that are assumed to have multiplicative effects in order to de-causalized the history.
- Hybrid—this model extracts the promotional lifts as additive effects in the de-causalized phase to decausalize the history. Then later in the forecast phase, it applies the promotional lifts as multiplicative effect to the forecast baseline. This is the system default value.

Causal - Seasonal Baseline Selection Criteria

This parameter allows you to choose whether to use minimum or average baseline criteria to determine baseline stability within Causal Forecasting Method.

- Average baseline criteria: If system calculated average baseline forecast < 4 , use Beta zero as System baseline forecast.
- Minimum baseline criteria: If system calculated minimum baseline forecast < 1 , then use Beta zero as System baseline forecast.

This parameter defaults to Average Baseline Criteria.

Causal - Smoothing Method

This parameter gives you the option to use either Indicator Based or Moving Average methods to smooth the sales history.

- Indicator Based: smoothes the periods that have active promotions (indicators) by using the sales of non-promoted periods in the time frame. This is the system default value.
- Moving Average: for each period t , the moving average method takes the average of the periods from $t-2$ to $t+2$ as the smoothed baseline.

Causal Aggregation Profile

Used only for Daily Causal Forecasting, the Causal Aggregation Profile is measure name of the profile used to aggregate promotions defined at day up to the week. The value entered in this field is the measure name of profile. If this profile is generated within Curve, the format of the measure name is apvp+level (for example, apvp01). Note that the only aggregation of promotion variables being performed here is along the Calendar hierarchy. RDF does not support aggregation of promotion variables along other hierarchies such as product and location hierarchies.

Causal Baseline Minimum Winters History Length

Used only in Causal forecasting, this parameter is the minimum number of historical time periods that are required for Winters to be used to produce the Causal baseline. The default setting is 104 periods, but it may be set to lower number, for example, 52 periods.

Causal Beta Scale

As part of the causal forecasting process, the average sales of non-promo periods in history are divided by the Beta Scale. If the resulting ratio is less than beta zero, the causal forecasting method succeeds. If not, the causal forecasting method fails and the Fallback Method is used to generate the forecast for the time series.

Causal Calculation Intersection

Used only for Daily Causal Forecasting, the Causal Calculation Intersection is the intersection at which the causal forecast is run. The format needs to match the hierarchy dimension names set in the RPAS Configuration Tools (such as itemstr_week). Each dimension must have only four characters. The order of the dimension does not matter. There is no validation of correct format of this intersection.

Causal Calculation Intersection Periodicity

Used only for Daily Causal Forecasting, the Causal Calculation Intersection Periodicity must be set to the periodicity of Causal Calculation Intersection. Periodicity is the number of periods within one year that correspond to the calendar dimension (for example, 52 if the Causal Calculation Intersection is defined with the week dimension).

Causal Capping Number

The value of this parameter is a measure based on the same product and location intersection as the forecast level. This measure contains the maximum value to use for calculating the causal forecast for time series that meet the Causal Capping conditions. If Use Causal Capping is set to **true**, the history for the time series is greater than or equal to the Minimum Causal Capping History, and the preliminary forecast is greater than or equal to the value in the Causal Capping Number measure, then the forecast is recalculated to be the value in the Causal Capping Number measure multiplied by the value in the Causal Capping Ratio measure.

Causal Capping Ratio

The value of this parameter is a measure based on the same product and location intersection as the forecast level. This measure contains the ratio to be used to calculate the forecasts for time series that meet the Causal Capping conditions. If Use Causal Capping is set to **true**, the history for the time series is greater than or equal to the value in the Minimum Causal Capping History measure, and the preliminary forecast is greater than or equal to the value in the Causal Capping Number measure, then the

forecast is recalculated to be the Causal Capping Number multiplied by the value in the Causal Capping Ratio measure.

Causal Data Source

Used only for Daily Causal Forecasting, the Causal Data Source is an optional setting that contains the measure name of the sales data to be used if the data for causal forecasting is different than the Data Source specified at the Final level. If needed, this field should contain the measure name of the source data measure (for example, DPOS).

If set to **true**, Seasonal Regression is excluded from the choice of methods for baseline calculation within Causal.

Note: This parameter does not impact the choice of forecast methods within SeasonalES.

Causal External Baseline

Enter the measure name that you want to serve as the baseline for the causal forecast at this level. The measure can be loaded or it can be generated in a different forecast run.

Causal Higher Intersection

An optional setting for Causal Forecasting, this intersection is the aggregate level to model promotions if the causal intersection cannot produce a meaningful causal effect. This intersection applies to promotions that have a Promotion Type is set to Override From Higher Level (set in the [Promotion Maintenance Workbook](#)). The format of this intersection needs to match the hierarchy dimension names set in the RPAS Configuration Tools such as sclsrgrn_ (Subclass/Region), and it must not contain the calendar dimension. Each dimension must have only four characters. The order of the dimension does not matter. There is no validation of correct format of this intersection.

Causal Short Seasonal Threshold

Used only in Causal forecasting, this parameter is the minimum number of historical time periods that are required for SeasonalES to be used to produce the Causal forecast. Otherwise, the baseline from Causal regression is used. The algorithm currently uses 52 weeks of history as the default for this parameter, but it may be set from 13 to 56 weeks. For an item with history less than the value in the Causal Short Seasonal Threshold and greater than 12 weeks, the Beta Zero is used only if Causal found the initial effects to be significant. Otherwise, the item is forecast using the Fallback Method.

Causal SLC - ROS

Enter the name of the measure that stores the rate of sale values for item/locations that are forecast using the Causal short lifecycle methodology.

Causal Spread Profile

Used only for Daily Causal Forecasting, the Causal Spread Profile is the measure name of the profile used to spread the causal baseline forecast from the Causal Calculation Intersection to the Final Level. If this profile is generated in Curve, this measure value is apvp+level (for example, apvp01).

Disable Causal

Place a check in this parameter (set to **true**) if causal forecasting is to be disabled for the forecast level. This parameter serves as a global override. This parameter can also help with performance when not running Causal forecast method.

Minimum Causal Capping History

Used only in Causal forecasting, this parameter is the minimum number of historical time periods that are required for Causal Capping to be used to produce the Causal forecast. Otherwise, no capping is applied. The algorithm currently uses 52 weeks of history as the default for this parameter.

Minimum Causal Capping History (Periods)

If Use Causal Capping is set to **true**, this parameter is used to set minimum number of historical time periods required before the system considers a time series for causal capping.

Use Causal Capping

Place a check in this parameter (set to **true**) if capping is to be applied to the causal forecast. Also required for causal capping are the following parameters:

- Causal Capping Number
- Causal Capping Ratio
- Minimum Causal Capping History

Note: More detailed information on the Causal forecasting algorithm is provided in [Forecasting Methods Available in RDF](#).

Final Level Parameters View - Advanced Settings

The Final Level view allows you to set the advanced parameters for the final-level forecasts. [Figure 5–6](#) provides an example of a view of this view in a master domain with three partitions/Local Domains, partitioned based on group.

Figure 5–6 Final Level Parameters View

1 - itm/str/week-Final	
Generate Baseline Forecasts	<input type="checkbox"/>
Generate Cumulative Intervals	<input type="checkbox"/>
Generate Intervals	<input type="checkbox"/>
Generate Methods	<input type="checkbox"/>
Generate Parameters	<input type="checkbox"/>
Updating Last Week Forecast	0
Updating Last Week Forecast	No Change
Days to Keep Forecasts	14
Like TS Duration (Periods)	10
Item End Date Action	Apply 0 After Spreading
Store Interim Forecast	<input type="checkbox"/>
Default Keep Last Changes	None
Interval Cap Lower Ratio	0.30
Interval Cap Upper Ratio	1.00

Measures: Final Level Parameters View - Advanced Settings

The Final Level Parameters view - Advanced Settings contains the following measures:

Components - Baseline

This measure stores the name of the measure that serves as baseline in the Components forecasting measure. If no measure is specified, the Components forecasting method assesses the value as one (1).

Components - Baseline Interval

This measure stores the name of the measure that serves as baseline interval in the Components forecasting measure. If no measure is specified, the Components forecasting method assesses the value as one (1).

Components - Regular Price Lifts

This measure stores the name of the measure that serves as regular price lifts in the Components forecasting measure. If no measure is specified, the Components forecasting method assesses the value as one (1).

Components - Promotional Lifts

This measure stores the name of the measure that serves as promotional lifts in the Components forecasting measure. If no measure is specified, the Components forecasting method assesses the value as one (1).

Days to Keep Forecasts

This field is used to set the number of days that the system stores forecasts based on the date/time the forecast is generated. The date/time of forecast generation is also referred to as birth date of the forecast. A forecast is deleted from the system if the birth date plus the number of days since the birth date is greater than the value set in the Days to Keep Forecast parameter. This process occurs when either the [Run Batch Forecast Workbook](#) wizard is used to generate the forecast or when `PreGenerateForecast` is executed.

When you start the [Run Batch Forecast Workbook](#), one of the first steps that RDF takes is to remove forecasts older than the value given in the Days to Keep Forecasts parameter. RDF uses the value of that parameter in your local domain. However, because forecasts are registered globally, the Forecast Delete has a global effect. That is, if you set the parameter to some low value, for example, one day, all forecasts in all local domains older than one day are removed. This can adversely impact other users of the system if done without consideration.

When the `PreGenerate` utility is used, typically for automated batch production of forecasts, RDF uses the value for the Days to Keep Forecasts parameter in the first local domain. Again, if this is set to a value that is not expected by the other users, it can cause unintended disruption.

Refer to the *Oracle Retail Demand Forecasting Implementation Guide* for more information on `PreGenerateForecast`.

Default Keep Last Changes

This field is a list from which you select the default change policy for forecast items. Valid values are:

Field	Description
Keep Last Changes (None)	There are no changes that are introduced into the adjusted forecast. The adjusted forecast is equal to the system forecast.
Keep Last Changes (Total)	Considers both the Last System Forecast and the Last Approved Forecast in determining change policy, but only uses the Last Approved Forecast to calculate the adjustment value. For each item/week combination that is forecast, RDF automatically introduces the same quantity that was approved in the Last Approved Forecast into the change only if that quantity differed from that in the Last System Forecast. If the quantities are the same, RDF introduces the current system-generated forecast into the adjusted forecast.
Keep Last Changes (Diff)	Considers both the Last System Forecast and the Last Approved Forecast in determining change policy and adjustment value. For each item/week combination that is forecast, RDF determines the difference between the Last System Forecast and the Last Approved Forecast. This difference (positive or negative) is then added to the current system forecast to update the adjusted forecast.
Keep Last Changes (Ratio)	Considers both the Last System Forecast and the Last Approved Forecast in determining change policy and adjustment value. For each item/week combination that is forecast, RDF determines the ratio between the Last System Forecast and the Last Approved Forecast. This ratio is expressed as a percentage. This same percentage is used to update the adjusted forecast.

Generate Baseline Forecasts

A check should be indicated in this field (set to **true**) if the baseline forecast is to be generated to be viewed in any task. This parameter should be set if the level is used for Causal forecasting and the baseline is needed for analysis purposes.

Generate Cumulative Interval

A check in this field (set to **true**) specifies whether you want RDF to generate cumulative intervals (this is similar to cumulative standard deviations) during the forecast generation process. Cumulative Intervals are a running total of Intervals and are typically required when RDF is integrated with the Oracle Retail Merchandising System. If you do not need cumulative intervals, you can eliminate excess processing time and save disk space by clearing the check box. The calculated cumulative intervals can be viewed within the [Forecast Approval Workbook](#).

Generate Intervals

A check in this field (set to **true**) indicates that intervals (similar to Standard Deviations) should be stored as part of the batch forecast process. Intervals can be displayed in the [Forecast Approval Workbook](#). If you do not need intervals, excess processing time and disk space may be eliminated by clearing the check box. For many forecasting methods, intervals are calculated as standard deviation, but for Simple, Holt, and Winters the calculation is more complex. Intervals are not exported.

Generate Methods

A check in this field (set to **true**) indicates that when an ES forecast method is used, the chosen forecast method for each fitted time series should be stored. The chosen method can be displayed in the [Forecast Approval Workbook](#).

Generate Parameters

A check in this field (set to **true**) indicates that the alpha, level, and trend parameters for each fitted time series should be stored. These parameters can be displayed in the [Forecast Approval Workbook](#).

Interval Cap Lower Ratio

The value entered in this field multiplied by the forecast represents the lower bound of the confidence interval for a given period in the forecast horizon.

Interval Cap Upper Ratio

The value entered in this field multiplied by the forecast represents the upper bound of the confidence interval for a given period in the forecast horizon.

Item End Date Action

This parameter allows the option for items with end dates within the horizon to have zero demand applied to time series before or after the interim forecast is calculated. The two options are:

Option	Description
Apply 0 After Spreading	This is the default value. Spreading ratios are calculated for time series with no consideration made to the end date of an item. It is after the source forecast is spread to the final-level when zero is applied to the System Forecast.
Apply 0 to Interim	For items that have an end date within the forecast horizon, zero is applied to the Interim Forecast before the spreading ratios are calculated. This ensures that no units are allocated to the final-level for time series that have ended.

Like TS Duration (Periods)

The Like TS Duration is the number of periods of history required, after which RDF stops using the substitution method and starts using the system forecast generated by the forecast engine. A value must be entered in this field if using Like-item/Sister-Store functionality. When the difference between the forecast start date and the history start date is greater than the Like TS Duration, RDF starts using the system forecast for the item.

The Like TS Duration determines for how long RDF applies the Like-item/Sister-Store functionality. In general, the difference between history start date and the forecast start date is compared to the Like TS Duration. If the difference is larger than the Like TS Duration, it means that the new item has been selling for a period long enough, such that RDF can create the forecast based on its own history.

The forecast start date has an override, which is usually used whenever the forecast start date is checked.

However, in the case of Like-item/Sister-Store, if the override references a date prior to the forecast start date, the forecast start date - instead of the override - is used when calculating the difference and comparing it to the Like TS Duration.

The reason behind this exception is that the override does not change unless it is initiated by a user. Hence, the difference between forecast start date override and history start date is always the same, so the Like-item functionality is always in effect.

If the override is ignored, the forecast start date is advancing every period, so the difference between it and history start date is always growing. Eventually it will become larger than the Like TS Duration, and thus turning off the Like-item functionality for the item.

Store Interim Forecast

A check should be placed in this field (set to **true**) if the interim forecast is stored. The Interim Forecast is the forecast generated at the Final Level. This forecast is used as the Source Data within Curve to generate the profile (spreading ratios) for spreading the source-level forecast to the final-level. The interim forecast should only be stored if it is necessary for any analysis purposes.

Updating Last Week Forecast

This field is a list from which you can select the method for updating the Approved Forecast for the last specified number of weeks of the forecast horizon. This option is valid only if the Approval Method Override (set in the [Forecast Administration Workbook](#)) is set to Manual or Approve by alert, and the alert was rejected. This parameter is used with the Updating Last Week Forecast Number of Weeks parameter.

Method	Description
No Change	When using this method, the last week in the forecast horizon does not have an Approved Forecast value. The number of weeks is determined by the value set in the Updating Last Week Forecast Number of Weeks parameter.
Replicate	When using this method the last week in the forecast horizon is forecast using the Approved Forecast for the week prior to this time period. To determine the appropriate forecast time period, the value set in Updating Last Week Forecast Number of Weeks is subtracted from the Forecast Length. For example, if your Forecast Length is set to 52 weeks and Updating Last Week Forecast Number of Weeks is set to 20, week 32's Approved Forecast is copied into the Approved Forecast for the next 20 weeks.
Use Forecast	When using this method, the System Forecast for the last weeks in the forecast horizon is approved.

Updating Last Week Forecast Number of Weeks

The Approved Forecast for the last weeks in the forecast horizon is updated using the method specified from the Updating Last Week Forecast list.

Final and Source Level Parameters View - Advanced Settings

The Final and Source Level view allows you to set the advanced parameters that are common to both the final and source-level forecasts.

Figure 5–7 Final and Source Level Parameters View

	6 - itm/str	7 -	8 -
Default Bayesian Alpha	1.00	1.00	1.00
Bayesian Cap Ratio	1.50	1.50	1.50
DD Duration (Periods)	0	0	0
Max Alpha (Profile)	1.00	1.00	1.00
Max Alpha (Simple Holt)	1.00	1.00	1.00
Max Alpha (Winters)	1.00	1.00	1.00
Crostons Min Gaps	5	5	5
Holt Min Hist (Periods)	13	13	13
Winters Min Hist (Periods)	104	104	104
Trend Damping Factor	0.50	0.50	0.50
Fall Back Method	Default	Default	Default
Seasonal Smooth Index	0.80	0.80	0.80
Default Winters Mode	Oracle Winters	Oracle Winters	Oracle Winters
Default Moving Average Window Length	0	0	0
Max Gamma (Winters)	0.20	0.20	0.20
Max Gamma (Holt)	0.20	0.20	0.20
Deseasonalized Demand Array			

Measures: Source Level View - Advanced Settings

The Source Level View - Advanced Settings contains the following measures:

Bayesian Alpha (range (0, infinity))

When using the Bayesian forecasting method, historic data is combined with a known sales plan in creating the forecast. As POS data comes in, a Bayesian forecast is adjusted so that the sales magnitude is a weighted average between the original plan's scale and the scale reflected by known history. This parameter displays the value of alpha (the weighted combination parameter). An alpha value closer to one (or infinity) weights the sales plan more in creating the forecast, whereas alpha closer to zero weights the known history more. The default is one (1).

Bayesian Cap Ratio

The Bayesian Cap ratio is used to cap the resulting Bayesian forecast if it deviates significantly from the sale plan. The Bayesian Cap ratio is used as follows:

Example 5–1

```
If forecast I > Bayesian Cap ratio * Max value of Past Sales AND
forecast I > Bayesian Cap ratio * Max value of Past Plan AND
forecast I > Bayesian Cap ratio * Plan I
then
forecast I = Plan I
```

It defaults to a value of 1.5.

Croston's Min Gaps

The Croston's Min Gaps is the default minimum number of gaps between intermittent sales for the batch forecast to consider Croston's (Simple/IntermittentES) as a potential

AutoES forecasting method for a time series. If there are not enough gaps between sales in a given product's sales history, the Croston's model is not considered a valid candidate model. The system default is five minimum gaps between intermittent sales. The value must be set based on the calendar dimension of the level. For example, if the value is to be 5 weeks, the setting for a final-level at day is 35 (5x7days) and a source-level at week is 5.

DD Duration (weeks)

Used with Profile Based forecast method, the DD Duration is the number of weeks of history required after which the system stops using the DD (De-seasonalized Demand) approach and defaults to the normal Profile-Based method. The value must be set based on the calendar dimension of the level. For example, if the value is to be 10 weeks, the setting for a final-level at day is 70 (10x7days) and a source-level at week is 10.

Default Moving Average Window Length

Used with Moving Average forecast method, this is the Default number of data points in history used in the calculation of Moving Average. This parameter can be overwritten at item/location from the [Forecast Maintenance Workbook](#).

Fallback Method

Set this parameter only if the Fallback Method is to vary from the default Fallback Methods used by the selected forecasting algorithm. If the method selected as the Default Forecast Method or Forecast Method Override does not succeed for a time series, this method is used to calculate the forecast and the default Fallback Methods in the forecasting process is skipped entirely. The default Fallback Methods are as follows:

Table 5–1 Fallback Methods

Fallback Method	Steps
If either the Causal, Bayesian, or Profile-Based are selected as the Default Forecast Method or Forecast Method Override and the method does not fit the data.	<ol style="list-style-type: none"> 1. RDF attempts to fit SeasonalES 2. RDF attempts to fit TrendES 3. RDF attempts to fit Simple/IntermittentES
If the SeasonalES is selected as the Default Forecast Method or Forecast Method Override and neither Multiplicative Seasonal or Additive Seasonal fits the data.	<ol style="list-style-type: none"> 1. RDF attempts to fit TrendES 2. RDF attempts to fit Simple/IntermittentES
If either the Multiplicative Seasonal or Additive Seasonal are selected as the Default Forecast Method or Forecast Method Override and the method does not fit the data.	<ol style="list-style-type: none"> 1. RDF attempts to fit TrendES 2. RDF attempts to fit Simple/IntermittentES
If the TrendES is selected as the Default Forecast Method or Forecast Method Override and the method does not fit the data.	<ol style="list-style-type: none"> 1. RDF attempts to fit Simple/IntermittentES

Holt Min Hist (Periods)

Used with the AutoES forecast method, Holt Min Hist is the minimum number of periods of historical data necessary for the system to consider Holt (TrendES) as a potential forecasting method. RDF fits the given data to a variety of AutoES candidate models in an attempt to determine the best method; if not enough periods of data are available for a given item, Holt is not be considered as a valid option.

The system default is 13 periods. The value must be set based on the calendar dimension of the level. For example, if the value is to be 13 weeks, the setting for a final-level at day is 91 (13x7days) and a source-level at week is 13.

Max Alpha (Profile) (range [0,1])

In the Profile-based model-fitting procedure, alpha, which is a model parameter capturing the level, is determined by optimizing the fit over the de-seasonalized time series. The time series is de-seasonalized based on a seasonal profile. This field displays the maximum value (that is, cap value) of alpha allowed in the model-fitting process. An alpha cap value closer to one (1) allows more reactive models (alpha = 1, repeats the last data point), whereas alpha cap closer to zero (0) only allows less reactive models. The default is one (1).

Max Alpha (Simple, Holt) (range [0,1])

In the Simple or Holt (TrendES) model-fitting procedure, alpha (a model parameter capturing the level) is determined by optimizing the fit over the time series. This field displays the maximum value (cap value) of alpha allowed in the model-fitting process. An alpha cap value closer to one (1) allows more reactive models (alpha = 1, repeats the last data point), whereas alpha cap closer to zero (0) only allows less reactive models. The default is one (1).

Max Alpha (Winters) (range [0,1])

In the Winters (SeasonalES) model-fitting procedure, alpha (a model parameter capturing the level) is determined by optimizing the fit over the time series. This field displays the maximum value (cap value) of alpha allowed in the model-fitting process. An alpha cap value closer to one (1) allows more reactive models (alpha = 1, repeats the last data point), whereas alpha cap closer to zero (0) only allows less reactive models. The default is one (1).

Max Gamma (Holt) (range [0,1])

In the Holt (TrendES) model-fitting procedure, gamma (a model parameter capturing the trend) is determined by optimizing the fit over the time series. This field displays the maximum value (cap value) of gamma allowed in the model-fitting process.

Max Gamma (Winters) (range [0,1])

In the Winters (SeasonalES) model-fitting procedure, gamma (a model parameter capturing the trend) is determined by optimizing the fit over the time series. This field displays the maximum value (cap value) of gamma allowed in the model-fitting process.

Seasonal Smooth Index

This parameter is used in the calculation of seasonal index. The current default value used within forecasting is 0.80. Changes to this parameter impacts the value of seasonal index directly and impact the level indirectly. When seasonal smooth index is set to one (1), seasonal index is closer to the seasonal index of last year sales. When seasonal smooth index is set to zero (0), seasonal index is set to the initial seasonal indexes calculated from history.

This parameter is used when the Winters Mode is set to Oracle Winters.

If the Winters Mode is Winters Standard, Winters Responsive, or Oracle Winters Decomposition, this parameter is optimized and your input value is ignored.

Trend Damping Factor (range [0,1])

This parameter determines how reactive the forecast is to trending data. A value close to zero (0) is a high damping, while a value of one (1) implies no damping. The default is 0.5.

Winters Min Hist (Periods)

Used with the AutoES forecast method, the value in this field is the minimum number of periods of historical data necessary for Winters to be considered as a potential forecast method. If not enough years of data are available for a given time series, Winters is not used. The system default is two years of required history. The value must be set based on the calendar dimension of the level. For example, if the value is to be 104 weeks/2 years, the setting for a final-level at day is 728 (104 weeks x 7 days) and a source-level at week is 104.

Winters Mode

When any forecast method calls multiplicative or additive Winters, the system executes the Winters forecasting approach indicated by the Winters Mode.

The Winters forecasting approaches are listed in [Table 5-2](#).

Table 5-2 Winters Forecasting Approaches

Winters Forecasting Approaches	Description
Oracle Winters (default approach)	Current or default seasonal forecasting approach, which uses a combination of Winters approach and decomposition. Decomposition allows level and trend to be optimized independently while maintaining a seasonal curve.
Oracle Winters Decomposition	Like Oracle Winters seasonal forecasting approach, Oracle Winters Decomposition uses a combination of Winters approach and decomposition. Key differences between Oracle Winters and Oracle Winters Decomposition include: <ul style="list-style-type: none"> ■ Seasonal smoothing factor is optimized. ■ The optimization places more importance on recent sales than historical sales. ■ Seasonal indices are recalculated based on optimization.
Winters Standard	The Standard Winters uses a standard Winters model, with no additional calculations. All three smoothing parameters are optimized.
Winters Responsive	Responsive Winters, like Standard Winters does not use decomposition and all three smoothing parameters are optimized. This approach is recommended for forecasting items with significant trends in the more recent historic sales. The difference between Winters Responsive and Winters Standard is: <ul style="list-style-type: none"> ■ Winters Responsive optimization places more importance on recent sales than historical sales.

[Table 5-3](#) illustrates the key differences between the models.

Table 5-3 Comparison of Winters Forecasting Models

Forecasting Approach	Optimized to Recent Sales	Decomposition	Seasonal Smoothing Parameter Optimized
Oracle Winters	No	Yes	No
Oracle Winters Decomposition	Yes	Yes	Yes
Winters Standard	No	No	Yes
Winters Responsive	Yes	No	Yes

Winters Mode Impact on AutoES and Causal

The Winters Mode measure determines which Forecasting Approach to use within AutoES and Causal. Also, the errors that are used to calculate the BIC may be different for different Forecasting Approaches and could impact the choice of Forecast Method within AutoES and Causal.

It is recommended that you choose a Forecasting Approach that best suits the nature of your business. The default forecasting approach is Oracle Winters. For additional information on the forecasting approaches, refer to the [Forecasting Methods Available in RDF](#) section.

Note: If patching this change into a domain, in order to view this measure in the [Forecast Administration Workbook](#), you must add it to the [Final and Source Level Parameters View - Basic Settings](#) by selecting it from the Show/Hide dialog within the RPAS Client.

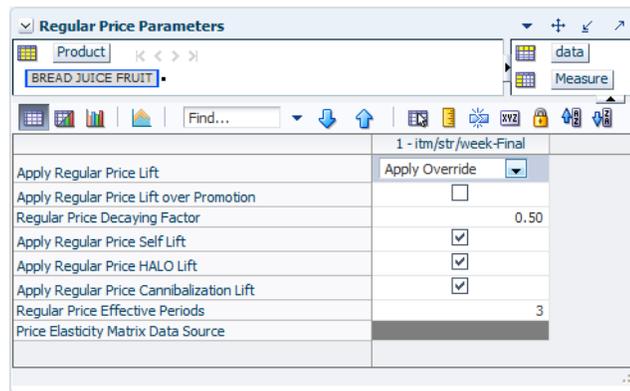
Deseasonalized Demand Array

Enter the name of the measure that serves as the deseasonalized data source in the profile-based forecasting method. If left blank, the system uses the forecast data source in combination with a seasonal profile to generate the deseasonalized demand.

Regular Price Parameters View- Advanced Settings

The Regular Price Parameters view allows you to set the parameters that allow you to incorporate the effects of regular price changes within the forecast.

Figure 5–8 Regular Price Parameters View



Measures: Regular Price Parameters View- Advanced Settings

The Regular Price Parameters View- Advanced Settings contains the following measures:

Apply Regular Price Lift

This parameter serves as a global setting, but can be overridden at lower levels.

Select a value from the list for this parameter:

- **Do Not Apply**— no regular price lifts are incorporated in the forecast
- **Apply System** — the effects calculated in RDF are incorporated in the forecast

- **Apply Override** — the lifts calculated in RDF are ignored and the override lifts are incorporated

There are three components to the Regular Price Lift:

- **Self Lift** — determined by the regular price self elasticity
- **Cannibalization** — determined by the regular price cannibalization cross elasticity
- **Halo** — determined by the regular price halo cross elasticity

Apply Regular Price Lift over Promotion

Is only effective if [Apply Regular Price Lift](#) is set to **Apply System** or **Apply Override**. Place a check in this parameter (set to **true**) if the application of regular price lifts is enabled for time periods that are promoted. If the parameter is not checked, regular price lifts will not be applied for periods with promotions. This parameter serves as a global setting, but can be overridden at lower levels.

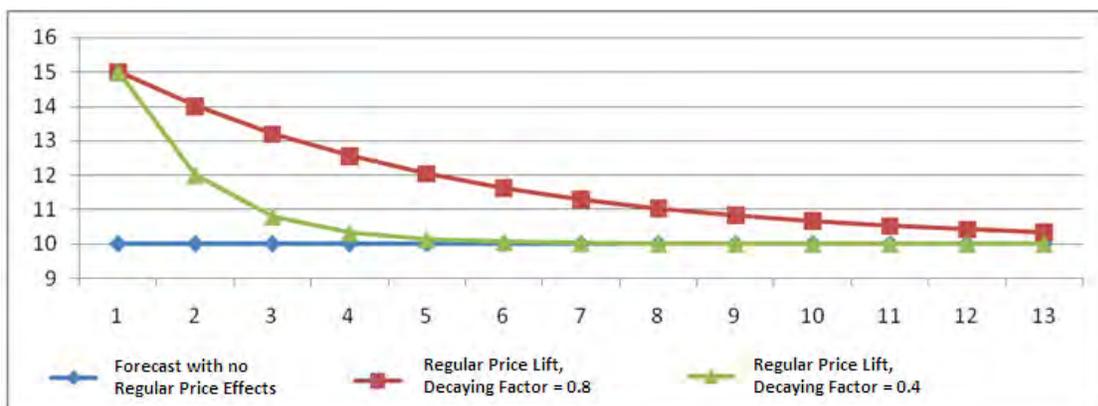
Regular Price Decaying Factor

This parameter determines how the magnitude of the price effects is applied over time.

A regular price change usually triggers a change in demand. However, the effect can be seen only for a few periods following the price change. An adjustable damping mechanism ensures that the regular price effect trends down as time goes by. The decaying factor is a number between zero (0) and one (1).

The closer the decaying factor is to zero (0), the faster the regular price effect decays.

Figure 5–9 *Decaying Factor Trend*



Apply Regular Price Self Lift

Is only effective if [Apply Regular Price Lift](#) is set to **Apply System** or **Apply Override**. Place a check in this parameter (set to **true**) if the application of regular price self lifts is enabled for the forecast level. This parameter serves as a global setting, but can be overridden at lower levels.

Apply Regular Price Halo Lift

Is only effective if [Apply Regular Price Lift](#) is set to **Apply System** or **Apply Override**. Place a check in this parameter (set to **true**) if the application of regular price Halo lifts is enabled for the forecast level. This parameter serves as a global setting, but can be overridden at lower levels.

Apply Regular Price Cannibalization Lift

Is only effective if [Apply Regular Price Lift](#) is set to **Apply System** or **Apply Override**. Place a check in this parameter (set to **true**) if the application of regular price cannibalization lifts is enabled for the forecast level. This parameter serves as a global setting, but can be overridden at lower levels.

Regular Price Effective Periods

Specifies the number of periods for which the regular price effect is applied.

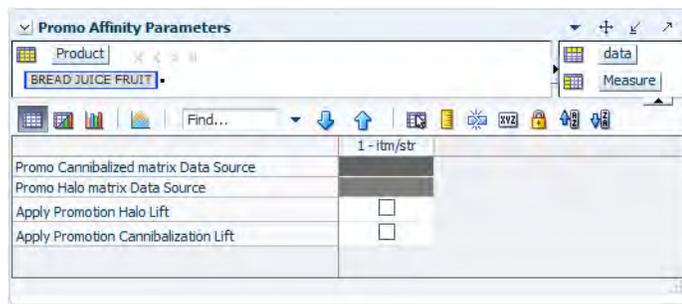
Price Elasticity Matrix Data Source

Specifies the name of the measures that stores the regular price elasticities, such as gammas (rdfgamma). The measure name is implementation specific.

Promo Affinity Parameters View

The Promo Affinity Parameters View allows you to set parameters to decide if cross promotional elasticities are incorporated in the forecast.

Figure 5–10 *Promo Affinity Parameters View*



Measures: Promo Affinity Parameters View

The Promo Affinity Parameters contains the following measures:

Apply Promotion Cannibalization Lift

Select this parameter (set to **true**) if the application of Cannibalization lifts is enabled for the forecast level. This parameter serves as a global setting, but can be overridden at lower-levels.

Apply Promotion Halo Lift

Select this parameter (set to **true**) if the application of Halo lifts is enabled for the forecast level. This parameter serves as a global setting, but can be overridden at lower-levels.

Promo Cannibalized Matrix Data Source

This parameter contains the name of the measure that is used to estimate the Cannibalization effects between items. The measure name is implementation specific.

Promo Halo Matrix Data Source

This parameter contains the name of the measure that is used to estimate the Halo effects between items. The measure name is implementation specific.

Forecast Approval Task

After the forecast is generated, the next steps in the forecasting process are analysis and approval. Approval of forecasts is required before the forecast data can be exported to other processes, such as replenishment programs. The Forecast Approval task allows you to view, analyze, adjust, and approve forecast results.

Some system forecasts may be set to be automatically approved by the system. The default approval method for items in a forecast is set in the [Forecast Administration Workbook](#), and these policies can be amended for individual product/location combinations in the [Forecast Maintenance Workbook](#). Any forecasts not set to Automatic Approval may require evaluation, adjustments, and ultimately approval before subsequent processes are executed.

You can view and analyze forecast data at multiple forecast levels (source-level and final-level) simultaneously. Revisions to and approvals of final-level forecast values are made on the appropriate views in the Forecast Approval task. The Forecast Approval task can contain these views:

Table 6–1 Views for the Forecast Approval Workbook

View	Description
Advanced Final Level View	This view displays all the regular price and promotional forecasting components
Final Level View	This view allows you to review final-level system-forecast quantities and make revisions to them if needed.
Approval View	This view allows you to specify the manual approval policy of forecasts by product and location.
Valid Forecast Run View	This view allows you to review the partition dimensions in which the generated forecast was run.
Source Level View	This view displays the system-generated source-level forecast and allows you to compare this data with final-level forecast values.
Final or Source System Parameters View	<p>The <i>Final System Parameters</i> view is only available if Generate System Parameters or Generate Methods is turned on in the Forecast Administration Workbook.</p> <p>The <i>Source System Parameters</i> view is only available if Generate System Parameters or Generate Methods is turned on in the Forecast Administration Workbook, and a Source Level was designated in Forecast Administration.</p>

When the Forecast Approval task is displayed, you may review the system-generated forecast and measures for any levels included in the task and make adjustments to forecast values at the final-level. Forecast values are overwritten in the Adjusted Forecast measure on the Final Forecast view. Approvals are made for each

product/location combination in the Approval Method measure of the Forecast Approval view.

After you complete your work, you can save the task using the Save function on the File menu. To update the master database with the approved forecast values, you must commit the task using the Commit Now or Commit Later option on the File menu. Once the task is committed, the forecast values are stored in the master database and can be used by other processes.

Forecast Approval Workbook

To build the Forecast Approval workbook, perform these steps:

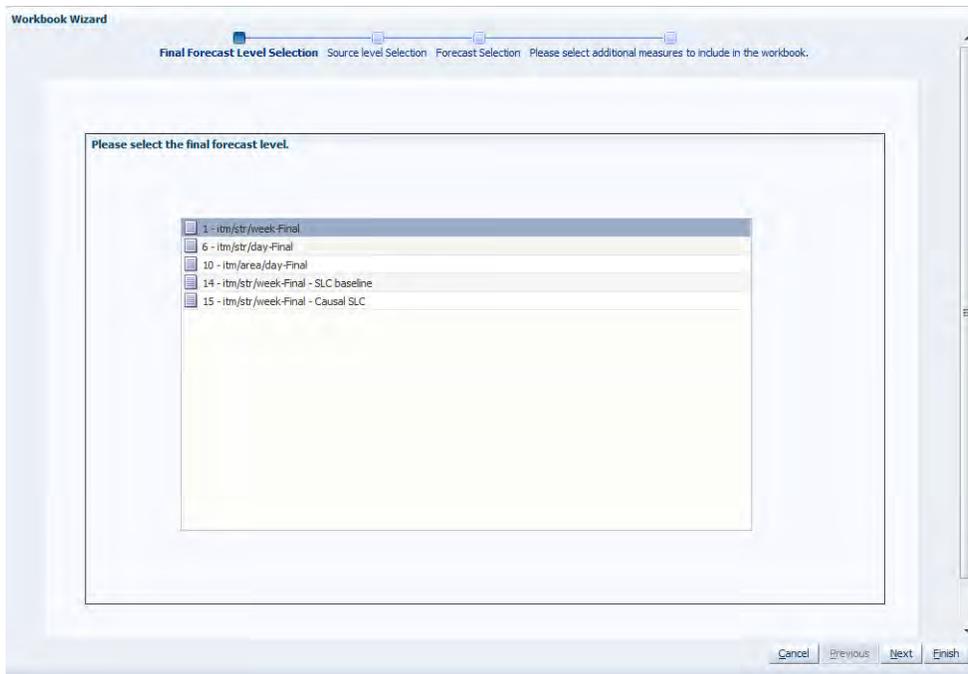
1. Click the New Workbook icon in the Forecast Approval task.

Figure 6–1 Forecast Approval Task



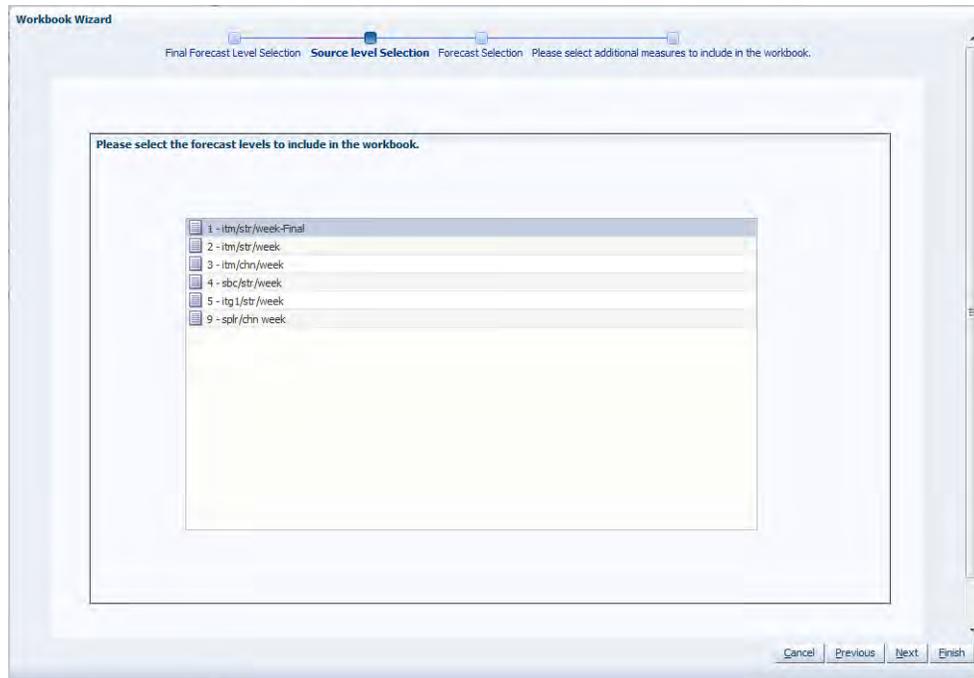
2. The Workbook wizard opens. Select the final forecast level you want to work with and click Next.

Figure 6–2 Workbook Wizard: Select Final Forecast Level



3. Select the forecast levels you want to work with and click **Next**. Select as many forecast levels as necessary for comparison.

Figure 6–3 Workbook Wizard: Select Forecast Levels

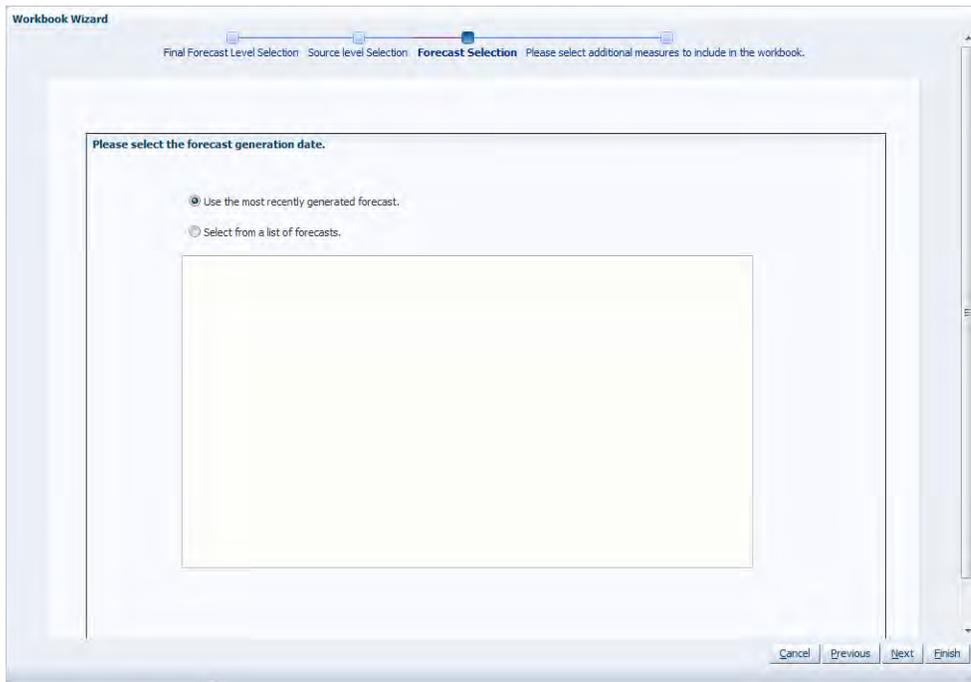


4. Select the forecast generation date you want to work with and click **Next**.
 - Select **Use the most recently generated forecast** to build a task containing the most recent forecast values.

Note: Use the most recently generated forecast - it must be selected if the task supports an AutoTask build.

- Click **Select from a list of forecast** to select from a list of previously generated forecasts stored in the system. Note that if this option is selected, the system automatically highlights the first forecast in the list, and if no further selections are made by you, this first forecast is used.
- If no selection is made by you, the birth date defaults to the latest generated generation ID.

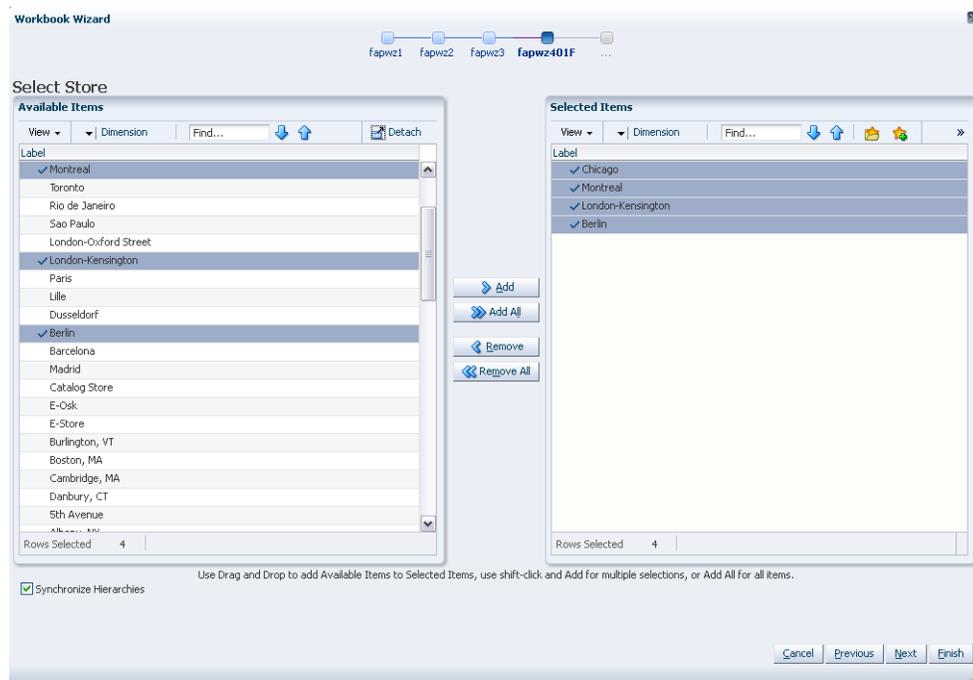
Figure 6–4 Workbook Wizard: Select Forecast Generation Date



5. Select the stores you want to work with and click **Next**.

Note: It is important to include all locations that are members of the location dimensions in the forecast levels to be analyzed. For example, if you select to view a forecast level that is defined at item/chain/week, you should include all locations that are members of the particular chain to be analyzed. It is recommended that Position Query functionality or selection from aggregate levels in the Location hierarchy is employed if the task is generated using AutoTask build.

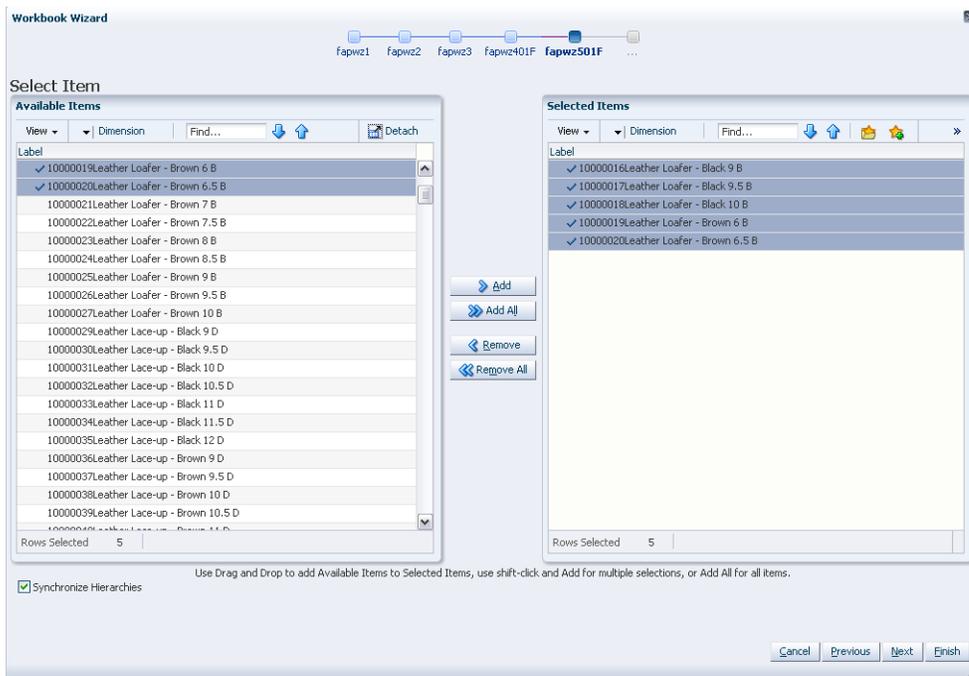
Figure 6–5 Workbook Wizard: Select Stores



6. Select the items you want to work with and click **Next**.

Note: It is important to include all products that are members of the Merchandise dimensions in the forecast levels to be analyzed. For example, if you select to view a forecast level that is defined at subclass/store/week, you must include all items that are members of the particular subclass to be analyzed. It is recommended that Position Query functionality or selection from aggregate levels in the Merchandise hierarchy is employed if the task is generated using AutoTask build.

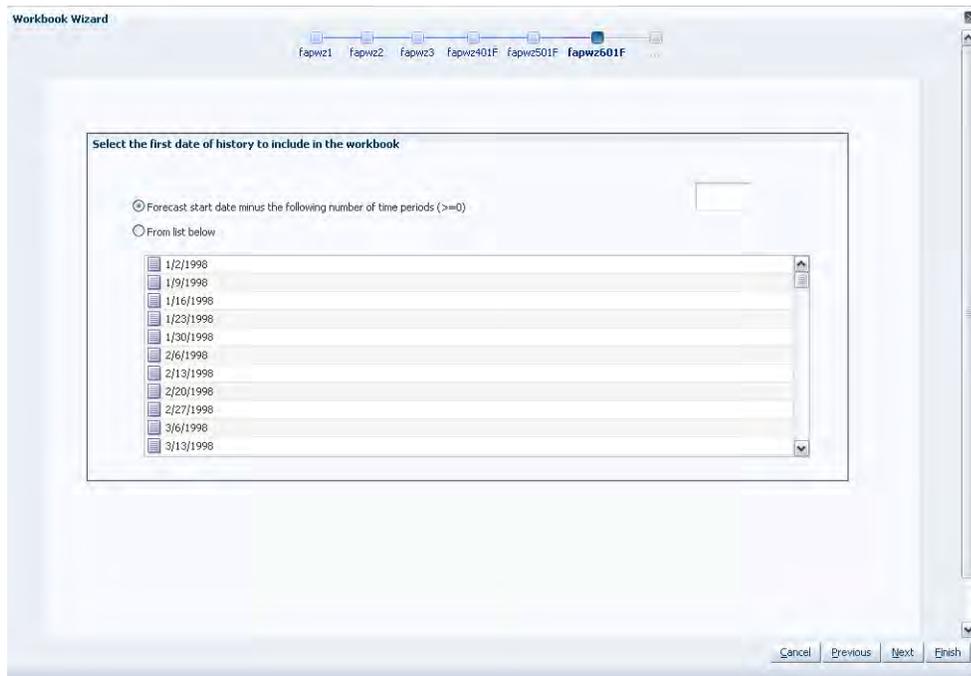
Figure 6–6 Workbook Wizard: Select Item



7. Select the first date of history you want to work with and click **Next**.

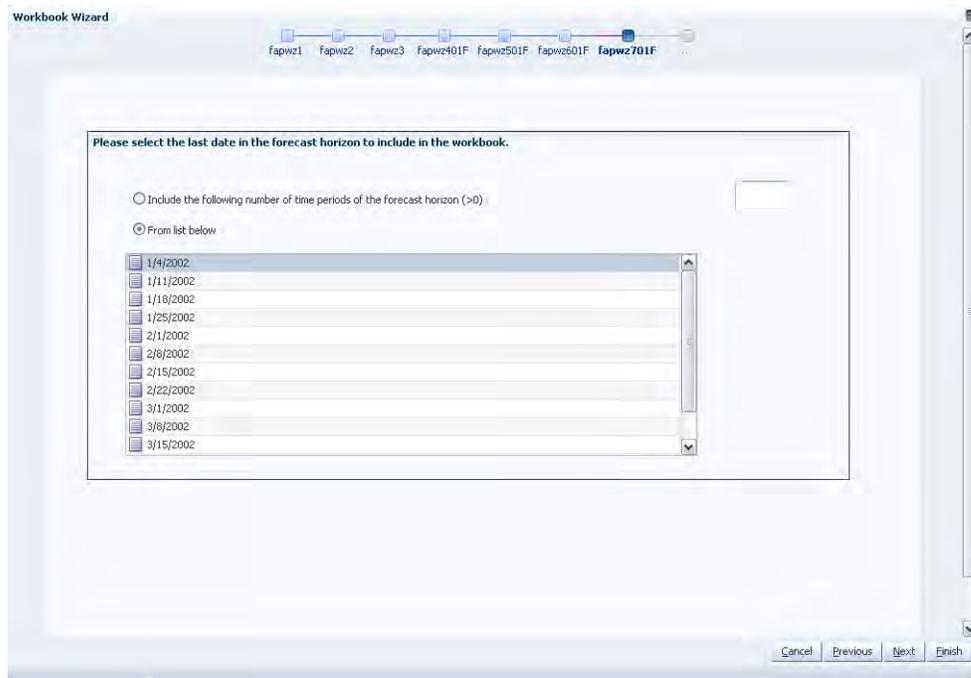
You may choose to either set the **Forecast Start Date minus the number of periods** or select the first date of history from the displayed list. It is recommended that the **Forecast Start Date minus the number of periods** is set if the task is generated using AutoTask build.

Note: If you make no selection, the system uses the forecast start date.

Figure 6–7 Workbook Wizard: History First Date

8. Select the last date in the forecast horizon you want to work with and click **Next**.

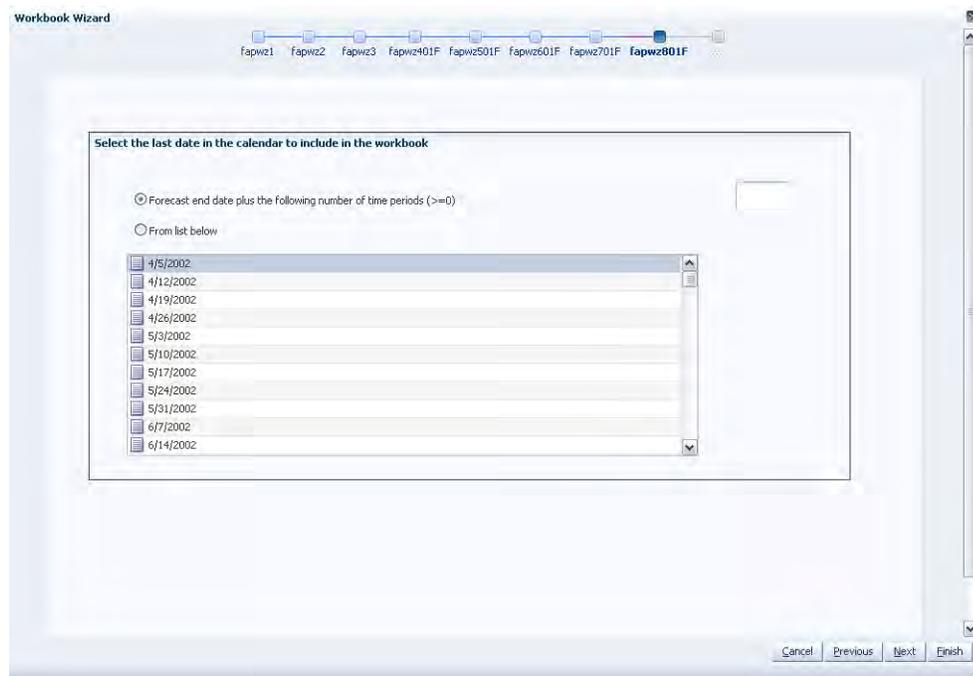
You may either choose to **Include the following number of time periods** of the forecast horizon or select the last date to include of the horizon from the displayed list. It is recommended that **Include the following number of time periods of the forecast horizon** is set if the task is generated using AutoTask build. The value entered must be within the forecast horizon. Click **Next**.

Figure 6–8 Workbook Wizard: Forecast Horizon Last Date

9. Select the last date in the calendar you want to work with and click **Next**.

You may choose to either set the **Forecast end date plus the following number of time periods** or select the last date to include of the post-horizon calendar from the displayed list. It is recommended that **Forecast end date plus the following number of time periods** is set if the task is generated using AutoTask build. Click **Next**.

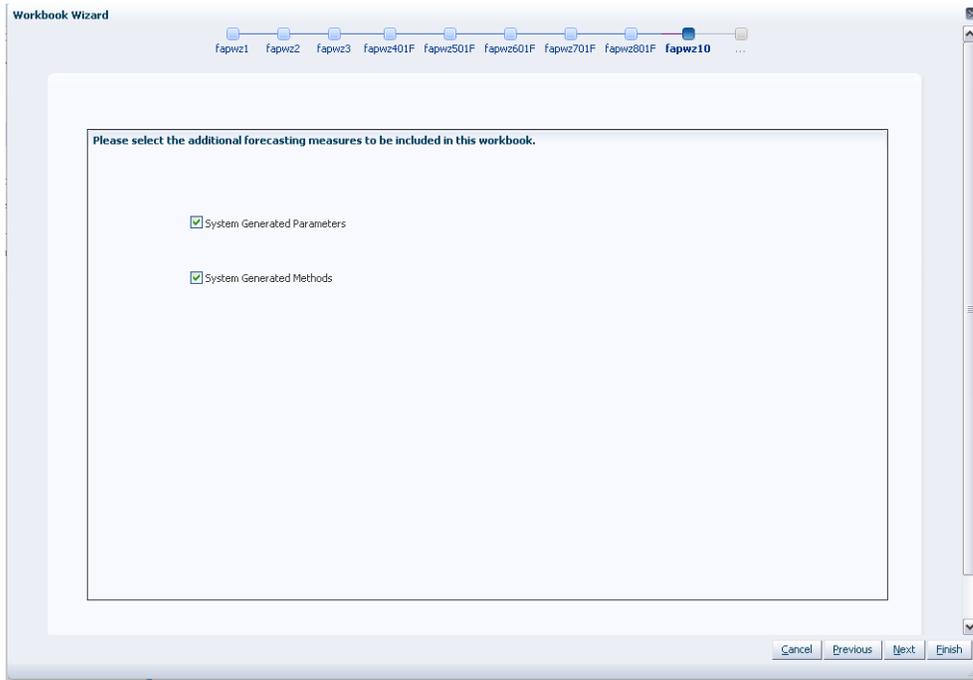
Note: If you make no selection, the system uses the last date in the forecast horizon chosen by you in the previous wizard window.

Figure 6–9 Workbook Wizard: Calendar Last Date

10. Select any additional forecast measures you want to work with in your task. The valid values of these measures may only be viewed if **Generate Intervals**, **Generate Cumulative Intervals**, **Generate Methods**, **Generate Parameters**, **Generate Baselines**, or **Store Interim Forecasts** were selected in the [Forecast Administration Workbook](#). Click **Finish**.

Note: If any of these parameters are included in the Forecast Approval task using AutoTask build and are later deselected for generation in the [Forecast Administration Workbook](#), AutoTask must be set again.

Figure 6–10 Workbook Wizard: Additional Forecasting Measures



The Forecast Approval workbook is built.

Advanced Final Level View

This view displays all the regular price and promotional forecasting components. These include Self Lifts, as well as Halo and Cannibalization Effects for Regular Price and Promotions. This information is read-only.

Figure 6–11 Advanced Final Level View

	12/21/2001	12/28/2001	1/4/2002	1/11/2002	1/18/2002	1/25/2002	2
Regular Price 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00	
Regular Price Cannibalization Lift 1 - itm/str/week-Final							
Regular Price HALO Lift 1 - itm/str/week-Final							
Regular Price Self Lift 1 - itm/str/week-Final							
Total Regular Price Lift Units 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00	
Promotion Self Lift Units 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00	
Promotion Halo Lift Units 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00	
Promotion Cannibalization Lift Units 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00	
Total Promotion Lift Units 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00	
Promotion Indicator 1 - itm/str/week-Final	<input type="checkbox"/>						

Measures: Advanced Final Level View

The Advanced Final Level view contains the following measures:

Promotion Cannibalization Lift Units

This measure displays the negative effect, measured in units, that the promotion of certain items has on this item.

Promotion Halo Lift Units

This measure displays the positive effect, measured in units, that the promotion of certain items has on this item.

Promotion Indicator

If this measure is checked (**true**), it indicates that a given item/location has any promotion active for a time period

Promotion Self Lift Units

This measure displays the effect that a promotion has on the sales units of an item.

Regular Price

This measure displays the price for which the item has sold at the given store. This is a time-phased measure, and the user can see how a change in price drives the regular price effects.

Regular Price Cannibalization Lift

This measure displays the negative effect, measured in units, that a price change of certain items has on this item.

Regular Price Halo Lift

This measure displays the positive effect, measured in units, that a price change of certain items has on this item.

Regular Price Self Lift

This measure displays the effect that a price change has on the sales units of an item.

Total Promotion Lift Units

This measure displays the total effect that a promotion can have on an item's sales units. It is the sum of Promotion Self Lift Units, Promotion Halo Lift Units, and Promotion Cannibalization Lift Units.

Note: If the measure, [Apply Promotional Lift Override](#), is selected in the [Advanced Final Parameter View](#), then RDF ignores the system calculated lifts and applies the [Aggregated Promotion Lift Override](#). It only applies the override in the week that has a promotion.

Total Regular Price Lift Units

This measure displays the total effect that a regular price change can have on an item's sales units. It is the sum of Regular Price Self Lift, Regular Price Halo Lift, and Regular Price Cannibalization Lift.

Final Level View

The Final Level view allows you to review the forecast quantities and make adjustments to forecasts if needed. The primary objective in the Forecast Approval task is to review and edit forecast values using the Adjusted Forecast field on the Final Level view and ultimately approve forecasts that have been user-adjusted or require manual approval.

Figure 6–12 Final Level View

	1/4/2002	1/11/2002	1/18/2002	1/25/2002	2/1/2002	2/8/2002	2/15/2002
Adjusted Forecast 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00	0.00
System Forecast 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Approved Forecast 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Last Approved Forecast 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Approved System Forecast 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00	0.00
History Data - Weekly Sales	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Total Regular Price Lift Units 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Total Promotion Lift Units 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00	0.00
System Baseline 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Adjusted Baseline 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Total Peak 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Adjusted Peak 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Forecast Interval	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Forecast Cumulative Interval	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Adjusted Cumulative Interval	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Interim Forecast	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Approved Cumulative Interval	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Measures: Final Level View

The Final Level view contains the following measures:

Adjusted Cumulative Interval

When changes are made to the Adjusted Forecast, the value of the Forecast Cumulative Interval is recalculated in this measure. The values in this measure are read-only. To view and store this measure:

- Generate Cumulative Intervals must be selected in the [Forecast Administration Workbook](#).
- Cumulative Intervals must be selected to be viewed in the Forecast Approval wizard.

Adjusted Forecast

The value in this field initially defaults to the System Forecast if the forecast is automatically approved by the system or through a Forecast Approval Alert. Otherwise, the value in the Adjusted Forecast is different than the System Forecast if:

- Adjusted Forecast has been updated by you.
- Keep Last Changes is set to Total, Difference or Ratio.
- Update Last Weeks Forecast is set to Replicate.

Changes to the Adjusted Forecast are automatically approved. The Approval view updates with the date of the adjustment and the name of the user who made the adjustment.

Note: Changes to the Adjusted Forecast for periods outside of the forecast horizon are not committed.

Edits to any non-committed values in the [Forecast Approval Workbook](#) are overwritten when data is refreshed.

Approved Cumulative Interval

The Approved Cumulative Interval is the cumulative interval that was approved at the time of the task build. The values contained in this measure are read-only. If changes are made to the Adjusted Forecast, the Approved Cumulative Interval reflects the recalculated values when the task is committed and data is refreshed.

Approved Forecast

The Approved Forecast is the forecast quantity that was approved at the time of the task build. The values contained in this measure are read-only. If changes are made to the Adjusted Forecast, then click **Calculate** to show the new value in the Approved Forecast.

Approved System Forecast

The Approved System Forecast is populated with the last System Forecast approved for a time series:

1. Approval Method set to Automatic Approval
2. Approval Method set to a Forecast Approval Alert and the alert is not triggered
3. Update Last Weeks Forecast is set to Use Forecast

The values in this measure are read-only.

Forecast Cumulative Interval

Cumulative Intervals are used in safety stock calculation within allocation and replenishment systems. This value is similar to a running total of the Forecast Interval and is read-only. To see this measure:

- Generate Cumulative Intervals must be selected in the [Forecast Administration Workbook](#).
- Cumulative Intervals must be selected to be viewed in the Forecast Approval wizard.

Forecast Interval

The Forecast Interval is calculated on the particular forecast region as capped standard deviation for some methods. It takes into consideration the system forecast for capping as well.

The forecast interval for any given week i is capped to a maximum of:

$$\text{Forecast}(i) * \text{upperCapRatio}$$

This cap is applied once the interval exceeds the system forecast, and it is applied for every value of the interval after that. The upperCapRatio (Interval Cap Upper Ratio) is a parameter set in the [Forecast Administration Workbook](#).

The forecast interval is capped to a minimum of $\text{Forecast}(i) * \text{lowerCapRatio}$ whenever the interval is less than this quantity. The `lowerCapRatio` (Interval Cap Lower Ratio) is another parameter set in the [Forecast Administration Workbook](#).

When forecasts are approved, the corresponding Intervals and Cumulative Interval measures are approved. To see this measure:

- Generate Intervals must be selected in the [Forecast Administration Workbook](#).
- Intervals must be selected to be viewed in the Forecast Approval wizard.

History Data

History Data (a read-only measure) is the sales data used to generate the forecast. This allows you to compare Actuals to forecast values. When the task is created, the Data Source measure is copied into History Data.

Interim Forecast

The Interim Forecast is the forecast generated at the final-level that is used as the Data Source in Curve to produce the Spreading Profile. This profile determines how the Source Forecast is spread down to the Final Forecast level. The values in this measure are read-only. To see this measure:

- Generate Interim Forecast must be selected in the [Forecast Administration Workbook](#).
- Interim Forecast must be selected to be viewed in the Forecast Approval wizard.

Last Approved Forecast

The Last Approved Forecast is the approved system forecast value when an approval occurred on a previous batch forecast for the time series. The values contained in this measure are read-only.

System Baseline

The System Baseline is a forecast generated on past sales data that contains no promotions (that is, normal demand given no causal effects). To see a generated Baseline Forecast:

- Promotional Forecasting must be implemented.
- Generate Baseline must be selected in the [Forecast Administration Workbook](#).
- The System Baseline must be selected to be viewed in the Forecast Approval wizard.

System Forecast

The System Forecast displays the system-generated forecast for the time series. The values contained in this field are read-only. This measure includes baseline as well as promotional lifts and regular price lifts, if applicable.

Total Promotion Lift Units

This measure displays the total effect that a promotion can have on an item's sales units. More information on the total lift and its components can be found in the Advanced Final Level view.

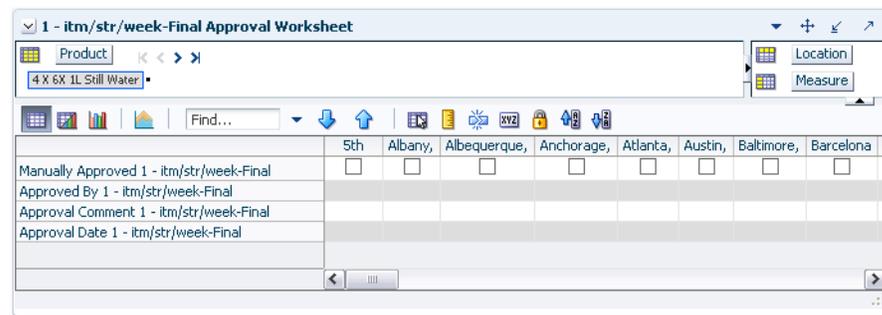
Total Regular Price Lift Units

This measure displays the total effect that a regular price change can have on an item's sales units. More information on the total lift and its components can be found in the Advanced Final Level view.

Approval View

The Forecast Approval view allows for non-adjusted System Forecast to be approved. This view can also be used to view the approval date of forecast values and to display the name of the user that manually approved forecast values for a given product/location combination. The default Approval Method is set in [Forecast Administration Workbook](#), and for product/location combinations that vary from the default, the [Forecast Maintenance Workbook](#) can be used.

Figure 6–13 Approval View



Measures: Approval View

The Approval view contains the following measures:

Approval Comment

Approval Comment is a field in which notes may be entered regarding the forecast values or any pertinent information for specified product/location combinations.

Approval Date

Approval Date is a read-only field that displays the date that the forecast quantity is approved either automatically during the batch forecast process or when changes are made to the Adjusted Forecast. This information is necessary for RDF to carry out any subsequent processes, such as replenishment procedures.

Approved By

Approved By is a read-only field that displays the name of the user to approve forecasts for an item/location. This field may be populated with *Sys* if the system was set to automatically approve forecasts during the batch forecast process.

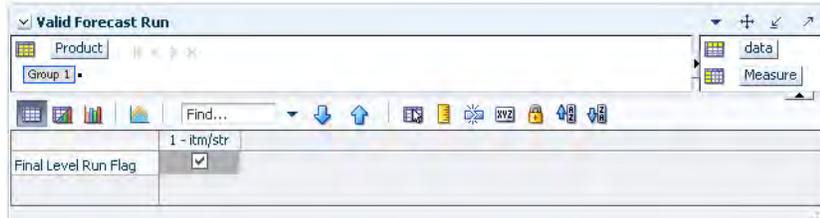
Manually Approved

Manually Approved is a Boolean flag. This field may be checked if you want to accept the System Forecast quantity for a time series that has yet to be approved. The flag is also activated when a change is made to the Adjusted Forecast and *Calculate* occurs. When this flag is activated, the time series for an item/location are approved and both the Approval Date and Approved By fields are updated.

Valid Forecast Run View

In a global domain environment, the Valid Forecast Run view allows you to identify which Local domains share the forecast birth date selected in the wizard when viewed in the Master domain. If in a Local domain, only the single position of the partition dimension is displayed.

Figure 6–14 Valid Forecast Run View



Measures: Valid Forecast Run View

The measure displayed on this view is viewed at the final-level/partition intersection. The Valid Forecast Run view includes the following standard measure:

Final Level Run Flag

The Final Level Run Flag is a read-only Boolean measure. A check (set to **true**) displayed in this field indicates which positions from the partition dimension had a successful forecast run for the birth date selected in the Forecast Approval wizard.

Note: Edits to the Adjusted Forecast can only be committed for partitions that have an activated Final Level Run Flag.

Source Level View

The Source Level view displays the system-generated source-level forecast. Final-level forecast values in the Final Level view can be viewed alongside and compared with their corresponding source-level forecasts.

Figure 6–15 Source Level View

	12/14/2001	12/21/2001	12/28/2001	1/4/2002	1/11/2002	1/18/2002	1/25/2002	2/1/2002	2/8/2002
System Forecast 3 - itm/chn/week	0.00	0.00	0.00	359.17	372.22	361.54	358.37	341.41	365.69
History Data - Weekly Sales	368.00	358.00	358.00	367.00	347.00	332.00	350.00	356.00	325.00
System Baseline 3 - itm/chn/week	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Measures: Source Level View

The Source Level view contains the following measures:

History Data

History Data (a read-only measure) is the sales data used to generate the forecast. This allows you to compare Actuals to forecast values. When the task is created, the Data Source measure is copied into History Data.

System Baseline

The System Baseline is a forecast generated on past sales data that contains no promotions (that is, normal demand given no causal effects). To see a generated Baseline Forecast:

- Promote must be implemented.
- Generate Baseline must be selected in the [Forecast Administration Workbook](#).
- The System Baseline must be selected to be viewed in the Forecast Approval wizard.

System Forecast

The System Forecast displays the system-generated forecast for the time series. The values contained in this field are read-only.

Final or Source System Parameters View

The Final or Source Parameters views allow you to view the alpha, level, and trend parameters for each fitted time series. These parameters are only available to be viewed if Generate Parameters or Generate Methods are available in the Advanced settings of the [Forecast Administration Workbook](#).

Figure 6–16 System Parameters View

	Bricks & Mortar	Catalog	E-osk	E-store	Retailer
Forecast Picked Method 3 - itm/chn/week	Multiplicative Seasonal	Multiplicative Seasonal	Multiplicative Seasonal	Multiplicative Seasonal	No Generated Forecast
System Generated Level 3 - itm/chn/week	391.84	44.95	4.46	13.33	0.00
System Generated Alpha 3 - itm/chn/week	0.05	0.05	0.00	0.01	0.00
System Generated Trend 3 - itm/chn/week	0.00	0.00	-0.00	0.00	0.00
System Generated Gamma 3 - itm/chn/week	0.00	0.00	0.05	0.04	0.00

Measures: Final or Source System Parameters View

The Final or Source Parameters view contains the following measures.

To see these measures:

- **Generate Parameters** must be selected in the [Forecast Administration Workbook](#).
- **System Generated Parameters** must be selected to be viewed in the Forecast Approval wizard.

Forecast Method Picked

This is the system-calculated alpha value (which is an internal optimization parameter that corresponds to the rate of decay of the weighting on the historical values) for the

corresponding product/location combination if the chosen method is one of the following methods: Simple, Holt, Additive Winters, Multiplicative Winters, and Profile Based.

System Generated Alpha

This is the system-calculated alpha value (which is an internal optimization parameter that corresponds to the rate of decay of the weighting on the historical values) for the corresponding product/location combination if the chosen method is one of the following methods: Simple, Holt, Additive Winters, Multiplicative Winters, and Profile Based.

System Generated Level

This is the system-calculated level (which is the constant baseline forecast) if the chosen method is one of the following methods: Simple, Holt, Additive Winters, Multiplicative Winters, and Seasonal Regression.

System Generated Gamma

This is the system-calculated gamma value for the corresponding product/location combination if the chosen method is one of the following methods: Simple, Holt, Additive Winters, Multiplicative Winters, and Profile Based.

System Generated Trend

This is the system-calculated trend (which is the rate of change of the baseline forecast with time) if the chosen method is one of the following methods: Holt, Additive Winters, Multiplicative Winters, and Seasonal Regression.

Forecast Delete Task

The Days to Keep Forecasts parameter (set in the [Forecast Administration Workbook](#)) supports the automatic deletion of old forecasts when the Run Batch wizard or PreGenerateForecast is executed. Occasionally, a user may need to manually delete a forecast. Some reasons might include:

- A forecast was run with the wrong source-levels selected.
- The forecast horizon was not properly set.
- Old forecasts need to be deleted to save space on the server.

The Delete Forecasts wizard guides you through the process of deleting unwanted forecasts from the system. Deletions of forecasts are permanent.

Forecast Delete Wizard

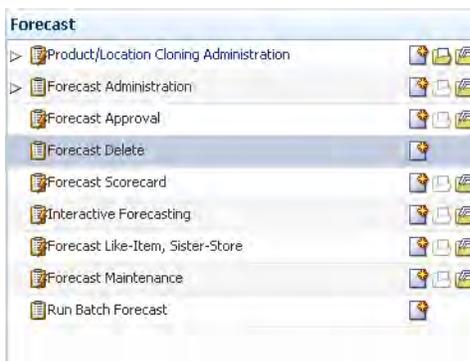
If the Forecast Delete wizard is used in a global domain environment. Deleting a forecast in the Master domain deletes the selected birth date within all domains. Deleting a forecast in the Local domain deletes the selected birth date only within the Local domain.

Note: The deletion of forecasts from RDF is permanent.

To use the Forecast Delete wizard, perform these steps:

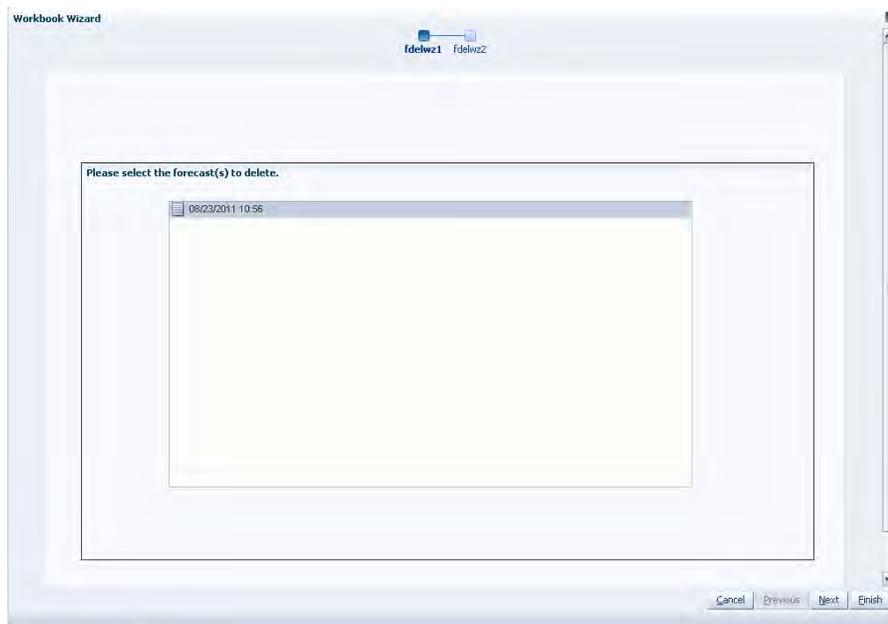
1. Click the New Workbook icon in the Forecast Delete task.

Figure 7-1 Forecast Delete Task



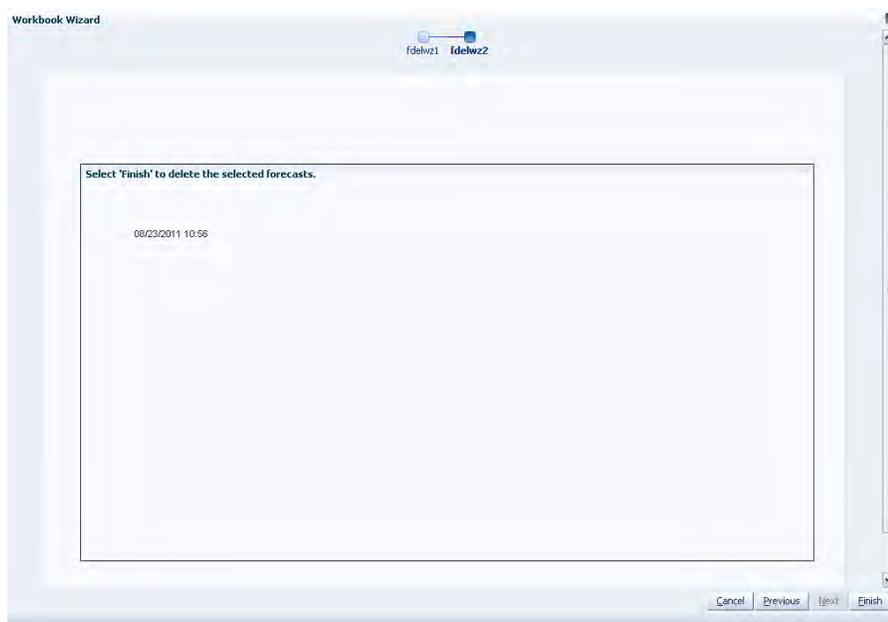
2. The Workbook wizard opens. Select the forecasts to delete and click **Next**.

Figure 7–2 Workbook Wizard: Select Forecasts to Delete



3. Select **Finish** to verify forecast deletion or **Cancel** to cancel deletion and exit the Forecast Delete wizard.

Figure 7–3 Workbook Wizard: Verify Delete



If **Finish** was selected on the final wizard window, the forecast is deleted.

Forecast Scorecard Task

This chapter describes the purpose and content of the Forecast Scorecard and the steps required in order to create and access this task. This chapter also discusses the two types of views contained in the Forecast Scorecard task, as well as the definitions of parameters that exist in each. Evaluating forecast accuracy through the use of error statistics is discussed, as is the process of comparing historical forecasts to actual sales data.

The purpose of the Forecast Scorecard is to monitor the accuracy of both system-generated forecasts and approved final forecasts. Once a forecast has been generated and actual point-of-sale data is received for the forecast period, statistical information can be reviewed to help you analyze the accuracy of forecasting models and methods.

The Forecast Scorecard consists of a wizard and two views.

Based on your selections in the wizard, the Forecast Scorecard provides statistical information and comparison data that allow you to monitor the accuracy of system-generated forecasts and final approved forecasts.

- The Final Error Measure view displays statistical information, such as mean error and root mean squared error, that reflects the accuracy of the forecast.
- The Final Actuals versus Forecasts view displays forecast results and actual point-of-sale values for specified product/location/calendar combinations.

Forecast Scorecard Workbook

To build the Forecast Scorecard workbook, perform these steps:

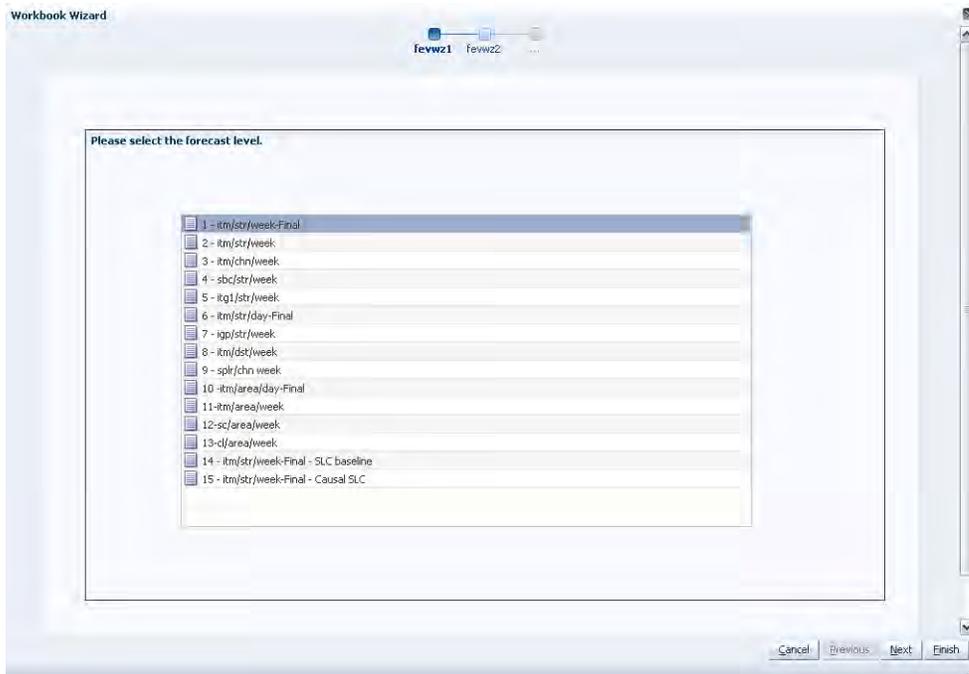
1. Click the New Workbook icon in the Forecast Scorecard task.

Figure 8–1 Forecast Scorecard Task

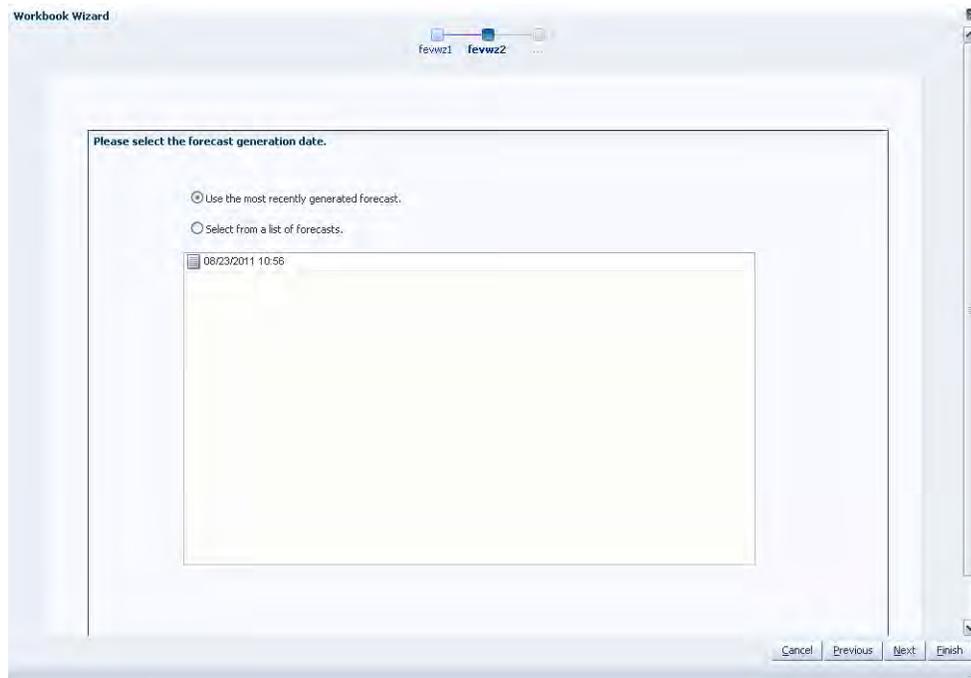


2. The Workbook wizard opens. Select the forecast level that you want to evaluate and click **Next**.

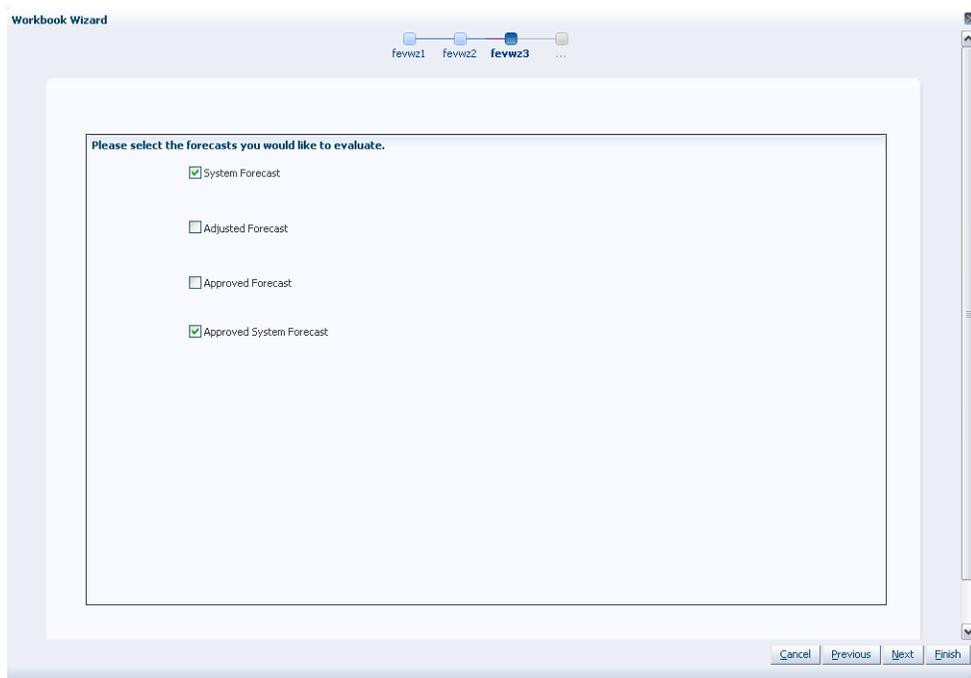
Figure 8–2 Workbook Wizard: Forecast Level



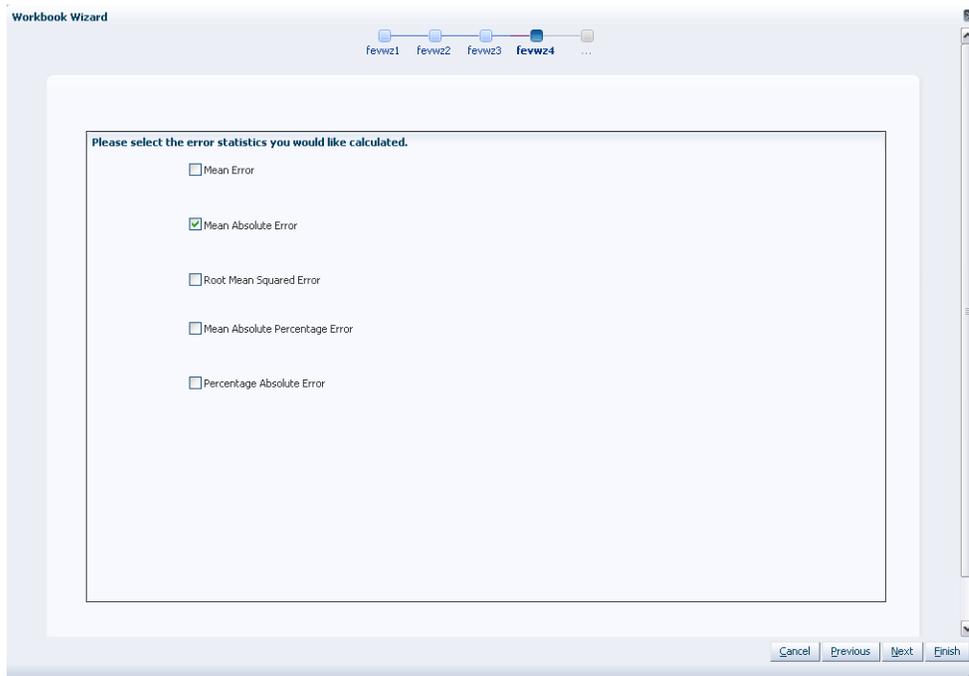
3. Select the generation date of the forecast you wish to review and click **Next**:
 - Click **Use the most recent generated forecast** to build a task containing the most recent forecast values.
 - Click **Select from a list of forecast** to select from a list of previously generated forecasts stored in the system.

Figure 8–3 Workbook Wizard: Generation Date

4. Select the measures associated with the forecast results to be evaluated and click **Next**.

Figure 8–4 Workbook Wizard: Select Error Statistics

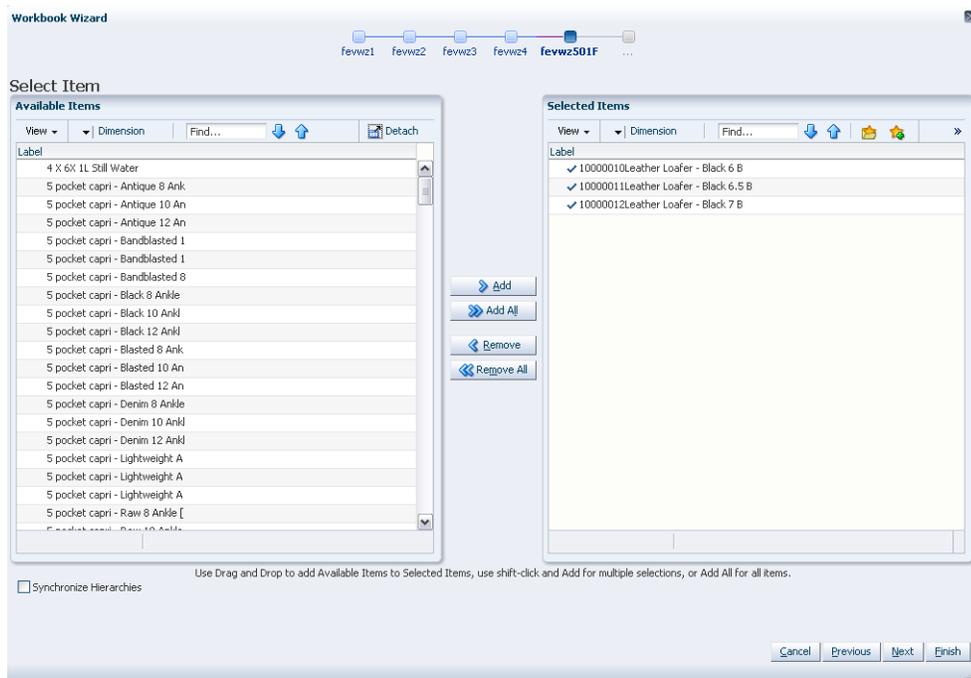
5. Check the boxes corresponding to the error statistics that you want calculated for the forecast data. and click **Next**.

Figure 8–5 Workbook Wizard: Calculate Error Statistics

6. Select the items you want to work with and click **Next**.

Note: It is important to include all products that are members of the Merchandise dimensions in the forecast levels to be analyzed. For example, if you select to view a forecast level that is defined at subclass/store/week, you must include all items that are members of the particular subclass to be analyzed. It is recommended that Position Query functionality or selection from aggregate levels in the Merchandise hierarchy is employed if the task supports an AutoTask build.

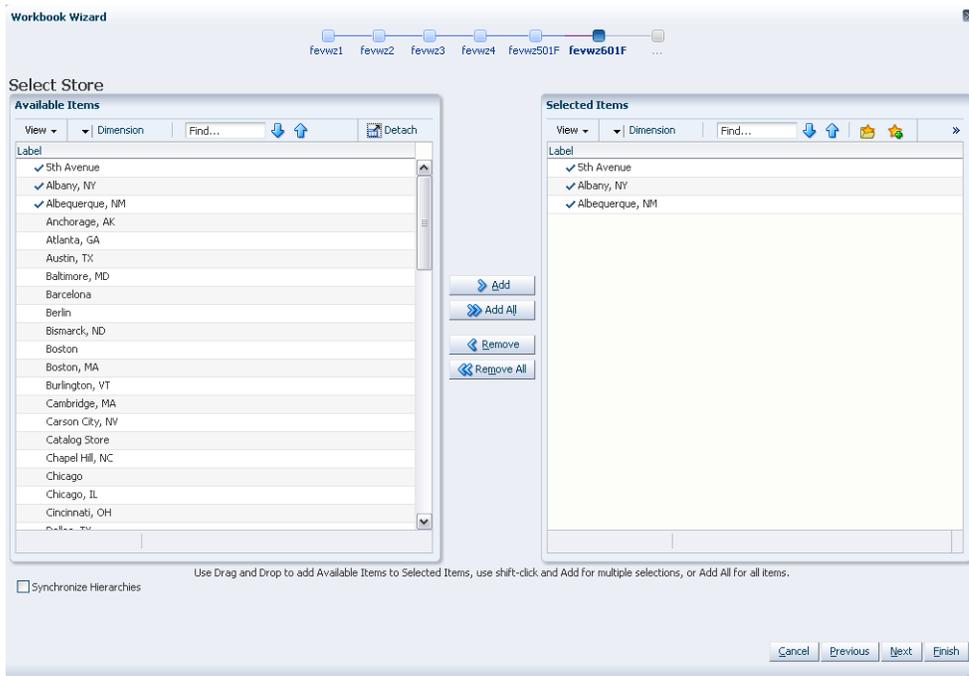
Figure 8–6 Workbook Wizard: Select Item



7. Select the specific locations that you want to view and click **Next**.

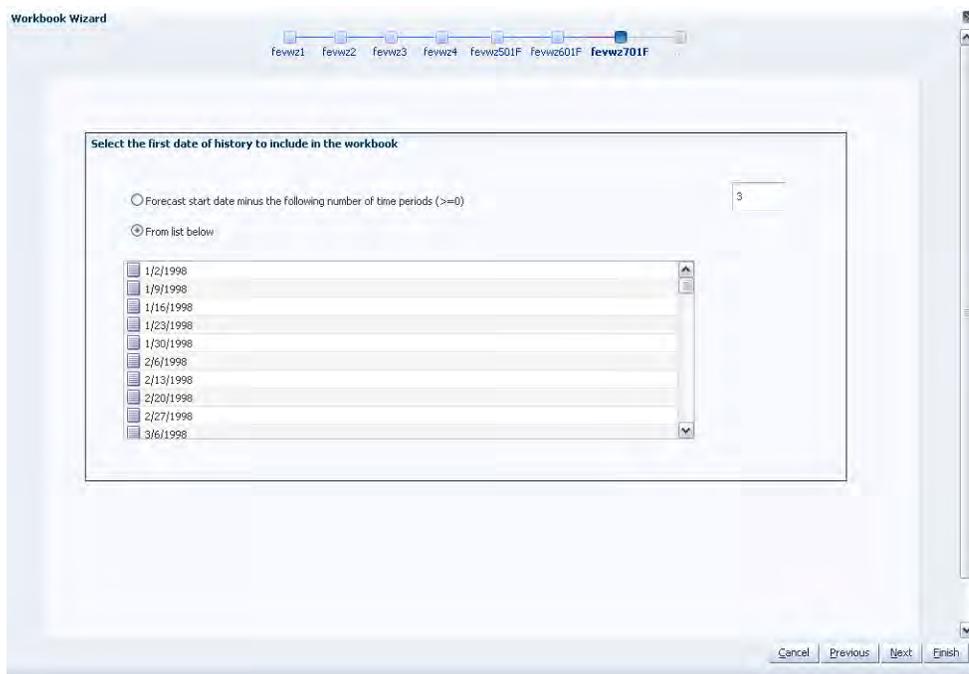
Note: It is important to include all locations that are members of the location dimensions in the forecast levels to be analyzed. For example, if you select to view a forecast level that is defined at item/chain/week, you should include all locations that are members of the particular chain to be analyzed. It is recommended that Position Query functionality or selection from aggregate levels in the location hierarchy is employed if the task supports an AutoTask build.

Figure 8–7 Workbook Wizard: Select Stores



8. Select the first date of history to include in the task. You may either choose to set the **Forecast Start Date minus the number of periods** or select the first date of history from the displayed list. It is recommended that the **Forecast Start Date minus the number of periods** is set if the task supports an AutoTask build. Make your selection and click **Next**.

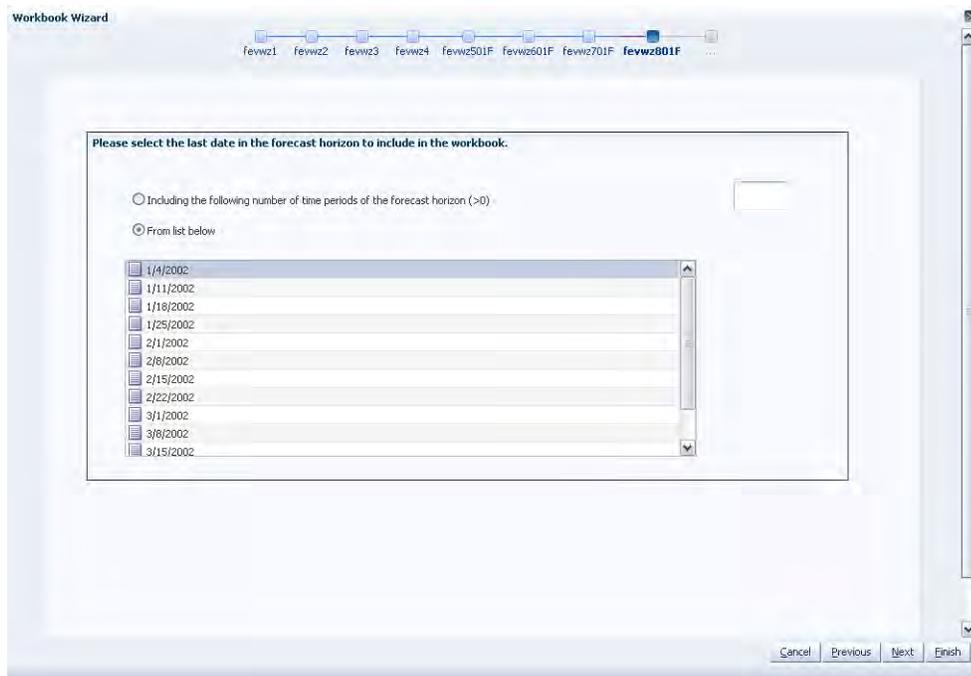
Figure 8–8 Workbook Wizard: Select First Date of History



9. Select the last date to include of the horizon from the displayed list and click **Next**.

Note: It is recommended that **Include the following number of time periods of the forecast horizon** is set if the task supports an AutoTask build.

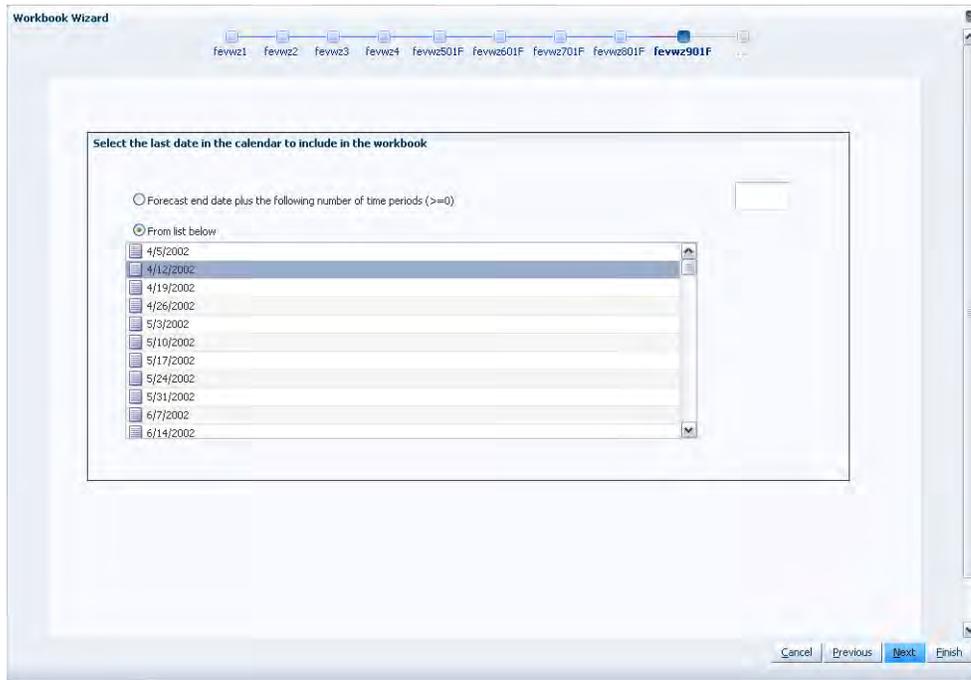
Figure 8–9 Workbook Wizard: Select Last Date of the Horizon



10. Select the last date in the calendar to include and click **Finish**.

Note: It is recommended that **Forecast end date plus the following number of time periods** is set if the task supports an AutoTask build.

Figure 8–10 Workbook Wizard: Select Last Date in the Calendar



The Forecast Scorecard workbook is built.

Error Measure View

The Final Error Measure view displays statistical information that reflects the accuracy of the forecast measures selected in the wizard process. You may need to evaluate a variety of such statistics to verify and compare forecast accuracy.

Given the situation, different levels of forecast accuracy can be useful. For example, in a situation with noisy data and no forecast previously available, a 200% error can be considered excellent and useful. In another situation with smooth data, if an old method of forecasting provided forecasts with a 10% error, a new method of forecasting with a 20% error would not be considered useful.

Different levels of forecasting accuracy are obtained at different levels of product aggregation. Item level forecasts with a 200% error can roll up into division-level forecasts with a 10% error. Therefore, the error measures are most useful when comparing different methods of forecasting or when looking at a particular model's accuracy over time. Accurate forecasts should lead to a reduction in inventory holding costs or to an increase in customer service levels; forecast errors should be converted into dollars saved.

Note: All error statistics are calculated over the first period in the forecast to the last period in the forecast for which forecast data is available.

Figure 8–11 Error Measure View

The screenshot shows a software window titled "1 - itm/str/week-Final Error Measure Worksheet". It features a navigation bar with "Product" and "Location" filters, a search bar, and a "Measure" button. Below the navigation bar is a table with the following data:

History Data for Forecast Window 1 - itm/str/week-Final	11.00
System Forecast 1 - itm/str/week-Final	9.37
System Forecast Mean Error 1 - itm/str/week-Final	0.54
Approved Forecast for Forecast Window 1 - itm/str/week-Final	
Interim Forecast 1 - itm/str/week-Final	0.00
Interim Forecast Mean Error 1 - itm/str/week-Final	3.67

Measures: Error Measure View

The Error Measure view contains the following measures:

History Data for Forecast Window

The sum of all actual sales from the first period in the forecast to the last period in the forecast for which point-of-sale data is available.

Forecast Measures

Your task may include one or all of the forecast measures: System Forecast, Adjusted Forecast, Approved Forecast, Approved System Forecast, or Interim Forecast. These measures are selected in the wizard process. This is the average value of these measures over the forecast horizon select in the wizard process.

Mean Error

The error of a forecast observation is the difference between the forecast value and the actual POS value. The Mean Error statistic is a measure of the average error over time. This is calculated by summing the errors for all observations, and then dividing by the number of observations to obtain the average. It measures forecast accuracy by calculating the error in units. Because a positive error in one period can cancel out a negative error in another period, this measure is useful when you are interested in how well the forecast predicts over the forecast horizon rather than on a period-to-period basis. Mean error is useful as a measure of forecast bias. A negative mean error suggests that overall the forecasting model overstates the forecast, while a positive mean error indicates forecasts that are generally too low.

Mean Absolute Error

The absolute error of a forecast observation is the absolute value of the difference between the forecast value and the actual POS value. The Mean Absolute Error statistic is a measure of the average absolute error. This is calculated by summing the absolute errors for all observations and then dividing by the number of observations to obtain the average. Mean Absolute Error gives you a better indication of how the forecast performed period by period because the absolute value function ensures that negative errors in one period are not canceled out by positive errors in another. Mean Absolute Error is most useful for comparing two forecast methods for the same series.

Root Mean Squared Error

This is the square root of the Mean Squared Error. The Root Mean Squared Error is one of the most commonly used measures of forecast accuracy because of its similarity to the basic statistical concept of a standard deviation. It evaluates the magnitude of

errors in a forecast on a period-by-period basis, and it is best used to compare alternative forecasting models for a given series.

Mean Absolute Percentage Error

The percentage error of a forecast observation is the difference between the actual POS value and the forecast value, divided by the actual POS value. The result of this calculation expresses the forecast error as a percentage of the actual value. The Mean Absolute Percentage Error statistic measures forecast accuracy by taking the average of the sum of the absolute values of the percentage error calculations across all observations. This method is useful when comparing the accuracy of forecasts for different volume products (it normalizes error by volume).

Percentage Absolute Error

The absolute error of a forecast observation is the absolute value of the difference between the forecast value and the actual POS value. The Percentage Absolute Error statistic measures forecast accuracy by calculating the total absolute error as a percentage of the total actual POS. It is calculated by summing the absolute errors for all observations, dividing this value by the absolute value of the sum of all Actuals, and dividing the result by the number of observations in the series. Finally, multiply the total by 100 to obtain a percentage result.

Actuals versus Forecasts View

This view displays forecast results and actual point-of-sale values for each product, location, and time period specified in the Forecast Scorecard wizard. This view allows you to compare the results of both system-generated forecasts and final approved forecasts to actual sales quantities.

Figure 8–12 Actuals versus Forecast View

	1/4/2002	1/11/2002	1/18/2002	4/5/2002	4/12/2002
History Data - Weekly Sales	3.00	4.00	4.00	4.00	4.00
System Forecast 1 - itm/str/week-Final	3.10	3.18	3.09	0.00	0.00
Interim Forecast 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00
Actual - System Forecast 1 - itm/str/week-Final	-0.10	0.82	0.91	4.00	4.00
Actual - Interim Forecast 1 - itm/str/week-Final	3.00	4.00	4.00	4.00	4.00

Measures: Actuals versus Forecasts View

The Actuals versus Forecasts view may contain the following measures:

Actuals

These Actuals do not include: System Forecast, Adjusted Forecast, Approved Forecast, Approved System Forecast, or Interim Forecast.

This displays the difference between the actual sales quantities and the committed forecast values for the forecast measures selected during the wizard process. Negative values indicate that the final forecast exceeded actual sales.

History Data

History Data displays the actual point-of-sale quantities for the product, location, and calendar combinations displayed.

Forecast Measures

Your task may include one or all of these forecast measures: System Forecast, Adjusted Forecast, Approved Forecast, Approved System Forecast, or Interim Forecast. These measures are selected in the wizard process. These measures display the forecast quantity for the product, location, and calendar combinations included in the task.

Interactive Forecasting Task

The Interactive Forecasting task is a forecast simulation tool that allows you to make changes to forecast parameters and see the results without having to wait for the batch run. In this workbook, you can edit various forecast parameters including sales history and forecast method. A new forecast is produced based on the changed parameters. In addition to forecasts, the Interactive Forecasting workbook can also generate fit in historical region and the system picked model if an ES forecast method is used.

Improve Performance for the Interactive Forecasting Workbook

Interactive Forecasting assumes that the historical data is contained in the measure *posxlsx*. All forecasts generated by interactive forecasting are based on this data.

Performance may be improved by utilizing a mask measure. Interactive forecasting uses the measures `mask<forecast level>xb`; for example *mask01xb*. Here, `<forecast level>` corresponds to the forecast level chosen in the wizard when opening the Interactive Forecasting workbook. The mask is a boolean measure. Only cells with the value **true** are included in the forecast. This is useful when desiring to focus on only a few product/locations.

The Interactive Forecast method is intended to be a quick way to view results produced by different forecast methods. It allows you to gauge, in a qualitative way, the suitability of a particular forecast method for a given set of sales history data. Even though it uses the exact same forecast methods as the full forecast, only the parameters listed previously will affect the results. No other user-input values are used. Users wishing to vary a larger number of parameters are encouraged to use the regular forecast functionality.

Interactive Forecasting Workbook

To build the Interactive Forecasting workbook, perform these steps:

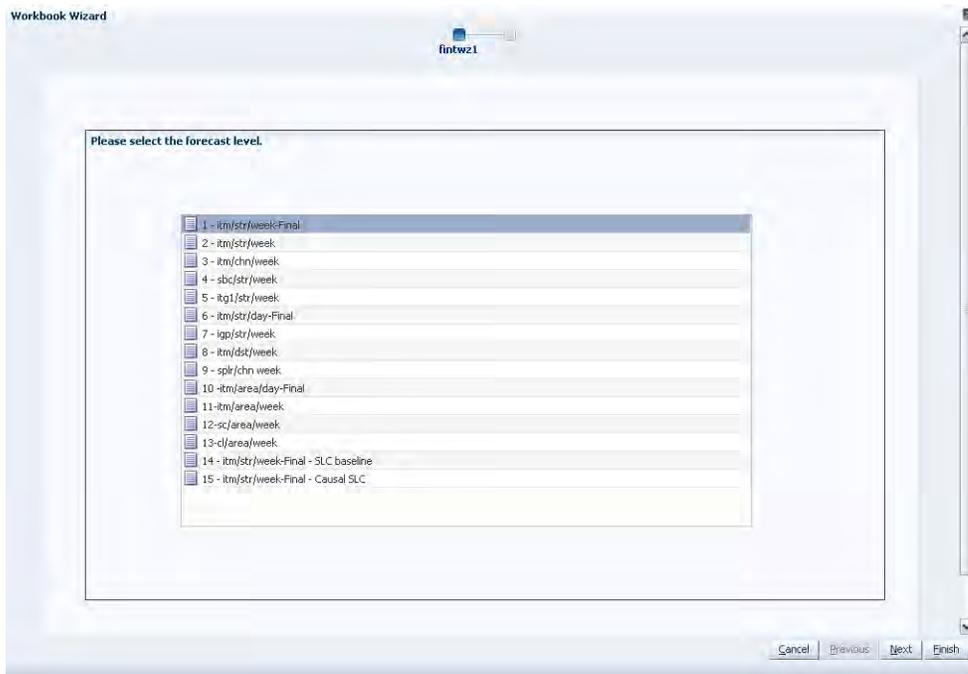
1. Click the New Workbook icon in the Interactive Forecasting task.

Figure 9–1 Interactive Forecasting Task

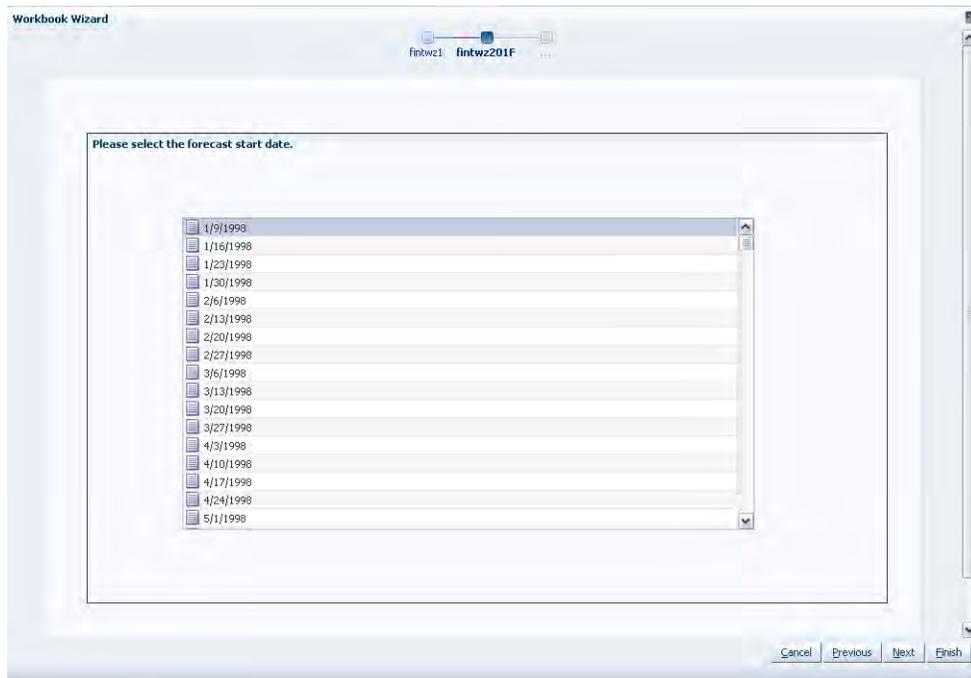


2. The Workbook wizard opens. Select the forecast level that you want to evaluate and click **Next**.

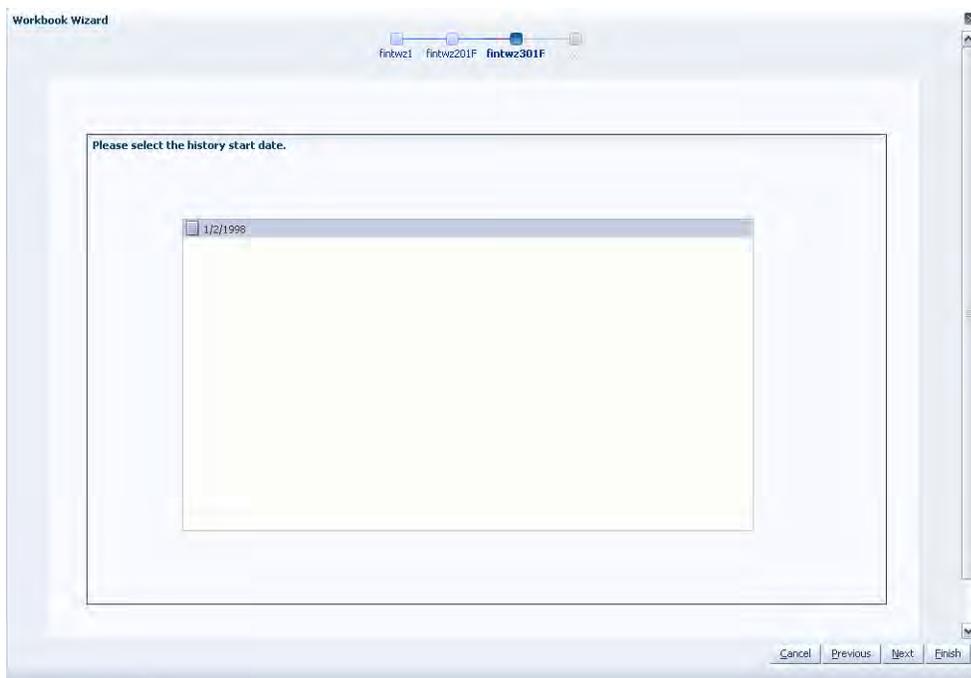
Figure 9–2 Workbook Wizard: Forecast Level



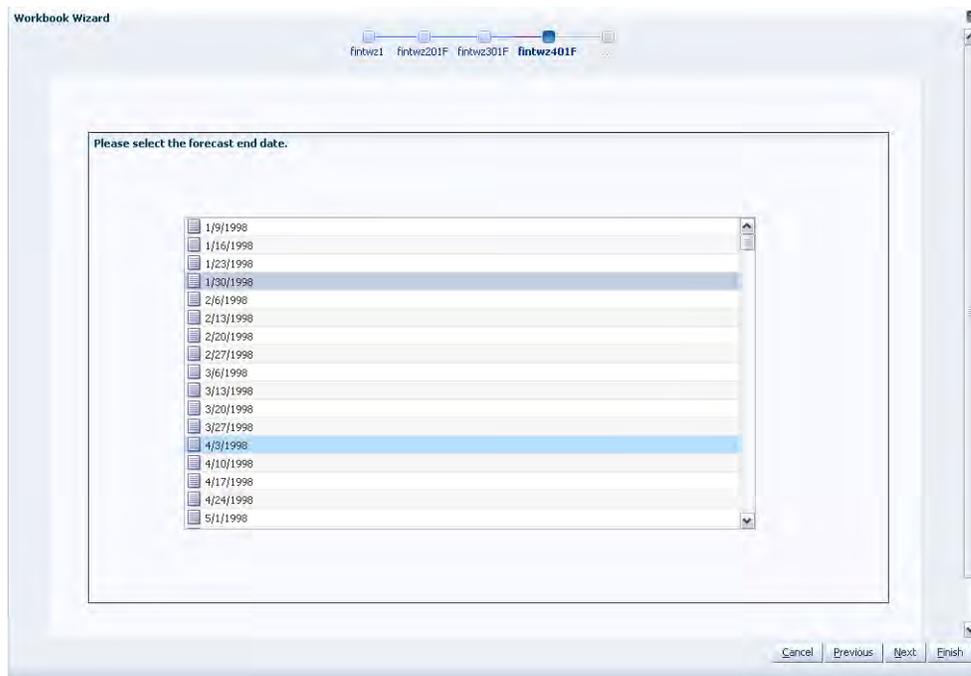
3. Select the forecast start date you wish to review and click **Next**:

Figure 9–3 Workbook Wizard: Forecast Start Date

4. Select the history start date and click Next.

Figure 9–4 Workbook Wizard: History Start Date

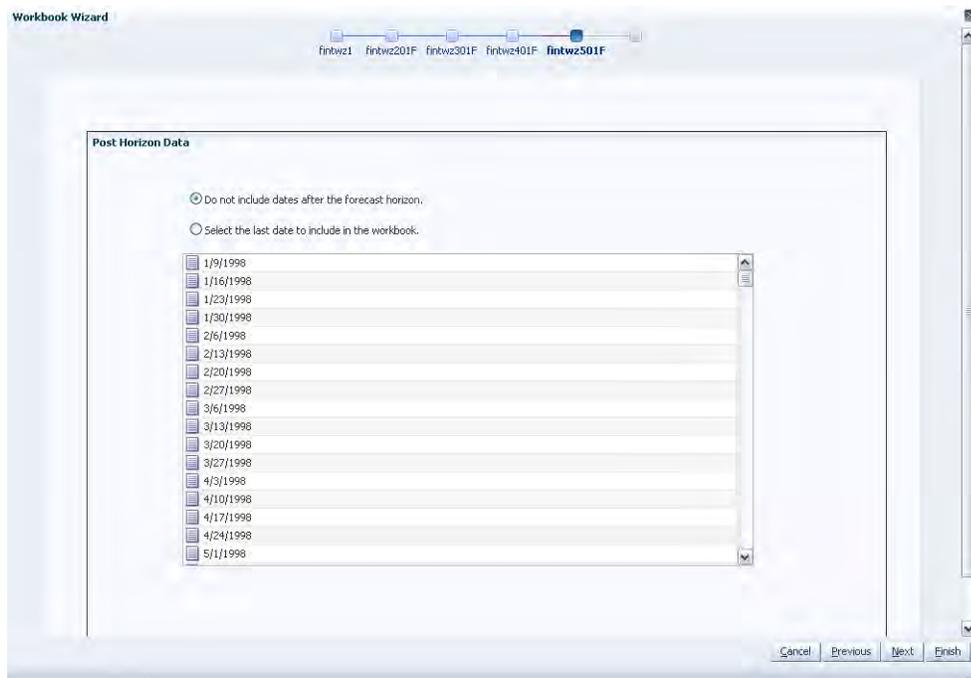
5. Select the end date for the forecast and click Next.

Figure 9–5 Workbook Wizard: Forecast End Date

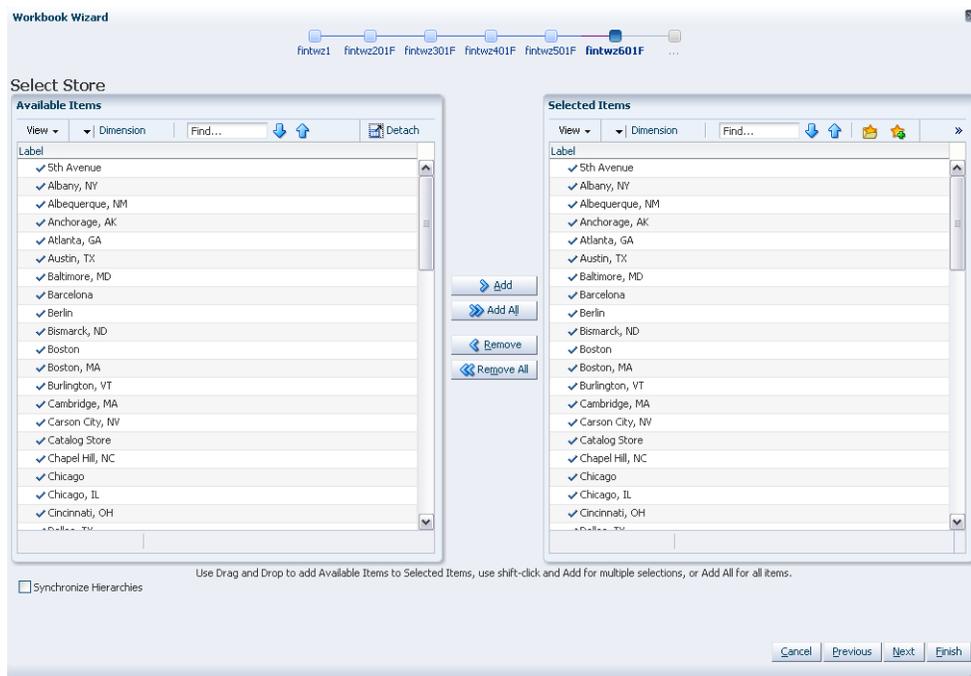
6. Select the end date for the forecast horizon and click **Next**.

Choose one of the following options:

- Specify future time periods beyond the forecast horizon to include in your workbook.
- Select **Do not include dates** after the horizon if no future dates are to be included in the workbook.

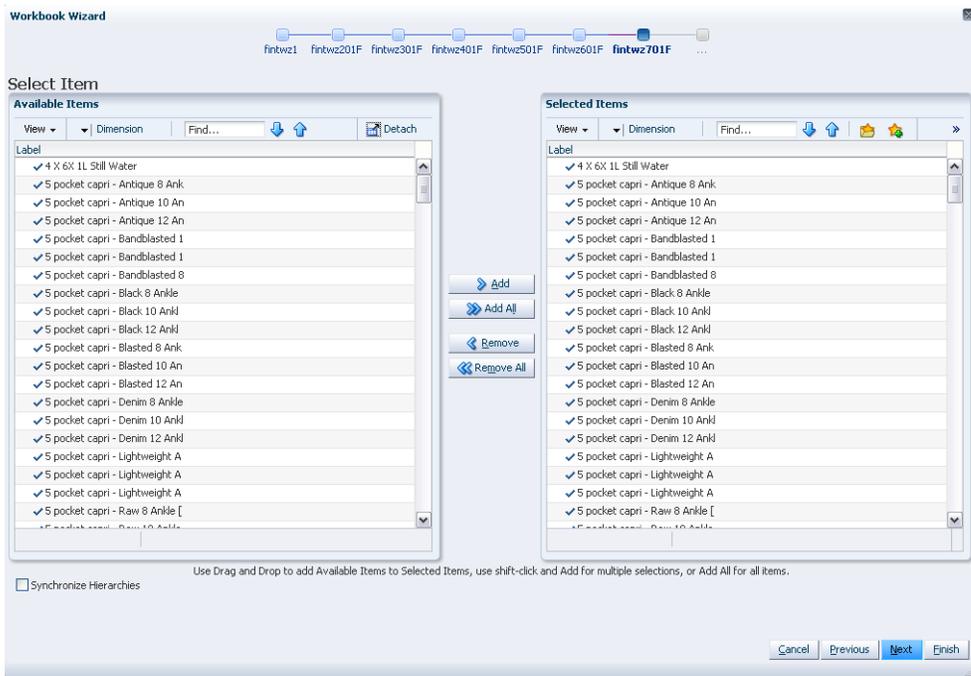
Figure 9–6 Workbook Wizard: Forecast Horizon End Date

7. Select the specific locations that you want to view and click Next.

Figure 9–7 Workbook Wizard: Select Stores

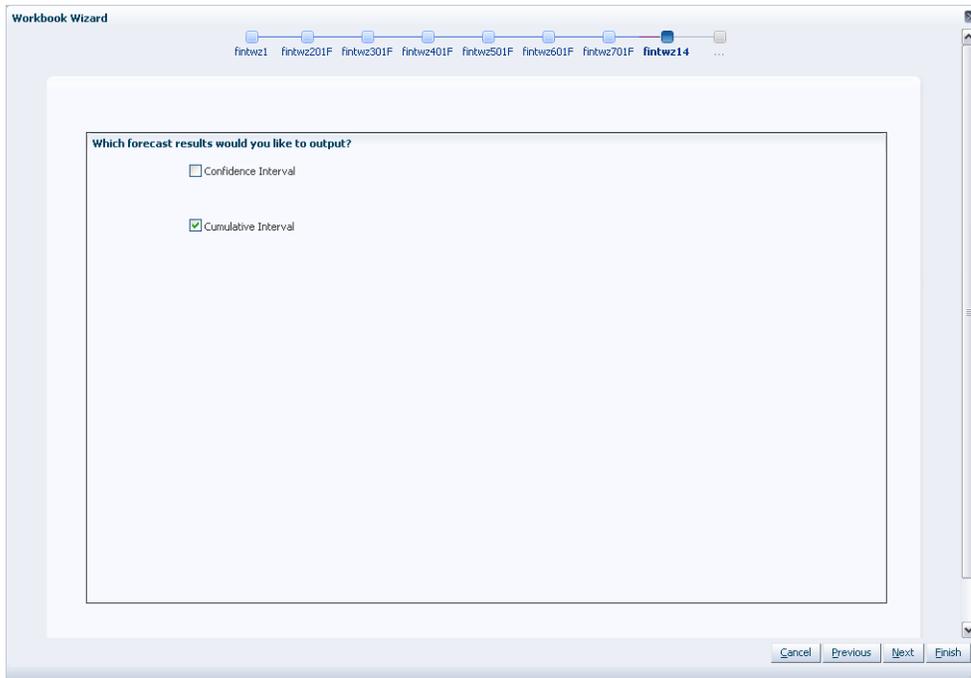
8. Select the products and click Next.

Figure 9–8 Workbook Wizard: Select Items



9. Select the forecast results to view and click **Finish**.

Figure 9–9 Workbook Wizard: Forecast Results



The Interactive Forecasting workbook is built.

Once the wizard is completed, the forecast is generated based on your selections in the wizard and the Default Forecast Method set for the specified forecast level in the [Forecast Administration Workbook](#).

Forecasting Parameter View

The Forecasting Parameter view is based on the intersection of the Product and Location dimensions for the select forecast level.

The Forecasting Parameter view contains the following measures. Of these measures, all but System picked model are editable. Changes can be made to the editable parameters, that is, History Start Date, Forecast Method, Forecast Start Date, and Forecast End Date, to regenerate a forecast of the time series.

Figure 9–10 Forecast Parameter View

Product	Location	Measure
5 pocket capri - Antique 8 Ank	London-Oxford Street	

Forecast Method 1 - itm/str/week-Final	Simple
Forecast Picked Method 1 - itm/str/week-Final	No Generated Forecast
History Start Date 1 - itm/str/week-Final	1/10/1998
Forecast Start Date 1 - itm/str/week-Final	3/14/1998
Forecast End Date 1 - itm/str/week-Final	5/15/1998

Measures: Forecasting Parameter View

The Forecasting Parameter view contains the following measures:

Forecast Method

A list from which you can select the method used to generate the forecast. The workbook defaults to the method selected in the [Forecast Administration Workbook](#) for the specified level if the method is supported by Interactive Forecasting. Otherwise, it defaults to *No Forecast*.

Forecast Picked Method

This is the method that was used to generate the forecast for the given product/location combination. This field is useful when combined methods are requested (for example, ES methods). Then this field displays the actual method the system picked from the combined methods. In case stand-alone methods are chosen, generally this field is the same as the method chosen in Forecast Administration or Forecast Maintenance. However, if the requested method is unable to produce a good fit, the system defaults to a simpler method, and that method is displayed here.

History Start Date

This is the starting date for historical sales data. For example, if your system start date is January 1, 2003, but you only want to use historical sales data from the beginning of 2011, you need to set your History Start Date to January 1, 2011. Only history after this date is used for generating the forecast. The default is the system start date unless otherwise specified. If sales data is collected weekly or monthly, RDF generates forecasts only using data from sales periods after the one containing the history start date.

It is also important to note that the system ignores leading zeros that begin at the history start date. For example, if your history start date is January 1, 2011, and an item/location does not have sales history until February 1, 2011, the system considers

the starting point in that item/location's history to be the first data point where there is a non-zero sales value.

Forecast Start Date

This is first date of the forecast horizon. The default is the start date selected in the wizard.

Forecast End Date

This is the last date of the forecast horizon. The default is the end date selected in the wizard.

Additional Information About the Interactive Forecasting View Parameters

Interactive Forecasting assumes that the historical data are contained in the measure *posxlb*. All forecasts generated by interactive forecasting are based on these data.

Performance may be improved by utilizing a mask measure. Interactive forecasting uses the measures *mask<forecast level>xb*; for example *mask01xb*. Here, *<forecast level>* corresponds to the forecast level chosen in the wizard when opening the Interactive Forecasting workbook. The mask is a boolean measure. Only cells with the value **true** are included in the forecast. This is useful when desiring to focus on only a few product/locations.

The Interactive Forecast method is intended to be a quick way to view results produced by different forecast methods. It allows the user to gauge, in a qualitative way, the suitability of a particular forecast method for a given set of sales history data. Even though it uses the exact same forecast methods as the full forecast, only the parameters listed previously will affect the results. No other user-input values are used. Users wishing to vary a larger number of parameters are encouraged to use the regular forecast functionality.

Also, the Winters Method has a pre-established minimum history of 104 periods.

Interactive Forecasting View

The Interactive Forecasting view is based on the intersection of the Product, Location, and Calendar dimensions for the forecast level selected in the wizard process. The Interactive Forecasting view contains the History Data and the System Forecast. Of these measures, only History Data is editable.

Figure 9–11 Interactive Forecasting View

	1/16/1998	1/23/1998	1/30/1998	2/6/1998	2/13/1998	2/20/1998
History Data - Weekly Sales	0.00	0.00	0.00	0.00	0.00	0.00
Forecast Cumulative Interval 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00
System Forecast 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00

Measures: Interactive Forecasting View

The Interactive Forecasting view contains the following measures:

History Data

This is the historical sales data set in the [Forecast Administration Workbook](#). This field is editable so you can change out-of-character sales if needed. For example, if your battery sales went up during a major power outage, you can lower the sales data back to a more normal level so the unusually high sales adversely affects your forecast. However these changes are for simulation purposes only and cannot be committed.

System Forecast

The quantity that the system predicts is required for the product, location, and calendar combination displayed. Changes to History Data, History Start Date, Forecast Start Date, Forecast End Date, and Forecast Method causes the System Forecast to re-calculate when *Calculate* is selected from the toolbar.

Forecast Like-item, Sister-Store Task

The Forecast Like-Item, Sister-Store task provides the ability to model a new product's demand after an existing product. Forecasts can thus be generated for the new product based on the selected history or the forecast of the existing product plus an Adjustment Ratio. Likewise, the sales history or the forecast of existing store locations can be used as the forecast foundation for new locations. This task includes the following views:

- [Like-Item View](#)
- [Sister-Store View](#)
- [Advanced Parameter View](#)

Note: This task is only valid at final-levels. Therefore, this task may include hierarchy dimensions that are at higher positions than item or store (for example, subclass or region).

Forecast Like-Item, Sister-Store Workbook

To build the Forecast Like-Item, Sister-Store workbook, perform these steps:

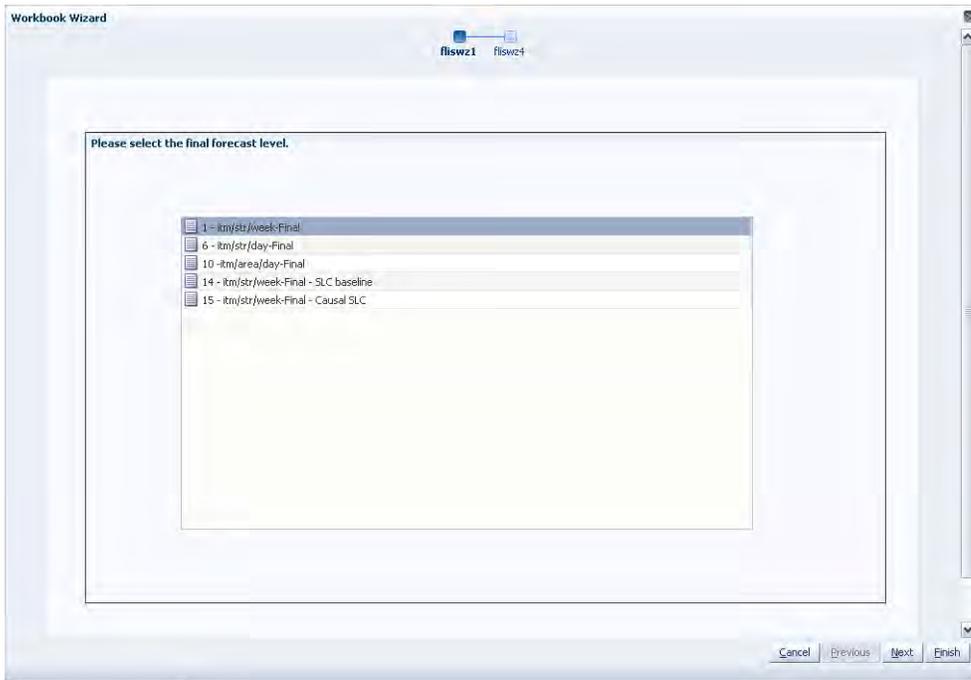
1. Click the New Workbook icon in the Forecast Like-Item, Sister-Store task.

Figure 10–1 Forecast Like-Item, Sister-Store Task



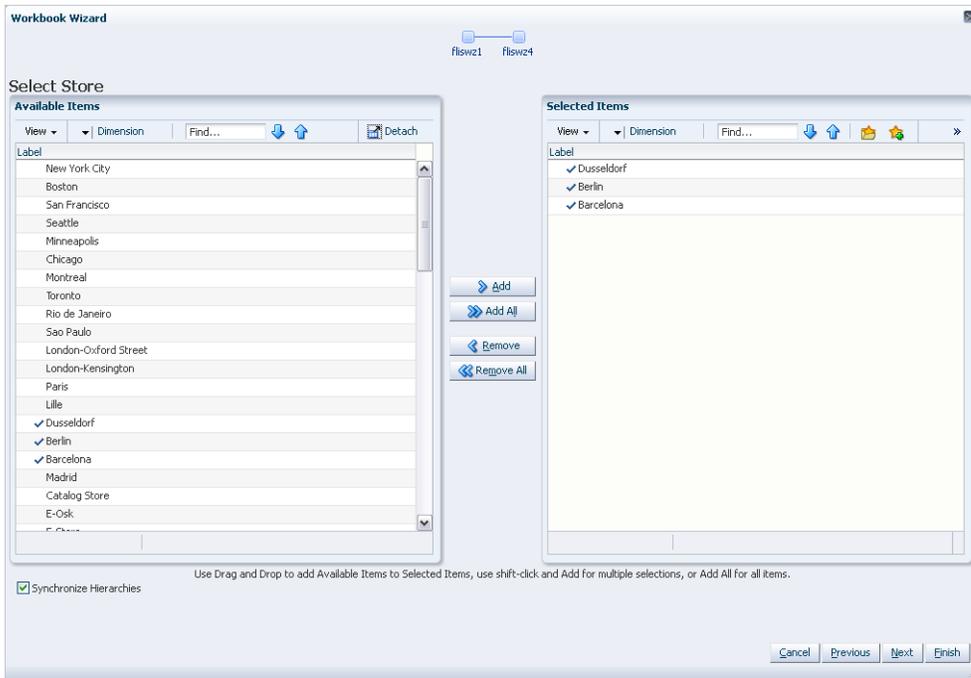
2. The Workbook wizard opens. Select the forecast level that you want to evaluate and click **Next**.

Figure 10–2 Workbook Wizard: Forecast Level



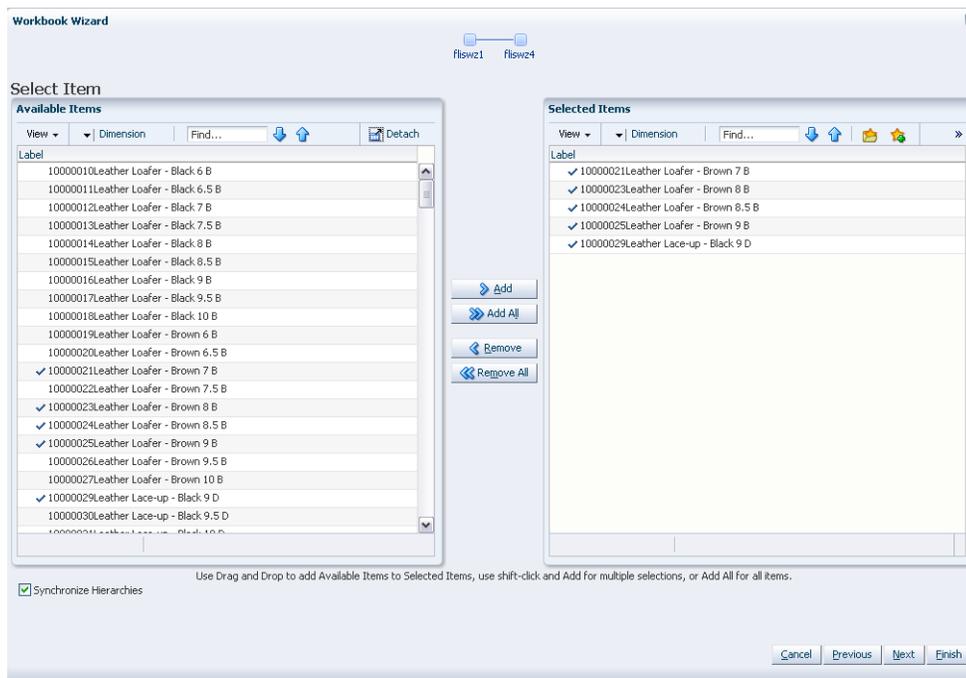
3. Select the specific locations that you want to view and click **Next**.

Figure 10–3 Workbook Wizard: Select Stores



4. Select the items you want to work with and click **Finish**.

Figure 10-4 Workbook Wizard: Select Item

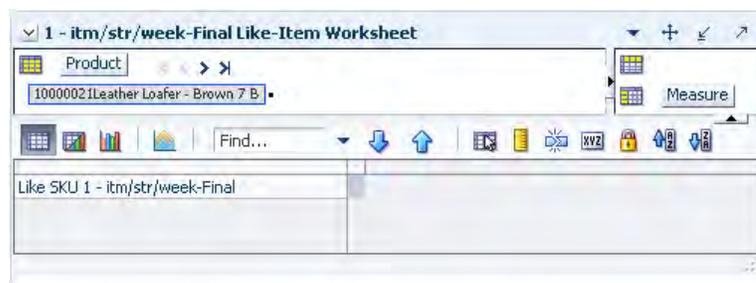


The Forecast Like-Item, Sister-Store workbook is built.

Like-Item View

The Like-Item view is used to forecast a new item by modeling it after an existing item.

Figure 10-5 Final Like-Item View



Measures: Like-Item View

The Like-Item view contains the following measure:

Like SKU

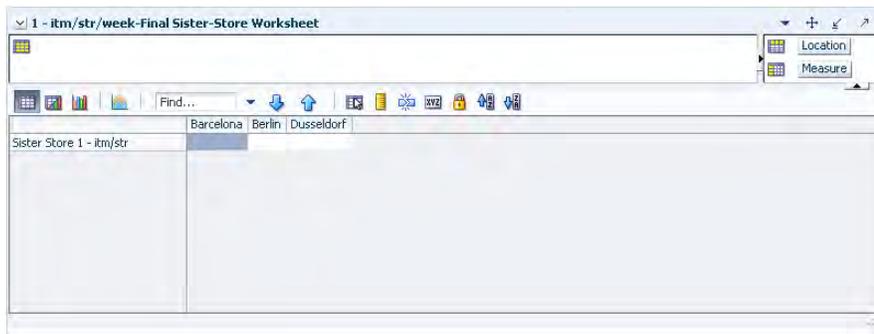
The Like SKU field displays the items selected during the wizard process. In the example shown in the [Final Like-Item View](#), the like-item (existing item) is selected on the right from a pick-list across from the new item. The like-item's forecast or sales is used for the new item based on the parameters selecting in the Advance Parameter view. In the previous example, Leather Loafer - Black 8 B is a new item, and Leather Lace-up Boat Shoes 8 B is being selected as the existing item that is used to Leather Loafer - Black 8 B model's forecast.

Note: When working with both the Like SKU and Sister-Store views, making a selection in one of the views requires Calculate to be run before the pick-list options are available in the other view.

Sister-Store View

The Sister-Store view is used to forecast demand for a new store by modeling it after an existing store.

Figure 10–6 *Final Sister-Store View*



Measures: Sister-Store View

The Sister-Store view contains the following measures:

Sister Store

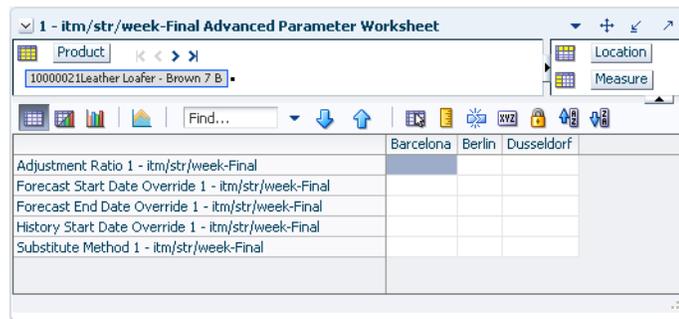
Displays the locations selected during the wizard process. In the example shown in [Figure 10–6](#), the Sister-Store (existing store) is selected on the right from a pick-list across from the new store. The sister store's forecast or sales is used for the new store based on the parameters selecting in the Advance Parameter view.

Note: When working with both the Like SKU and Sister-Store views, if a selection is made in one of the views, Calculate must be run before the pick-list options are available in the other view.

Advanced Parameter View

The Advanced Parameter view is used to manage item/locations in which the Forecast Start Date, Forecast End Date, or History Start Date varies from the default settings in the [Forecast Administration Workbook](#) and to set the Adjustment Ratio for the new item or location being forecast.

Figure 10–7 Advanced Parameter View



Measures: Advanced Parameter View

The Advanced Parameter view contains the following measures:

Adjustment Ratio

You may enter an Adjustment Ratio to apply to the forecast for the new product/location combination displayed. This is a real number between [zero (0), infinity]. The default (NA) value is 1.00, (in other words 100%) which translates to no adjustment.

Note: If demand for a new item is expected to be 30% greater than its like-item, the Adjustment Ratio would be set to 1.30. If demand for a new item is expected to be 30% less than its like-item, the Adjustment percent would set to 0.70.

Forecast Start Date Override

This parameter represents the date to start forecasting for an item/location combination. This parameter can be set in the future if using like-item or Sister-Store functionality, and, upon reaching that time, the forecast is generated. If this date is set to the past, it is ignored in favor of the Forecast Start Date from the [Forecast Administration Workbook](#). This means that the Forecast Start Date for this intersection needs to be edited once it is no longer in the future. For like-item or sister store, the Forecast Start Date and the History Start Date should be set to the same date. It is important to understand how Forecast Start Date should be used in conjunction with Forecast End Date (Refer to the [Forecast End Date Override](#) section). No value is in this measure if the system default set in the [Forecast Administration Workbook](#) is to be used.

Note: This measure can also be set in the [Forecast Maintenance Workbook](#). Changes to this measure can be seen in the Forecast Maintenance task and the [Forecast Like-Item, Sister-Store Workbook](#). The most recent commit (between either of the tasks) is the value used by the system.

Forecast End Date Override

This parameter represents the last point in time for which the Forecasting Engine forecasts for an item/location combination. Should this parameter be set to a date less than the Forecast Start Date plus the Forecast Length (in the [Forecast Administration Workbook](#)), the engine forecasts zero (0) past this date. If Forecast End Date is more

than Forecasting Start Date plus Forecasting Length, no forecast is generated outside Forecasting Start Date plus Forecasting Length window. In other words, both Forecast Start Date and Forecasting End Date are relevant for time periods within the forecast horizon set at the global level. No value is in this measure if the system default set in the [Forecast Administration Workbook](#) is to be used.

Forecast End Date can be used for new item or location forecasting if the item or location needs to be forecast for a period shorter than the Like TS Duration (set globally in [Forecast Administration Workbook](#)).

Note: This measure can also be set in the [Forecast Maintenance Workbook](#). Changes to this measure can be seen in the [Forecast Maintenance Workbook](#) and the [Forecast Like-Item, Sister-Store Workbook](#). The most recent commit (between either of the tasks) is the value used by the system.

History Start Date Override

This parameter represents the first point in time from which the Forecasting Engine begins training and modeling (that is, if there are two years of history, but only one year is required, set the start date to a year ago). This parameter overrides at the item/store level from the global settings in the [Forecast Administration Workbook](#). This can be used to level out past sales. For example, if there is a large spike in the first three weeks of sales for an item was on sale, set the Historical Start Date to one week past that period. Those first few weeks are not used when generating the forecast.

It is also important to note that the system ignores leading zeros that begin at the history start date. For example, if your history start date is January 1, 2011 and an item/location does not have sales history until February 1, 2011, the system considers the starting point in that item/location's history to be the first data point where there is a non-zero sales value.

The History Start Date for the new item or new store should be set with the same date as the Forecast Start Date.

Note: When using any of the Lifecycle Methods (refer to [Substitute Methods](#)), the History Start Date for the substitute item or location must be set to the point in the sales history that the new item or location begins using as its sales.

Note: This measure can also be set in the [Forecast Maintenance Task workbook](#). Changes to this measure can be seen in the [Forecast Maintenance Task workbook](#) and the [Forecast Like-Item, Sister-Store Workbook](#). The most recent commit (between either of the tasks) is the value used by the system.

Substitute Methods

Displays a list from which you can select the substitute method. When a Substitute Method is used to forecast, the method set for an intersection is cleared once the Default Forecast Start Date is greater than the Forecast Start Date Override plus the Like TS Duration for the intersection. Valid options are:

Table 10–1 Substitute Methods for Forecasting

Option	Description
None	There is no substitution for this product/location combination. This is the default value.
Seasonality/New Item	You provide a like-item that has a similar seasonality pattern that sells at the same store. The new product's forecast is the like-item's demand forecast with the applied adjustment. The forecast is set to zero (0) for all dates before the new product's start date.
Seasonality/New Store	You provide a Sister-Store that has a similar seasonality pattern that sells the same product. The product's forecast at the new store is the demand forecast of the same product at the Sister-Store with the applied adjustment. The forecast is set to zero (0) for all dates before the new store's open date.
Seasonality/New Item /New Store	You provide a like-item that sells at a Sister-Store that has a similar seasonality pattern. The new product's forecast at the new store is the demand forecast of the like-item at the Sister-Store with the applied adjustment. The forecast is set to zero (0) for all dates before the new product's start date after the new store opens.
Lifecycle/New Item	You provide a like-item that had a similar lifecycle pattern that sells at the same store. The new product's forecast is the like-items actual sales with the applied adjustment shifted such that the like-item's first sales matches the new product's start date.
Lifecycle/New Store	You provide a Sister-Store that had a similar lifecycle pattern that sells the same product. The product's forecast at the new store is the products actual sales at the Sister-Store with the applied adjustment shifted such that the sister store's first sales matches the new store's open date.
Lifecycle/New Item/New Store	You provide a like-item that sells at a Sister-Store that has a similar lifecycle pattern. The new product's forecast at the new store is like-item's actual sales at a Sister-Store with the applied adjustment shifted such that the like-item's first sales at the Sister-Store matches the new product's start date after the new store opens.

Required Steps for Forecasting Using Each of the Like SKU/Sister-Store Methods

To support Like SKU /Sister-Store functionality, the Like TS Duration must be set in the [Forecast Administration Workbook](#) - Advance tab. This parameter sets the number of weeks of history required after which RDF stops using the substitution method and starts using the system forecast generated by the forecast engine.

Procedures

The following procedures outline the steps required for using each of the methods described in [Table 10–1, "Substitute Methods for Forecasting"](#):

Seasonality/SKU

Introduces a new item at an existing store (Like-item with a similar forecast):

1. From the Like-Item view, select a like-item from the list across from the new item.

2. From the Advanced Parameter view, Set the Forecast Start Date for the new item at an existing store.
3. From the Advanced Parameter view, Set the History Start Date for the new item at the existing store to the same date as the Forecast Start Date.
4. From the Advance Parameter view, Set the Adjustment percent (optional) for the new item at the existing store.

Seasonality/STR

Introduces an existing item at a new store (Sister-Store with similar forecast):

1. From the Sister-Store view, Select a Sister-Store from the list across from the new store.
2. From the Advanced Parameter view, Set the Forecast Start Date for the existing item at the new store.
3. From the Advanced Parameter view, Set the History Start Date for the existing item at the new store to the same date as the Forecast Start Date.
4. From the Advance Parameter view, Set the Adjustment percent (optional) for the existing item at the new store.

Seasonality/SKU_STR

Introduces a new item at a new store (Like-item and Sister-Store with a similar forecast):

1. From the Like-Item view, Select a like-item from the list across from the new item.
2. From the Sister-Store view, Select a Sister-Store from the list across from the new store.
3. From the Advanced Parameter view, Set the Forecast Start Date at the intersection of the new item and the new store.
4. From the Advanced Parameter view, Set the History Start at the intersection of the new item and new store.
5. From the Advance Parameter view, Set the Adjustment percent (optional) at the intersection of the new item and new store.

Lifecycle/SKU

Introduces a new item at an existing store (Like-item's sales history to be used as the forecast for the new item):

1. From the Like-Item view, Select a like-item from the drop-down list across from the new item.
2. From the Advanced Parameter view, Set the Forecast Start Date for the new item at the existing store.
3. From the Advanced Parameter view, Set the History Start Date for the new item at the existing store to the same date as the Forecast Start Date.
4. From the Advanced Parameter view, Set the History Start Date for the like-item at the existing store to the point in its sales history that maps to the new item's forecast.
5. From the Advance Parameter view, Set the Adjustment percent (optional) for the new item at the existing store.

Lifecycle/STR

Introduces an existing item at new store (Sister store's sales history to be used as the forecast for the new store):

1. From the Sister-Store view, Select a Sister-Store from the list across from the new store.
2. From the Advanced Parameter view, Set the Forecast Start Date for the existing item at the new store.
3. From the Advanced Parameter view, Set the History Start Date for the existing item at the new store to the same date as the Forecast Start Date.
4. From the Advanced Parameter view, Set the History Start Date at the intersection of the Sister-Store and existing item to the date in its sales history that maps to the new store's forecast.
5. From the Advance Parameter view, Set the Adjustment percent (optional) for the existing item at the new store.

Lifecycle/SKU_STR

Introduces a new item at a new store (Like-item's and sister store's sales history to be used as the forecast for a new item at a new store):

1. From the Like-Item view, Select a like-item from the list across from the new item.
2. From the Sister-Store view, Select a Sister-Store from the list across from the new store.
3. From the Advanced Parameter view, Set the Forecast Start Date at the intersection of the new item and new store.
4. From the Advanced Parameter view, Set the History Start Date at the intersection of the new item and new store to the same date as the Forecast Start Date.
5. From the Advanced Parameter view, Set the History Start Date at the intersection of the Like-item and Sister-Store to the date in its sales history that maps to the new item and new store's forecast.
6. From the Advance Parameter view, Set the Adjustment percent (optional) at the intersection of the new item and new store.

Forecast Maintenance Task

The Forecast Maintenance task allows you to select and modify forecasting parameters for product/location combinations when the values of these parameters differ from the default values that are assigned in the [Forecast Administration Workbook](#).

Suppose, for example, that the default forecast method of all the products in the database was set in the [Forecast Administration Workbook](#) to be AutoES. For a particular product, however, you know that a SeasonalES model is a better predictor of sales performance. To make this change, you must access the Forecast Maintenance task, select the product/location intersection to be reviewed, and make the appropriate change to the forecast method.

The Forecast Maintenance task is split into two workflow activities:

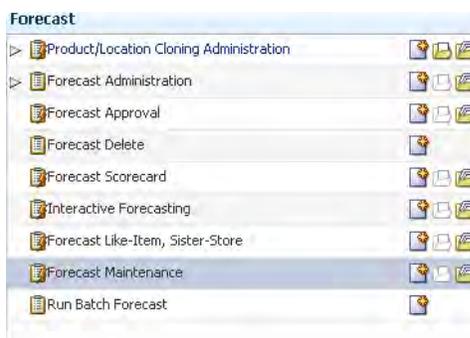
- Basic Settings — Includes all Final Level views and their respective Source Level views.
- Advanced Settings — Includes the Forecast Start Date, End Date, and History Start Date measures for the selected Final Level.

Forecast Maintenance Workbook

To build the Forecast Maintenance workbook, perform these steps:

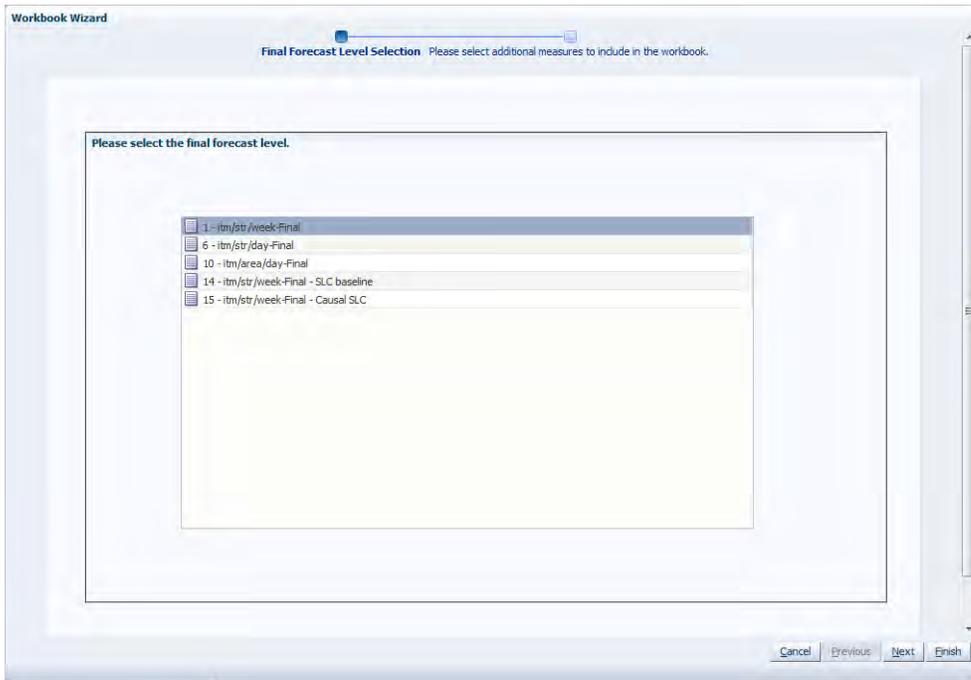
1. Click the New Workbook icon in the Forecast Maintenance task.

Figure 11–1 Forecast Maintenance Task



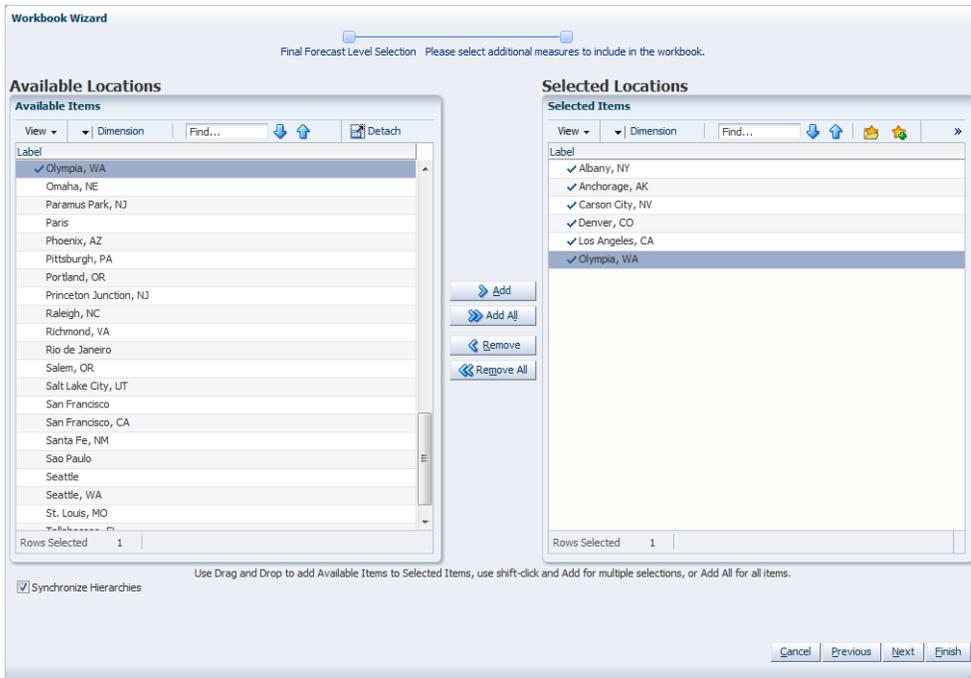
2. The Workbook wizard opens. Select the forecast level that you want to evaluate and click **Next**.

Figure 11–2 Workbook Wizard: Causal Level

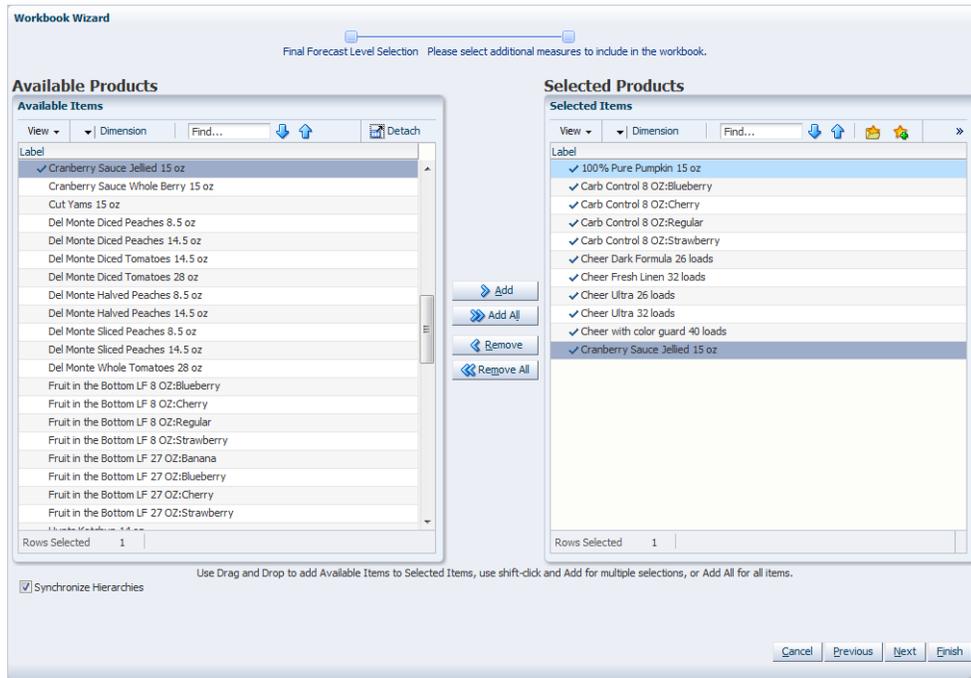


3. Select the specific locations that you want to view and click **Next**.

Figure 11–3 Workbook Wizard: Select Locations



4. Select the items you want to work with and click **Finish**.

Figure 11–4 Workbook Wizard: Select Item

The Forecast Maintenance workbook is built.

Forecast Maintenance Workbook: Basic Settings

The Forecast Maintenance Basic Settings workbook includes [Final and Source Level Views - Basic Settings](#).

Final and Source Level Views - Basic Settings

The Final Level View and Source Level view allows for certain parameters set at a global level (in Forecast Administration) to vary at different item/locations. [Figure 11–5](#) and [Figure 11–6](#) show the Final Level and Source Level views.

Figure 11–5 Final Level View - Basic Settings

	10000019Leather	10000024Leather	10000026Leather	10000029Leather
Approval Method Override 1 - itm/str/week-Final				
Forecast Method Override 1 - itm/str/week-Final	No Override	No Override	No Override	No Override
Keep Last Changes Override 1 - itm/str/week-Final				
Source Level Override 1 - itm/str/week-Final				
Optimal Source Levels 1 - itm/str/week-Final				
Pick Optimal Level 1 - itm/str/week-Final	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Winters Mode Override 1 - itm/str/week-Final	No Override	No Override	No Override	No Override
Moving Average Window Length Override 1 - itm/str/week-Final	0	0	0	0
Bayesian Alpha Override 1 - itm/str/week-Final	1.00	1.00	1.00	1.00

Figure 11–6 Source Level View - Basic Settings

	122Loafer	223Lace Up
Forecast Method Override 4 - sbc/str/week	No Override	No Override
Winters Mode Override 4 - sbc/str/week	No Override	No Override
Moving Average Window Length Override 4 -	0	0
Bayesian Alpha Override 4 - sbc/str/week	1.00	1.00

Measures: Final and Source Level Views - Basic Settings

The Final and Source Level Views - Basic Settings contains the following measures:

Approval Method Override

Set only at the final-level, the Approval Method Override is a list from which you select the approval policy for individual product/location combinations. No value is in this field if the system default set in the [Forecast Administration Workbook](#) is to be used. Valid values are:

Field	Description
Manual	The System Forecast and Adjusted Forecast are not automatically approved. Forecast values must be manually approved by accessing and amending the Forecast Approval Workbook .
Automatic	The Adjusted Forecast is automatically approved as is.
By Alert <name of the alert>	This list of values may also include any Forecast Approval alerts that have been configured for use in the forecast approval process. Alerts are configured during the implementation. Refer to the <i>Oracle Retail Predictive Application Server Configuration Tools User Guide</i> for more information on the Alert Manager and the <i>Oracle Retail Demand Forecasting Configuration Guide</i> for more information on configuring Forecast Approval Alerts.

Note: If you select a specific alert as your approval method and later on you delete the alert, the approval works as manual. The same happens if the alert is on a wrong intersection.

Bayesian Alpha Override

This is the override of the Bayesian Alpha parameter that is specified globally in the [Forecast Administration Workbook](#). The override can be specified at every forecasting level, that is, final and source-levels.

An alpha value closer to one (or infinity) weights the sales plan more in creating the forecast, whereas alpha closer to zero weights the known history more. The default is one (1).

Forecast Method Override

Set at both final- and source-levels, the Forecast Method Override is a list from which you can select a different forecast method than the Default Forecast Method set in the [Forecast Administration Workbook](#). *No Override* appears in this field if the system default set in the [Forecast Administration Workbook](#) is to be used. Valid options depend on your system setup.

Keep Last Changes Override

Set only at final-levels, Keep Last Changes Override field may be used to override the default setting at a product/location intersection. *None* appears in this field if there is no override applied to the intersection.

Moving Average Window Length Override

Used with Moving Average forecast method, this is the number of data points in history used in the calculation of Moving Average. This parameter, when set to a non-zero value, overrides the value in the [Forecast Administration Workbook](#).

Optimal Source Levels

Displayed only at final-levels, a value is populated in this field if AutoSource has been run on the final-level. The AutoSource executable evaluates all levels associated to a final-level and returns the source-level that yields the optimal forecast results. For more information on AutoSource, refer to the *Oracle Retail Demand Forecasting Implementation Guide*.

Note: For item/stores that are new or highly seasonal, AutoSource may not return the best recommendation since new items may not have an adequate sales history length and highly seasonal items may only sell for a short period during the year. For these items, you should not set the AutoSource recommendation as default at the final-level. Only use AutoSource recommendations for item/stores that have an adequate sales history.

Pick Optimal Level

Set only at final-levels, a check in this field (set to **true**) indicates that the batch forecast should use the Optimal Source Level selected by AutoSource. For more information on AutoSource, refer to the *Oracle Retail Demand Forecasting Implementation Guide*.

Source Level Override

Set only at final-levels, the Source Level Override is the level at which the aggregate, more robust forecast is run. Forecast data from this level is spread down to the lowest level based on the relationship between the two levels in the hierarchy. *No Override* appears in this field if the system default set in the [Forecast Administration Workbook](#) is to be used.

Winters Mode Override

Used only by Forecasting Methods calling Additive or Multiplicative Winters. This parameter allows you to override the Winters Mode setting assigned in the [Forecast Administration Workbook](#). The default setting for this parameter is No Override.

Forecast Maintenance Workbook: Advanced Settings

The Advanced Settings workflow activity is used to override the dates that are used in the forecast generation process as well as historical start dates for any intersection at the final-level that varies from the default settings in the [Forecast Administration Workbook](#).

Advanced Final Parameter View

[Figure 11-7](#) shows the Advanced Final Parameter view.

Figure 11–7 Advanced Final Parameter View

Parameter	Value
Forecast Start Date Override 1 - itm/str/week-Final	100% Pure
Forecast End Date Override 1 - itm/str/week-Final	
History Start Date Override 1 - itm/str/week-Final	
Apply Promotional Lift Override 1 - itm/str/week-Final	<input type="checkbox"/>
Apply Regular Price Lift Override 1 - itm/str/week-Final	No Override
Apply Regular Price Lift over Promotion Override 1 - itm/str/week-Final	No Override
Regular Price Decaying Factor Override 1 - itm/str/week-Final	
Apply Regular Price Self Lift Override 1 - itm/str/week-Final	No Override
Apply Regular Price HALO Lift Override 1 - itm/str/week-Final	No Override
Apply Regular Price Cannibalization Lift Override 1 - itm/str/week-Final	No Override
Regular Price Effective Periods Override 1 - itm/str/week-Final	
Apply Promotion Halo Lift Override 1 - itm/str/week-Final	No Override
Apply Promotion Cannibalization Lift Override 1 - itm/str/week-Final	No Override

Measures: Advanced Final Parameter View

The Advanced Final Parameter view contains the following measures:

Apply Promotion Cannibalization Lift Override

This parameter is only effective if the Apply Promotional Lift override is selected (set to **true**).

For this parameter, select a value from the list to instruct RDF to apply or override the system generated Cannibalization effects.

- **Yes** — applies the system generated Cannibalization effects
- **No** — does not apply any Cannibalization lifts
- **No Override** — uses the default setting from Forecast Administration

This is an override of the global setting.

Apply Promotion Halo Lift Override

This parameter is only effective if the Apply Promotional Lift override is selected (set to **true**).

For this parameter, select a value from the list to instruct RDF to apply or override the system generated Halo effects.

- **Yes** — applies the system generated Halo effects
- **No** — does not apply any Halo lifts
- **No Override** — uses the default setting from Forecast Administration

This is an override of the global setting.

Apply Promotional Lift Override

Select this parameter (set to **true**) if you want RDF to ignore the system calculated lifts, and apply the overrides.

Apply Regular Price Cannibalization Lift Override

This parameter is only effective if:

- **Apply Regular Price Lift Override** is set to Apply System or Apply Override

or

- **Apply Regular Price Lift Override** is set to No Override and [Apply Regular Price Lift](#) is set to Apply System or Apply Override

For this parameter, select a value from the list to instruct RDF to override the system generated Cannibalization effects.

- **Yes** — overrides the system generated Cannibalization effects
- **No** — does not apply any Cannibalization lifts
- **No Override** — applies the system generated Cannibalization effects

This is an override of the global setting.

Apply Regular Price Halo Lift Override

This parameter is only effective if:

- **Apply Regular Price Lift Override** is set to Apply System or Apply Override

or

- **Apply Regular Price Lift Override** is set to No Override and [Apply Regular Price Lift](#) is set to Apply System or Apply Override

For this parameter, select a value from the list to instruct RDF to override the system generated Halo effects.

- **Yes** — overrides the system generated Halo effects
- **No** — does not apply any Halo lifts
- **No Override** — applies system generated Halo effects

This is an override of the global setting.

Apply Regular Price Lift Over Promotion Override

This parameter is only effective if:

- **Apply Regular Price Lift Override** is set to Apply System or Apply Override

or

- **Apply Regular Price Lift Override** is set to No Override and [Apply Regular Price Lift](#) is set to Apply System or Apply Override

For this parameter, select a value from the list to instruct RDF to override the system generated Cannibalization effects.

- **Yes** — overrides the system generated Cannibalization effects
- **No** — does not apply any Cannibalization lifts
- **No Override** — applies the system generated Cannibalization effects

This is an override of the global setting.

Apply Regular Price Lift Override

For this parameter, select a value from the list to instruct RDF to apply or override regular price lifts:

- **Do Not Apply** — does not apply regular price effects (including Self, Cannibalization, and Halo effects)
- **Apply System** — applies system calculated effects
- **Apply Override** — ignores system calculated effects and applies overrides
- **No Override** — applies the default setting from Forecast Administration

This is an override of the global setting.

Apply Regular Price Self Lift Override

This parameter is only effective if:

- **Apply Regular Price Lift Override** is set to Apply System or Apply Override
- or*
- **Apply Regular Price Lift Override** is set to No Override and [Apply Regular Price Lift](#) is set to Apply System or Apply Override.

For this parameter, select a value from the list to instruct RDF to apply or override the system generated Cannibalization:

- **Yes** — overrides the system generated Cannibalization effects
- **No** — does not apply any Cannibalization lifts
- **No Override** — applies the system generated Cannibalization effects

This is an override of the global setting.

Forecast End Date Override

This parameter represents the last point in time for which the forecasting engine forecasts for a particular intersection. Should this parameter be set to a date less than the Forecast Start Date plus the Forecast Length (in Forecast Administration), the engine forecasts zero (0) past this date. If Forecast End Date is more than Forecasting Start Date plus Forecasting Length, you do not get a forecast outside Forecasting Start Date plus Forecasting Length. In other words, both Forecast Start Date and Forecasting End Date are relevant for time periods within the forecast horizon set at the global level. No value is in this measure if the system default set in the [Forecast Administration Workbook](#) is to be used.

Note: This measure can also be set in the [Forecast Like-Item, Sister-Store Workbook](#). Changes to this measure can be seen in the [Forecast Maintenance Workbook](#) and the [Forecast Like-Item, Sister-Store Workbook](#). The most recent commit (in either task) is the value used by the system.

Forecast Start Date Override

This parameter represents the date to start forecasting at a particular intersection. If this date is set to the past, it is ignored in favor of the Forecast Start Date from the [Forecast Administration Workbook](#). This means that you do not need to change the Forecast Start Date once it is no longer in the future. It is important to understand how Forecast Start Date should be used in conjunction with [Forecast End Date Override](#).

No value is in this measure if the system default set in the [Forecast Administration Workbook](#) is to be used.

Note: This measure can also be set in the [Forecast Like-Item, Sister-Store Workbook](#). Changes to this measure can be seen in the [Forecast Maintenance Workbook](#) and the [Forecast Like-Item, Sister-Store Workbook](#). The most recent commit (between either task) is the value used by the system.

History Start Date Override

This parameter represents the first point in time from which the Forecasting Engine begins training and modeling (that is, if there are two years of history, but you only want to use one year, you set the start date to a year ago). This parameter overrides the History Start Date set in the [Forecast Administration Workbook](#) to the desired item/location intersection. For example, if you have a large spike in the first three weeks of sales for an item on sale, you can set the Historical Start Date to one week past that period, and those first few weeks are not used when generating the forecast.

It is also important to note that the system ignores leading zeros that begin at the history start date. For example, if your history start date is January 1, 2003, and an item/location does not have sales history until February 1, 2003, the system considers the starting point in that item/location's history to be the first data point where there is a non-zero sales value.

If this parameter is set into the future, there would be no forecast, as the history training window is read as zero.

Note: This measure can also be set in the [Forecast Like-Item, Sister-Store Workbook](#). Changes to this measure can be seen in the [Forecast Maintenance Workbook](#) and the [Forecast Like-Item, Sister-Store Workbook](#). The most recent commit (between either task) is the value used by the system.

Regular Price Decaying Factor Override

This parameter determines how the magnitude of the price effects is applied over time. A regular price change usually triggers a change in demand. However, the effect can be seen only for a few periods following the price change. An adjustable damping mechanism ensures that the regular price effect trends down as time goes by. The decaying factor is a number between zero (0) and one (1). The closer the decaying factor is to zero, the faster the regular price effect decays. This is an override of the global setting.

Regular Price Effective Periods Override

Specifies the number of periods for which the regular price effect is applied. This is an override of the global setting.

Run Batch Forecast Task

The Run Batch Forecast Task wizard allows you to manually execute the forecast generation process at a time other than the regularly scheduled batch job. If a global domain environment is implemented, forecasts generated in the Local domain can be viewed in the Master domain; however, this forecast is isolated to the data in the Local domain. The execution of PreGenerateForecast at the Master domain, then passing the output of the process to Generate from the back-end of each Local domain allows for the Local domains to share a birth date, thus supporting a view to forecast data across Local domains in the Master domain. Refer to the *Oracle Retail Demand Forecasting Implementation Guide* for more information on PreGenerateForecast and Generate.

Running a Batch Forecast Manually

The following sections describe the procedures to manually run a batch forecast.

Before You Begin

Prior to using the Run Batch Forecast Task wizard, at minimum the following tasks must be performed:

1. Create or access a [Forecast Administration Workbook](#).
2. In the Forecast Start Date measure, enter the starting date of the forecast horizon; otherwise the system defaults to the system date (today).
3. Set a Default Forecasting Method for the Final and Source Level.
4. Set a Default Source Level for the Final Level.
5. Set the Forecast Length if there is no value already in this measure.
6. Set the Spreading Profile if the Source Levels are at aggregate dimensions above the Final Level.
7. Commit any changes by selecting from the **File** menu, then select **Commit Now**.
8. Close the task.

Procedure

Follow these steps to manually run a batch forecast:

1. Within the Local or Simple domain, select **New** from the File menu.
2. Select the Forecast tab to display a list of task workbooks for statistical forecasting. Highlight Run Batch Forecast and click **OK**.

3. The Run Batch Forecast wizard opens and prompts you to select the Final Levels to forecast. Select **Next** or **Finish**.

The Run Batch wizard automatically executes PreGenerateForecast and Generate within the Simple or Local Domain. If **Next** is selected from the last wizard window, the wizard does not advance to the completion message until the forecast has been generated. Depending on the amount of product/locations to be forecast and the forecast horizon, it may take a several minutes before the system advances to the final window. When the forecast generation is completed, the wizard displays a window that notifies you of the forecast generation ID.

After a forecast is generated, the Forecast Start Date field is cleared. This ensures that the same forecast is not generated again on the same date. The Next Run Date field is also updated based on the birth date of the forecast plus the Forecast Cycle. Both fields can be viewed in the [Forecast Administration Workbook](#).

Run Batch Forecast Workbook

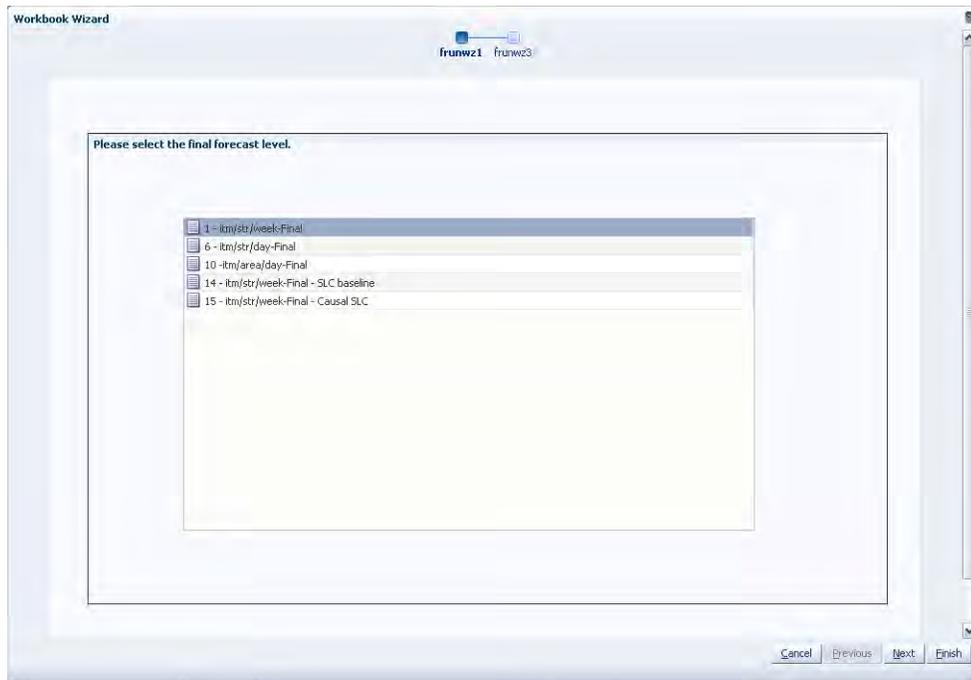
To build the Run Batch Forecast workbook, perform these steps:

1. Click the New Workbook icon in the Run Batch Forecast task.

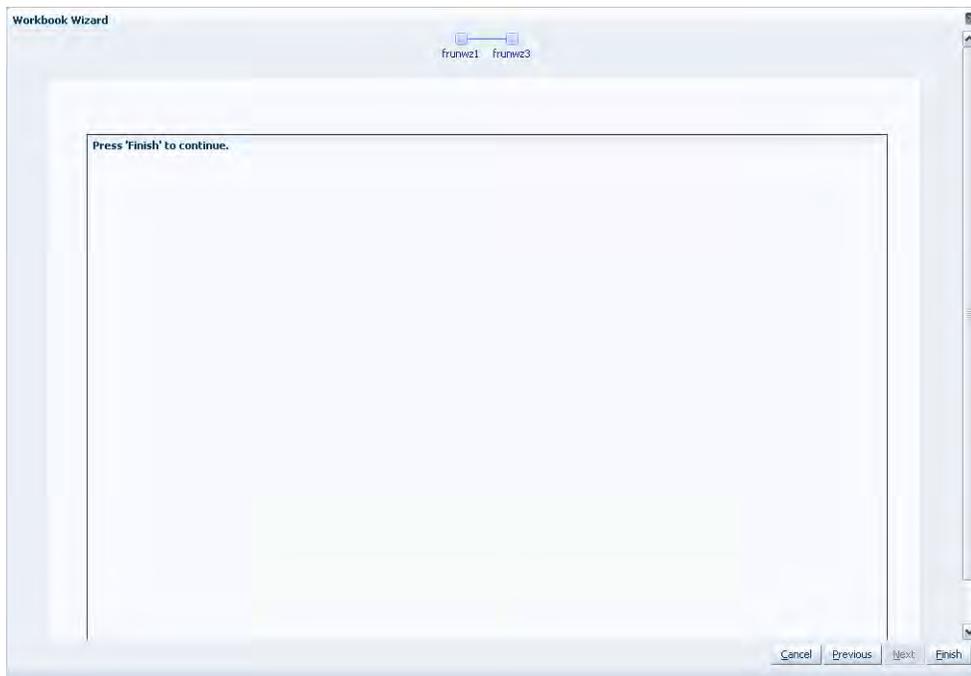
Figure 12–1 Run Batch Forecast Task



2. The Workbook wizard opens. Select the levels to work with and click **Next**.

Figure 12–2 Workbook Wizard: Select Final Forecast Level

3. Verify and click **Finish**.

Figure 12–3 Workbook Wizard: Verify

The Run Batch wizard automatically executes PreGenerateForecast and Generate within the Simple or Local Domain. Depending on the amount of product/locations to be forecast and the forecast horizon, it may take a several minutes before the system advances to the final window. When the forecast generation is completed, the wizard displays a window that notifies you of the forecast generation ID.

After a forecast is generated, the Forecast Start Date field is cleared. This ensures that the same forecast is not generated again on the same date. The Next Run Date field is also updated based on the birth date of the forecast plus the Forecast Cycle. Both fields can be viewed in the [Forecast Administration Workbook](#).

Promote (Promotional Forecasting)

This chapter provides an introduction to promotional forecasting and explains how it differs from the traditional statistical forecasting methodology. It discusses the advantages and limitations of both statistical and promotional forecasting models, and it outlines the use of Oracle Retail's Causal method of forecasting demand. It describes terminology used in the context of promotional forecasting, and it concludes with detailed descriptions of the tasks contained in Promote.

About Promote

Promote is an optional add-on automated predictive solution that allows you to incorporate the effects of promotional and causal events (such as radio advertisements and holiday occurrences) into your time-series forecasts. The promotional forecasting process uses both past sales data and promotional information to forecast future demand.

What Is Promotional Forecasting?

Traditional statistical forecasting methods provide significant benefits to the process of forecasting consumer demand because they are good at predicting level, trend, and seasonality based on sales history. The limitation of traditional statistical methods is that they forecast with less accuracy when there are special events that cause significant deviations in selling patterns.

For example, the Easter holiday, for which companies often run promotions, occurs on a different date each year. Traditional statistical forecast methods can identify seasonality in sales history, but this seasonality is based on periodic similarities in the sales pattern. Since Easter occurs on different dates from year to year (that is, its period of recurrence is not regular), manual intervention is required to predict change in demand using the traditional statistical forecasting method. Events like this are called promotion events. Promotion events, such as advertisements, irregularly occurring holidays, competitor information, free gift with purchase offers, and so on are events that drive businesses from the normal selling cycle.

The goal of a promotional forecasting system is to improve time series forecasting by:

- Providing the forecasting system with visibility as to when certain promotion events were active in the past (for example, identifying which weeks of a given year were affected by an Easter promotion).
- Automatically determining the statistical effect, if any, of these events.
- Incorporating significant effects into the future forecasts for time periods also associated with the observed promotion event.

The Promote module combines the automation of statistical time series forecasting with the improved accuracy found in customized causal forecasting. Promote uses both past sales data and promotional information (for example, advertisements, holidays) to forecast future demand. In order to understand the underlying rationale for the promotional forecasting process, it is important to understand the advantages and limitations of its underlying components.

Comparison Between Promotional and Statistical Forecasting

Statistical time series forecasting uses past demand to predict future demand. The most basic component of the time series forecast is the level of sales. This is usually determined by looking at demand in the recent past. There often exists an underlying trend that can be detected within sales history. This is usually determined by looking at the change in demand in the recent past. A third factor influencing retail demand is seasonality. A forecasting algorithm trying to determine the effects of seasonality can only look for periodic similarities in the sales pattern. For example, December sales from previous years can be used to adjust the forecast for December only because December occurs regularly every 12 months. At every step, the time series approach is limited to using historical demand to predict future demand without regard to the underlying causes that are driving demand in the first place.

Promotional events, however, can create problems in estimating level, trend, and seasonality. Certain events, such as irregularly occurring holidays, newspaper/radio advertisements, free gift with purchase offers, and special discounts can cause significant deviations from the selling pattern.

Promotional forecasting, unlike statistical forecasting, attempts to predict future demand by identifying the factors that influence spikes and dips of past demand. Once these factors are known, the magnitude and direction of their effect can be evaluated. Their presence can then be incorporated into forecasting models for use during times when the causal factors are again expected to be present.

Developing Promotional Forecast Methods

This section describes how custom promotional forecast models have been developed in the past, leading to the discovery of several consistent findings. These findings have been incorporated into Oracle Retail's development of the Promote forecasting module.

Promotional forecasting uses promotional factors to predict future demand. The first step is to determine all of the pertinent information affecting sales and transform this information into variables that the system understands. Seasonality, for instance, can be represented by a single seasonal continuous variable, such as the number of daylight hours or average daily temperature.

Alternatively, it can be represented by 12 different indicator variables representing each of the months. An indicator variable consists only of 0's and 1's (1 indicates that the event is on). For example, a monthly indicator variable for January would consist of a 1 during the first month of the year and 0's for the remaining months.

Once a list of variables is determined, the model needs to represent the promotion events in terms of their influence on overall demand. For example, if a set of promotional variables has a multiplicative promotional effect on demand, a log transformation may be needed to improve the model. After a suitable model is developed, it must be implemented using multivariate linear regression or neural network architecture with custom code handling the data loading and variable transformations. The final custom model may be quite accurate over the data set on which it was developed.

However, this model may not be general enough to be used universally across all data sets, thus requiring the development of multiple custom models to cover a client's entire domain. This has been found to be very time consuming and costly.

The process of developing custom promotional models has, however, brought to light a number of consistent patterns:

- Level, trend, and seasonality are universal components of almost any forecast.
- Including a time-series forecast as an input variable often improves promotional models.
- Indicator variables are robust in that they can represent both additive and multiplicative effects.

These findings have led Oracle Retail to develop a novel approach to promotional forecasting that combines the automation and generalization of time series forecasting with the improved (albeit, dataset-specific) accuracy met through customized causal forecasting.

Promotional Forecasting Approach

Oracle Retail combines time series forecast methods with causal forecast methods, resulting in a new forecast method supported by the promotional forecasting module. Promote uses the AutoES method of forecast generation to determine a baseline time series forecast and then uses indicator variables to represent promotional events of interest. By giving the forecasting routine visibility as to when certain events occurred in the past, the system can correlate changes in the sales demand with these events and determine each promotional event's effect on sales. Then, whenever these events occur in the future, the promotional effects can be incorporated into the forecast.

The Promotional Forecasting module has been developed to produce generalized promotional models automatically with little human intervention. Combined with the system's ability to allow you to develop your own data loading routines, Promotional Forecasting provides a cost-effective means for producing forecasts using promotional information.

Promotional Forecasting Terminology and Workflow

Promotional Forecasting is designed to produce sales forecasts using both past sales history and event on/off information, both of which you provide. Using the sales data, the system first determines a seasonal time series model to describe the purchasing behavior of consumers. Differences between the seasonal model and the actual sales are then correlated with known events. Events that are found to have a statistically significant impact on sales are then included in a promotional forecast model as promotion events. For each promotion event, its promotion effect on sales is determined. The final promotional model consists of the seasonal model, promotion event on/off information, and each promotion event's resulting effect on sales. By combining these three, a final promotional forecast is computed.

Note: More detailed information on the Causal forecasting algorithm is provided in the [Forecasting Methods Available in RDF](#) section.

Examples of Promotion Events

The following are examples of promotional variables that could be created and the manner in which their associated on/off event status is specified:

Example 13–1 Example 1

Christmas Day applies to all products/locations. The Christmas promotional variable therefore has only one dimension, Day." Because Christmas Day falls on the 359th day of each non-leap year, the Day359 variable is set to True for every such year (all other days it is set to False).

Example 13–2 Example 2

A television advertisement is run locally in the New York/New Jersey area for the four weeks at the beginning of the spring fashion season. The TVAD promotional variable has two associated dimensions, State and Week. Week13, Week14, Week15, and Week16 is set to True only for states NY and NJ (all other states/weeks is set to False). Since no product dimension exists, the TV ad is assumed to have an effect on all products.

Example 13–3 Example 3

A holiday promotion is run involving all sporting goods items for the two weeks prior to Father's Day. The Father's Day promotional variable has two associated dimensions - item and Week. For this year, Week23 and Week24 is set to True only for individual items related to sporting goods items (all other weeks and all other items is set to False). Since no location dimension exists, the Father's Day promotion is assumed to apply to all stores.

After promotional variables have been loaded into RDF, you can use the Promotion Planner task to view, edit, and update associated values without having to reload new data.

Promote Tasks and Wizards

The Promote tasks and wizards allow you to manage the promotion events used in the system's promotional forecasting processes, and view/edit the system's analysis of the effects of these events on demand.

Note: External to Promote, the [Forecast Administration Workbook](#) includes several parameters that may be used as additional configuration options for promotional forecasting. For additional information, refer to [Procedures in Promotional Forecasting](#).

[Table 13–1](#) lists the Promote tasks and wizards.

Table 13–1 Promote Tasks and Wizards

Task Workbooks	Description	Views
Promotion Effectiveness Task	Our promotional What-if task. This allows you to analyze the performance of previous promotions and simulate future promotions by editing when promotional events are active for an item/location and by modifying the promotion lift effects.	Promotion Model Type View Promotions View Aggregated Promotion Lift Override View Promotional Forecast View Promotion Parameters View
Promotion Maintenance Task	Allows you to review the system-calculated promotional lift effects, edit these effects, and determine how changes are factored into the promotional model.	Promotion Parameters View
Promotion Management Task	Allows you to enable and disable promotions for causal forecast levels and to specify whether promotions can have a negative effect.	Promo Variable Model Type View Promotion Enable View Accept Negative Lift View
Promotion Planner Task	Allows you to specify when certain promotional events were active in the past, and when they are active in the future.	Promotions View

Promotion Effectiveness Task

The Promotion Effectiveness task is a historical and future view to the effects of a planned promotion.

This task has these views:

- [Promotion Model Type View](#)
- [Promotions View](#)
- [Aggregated Promotion Lift Override View](#)
- [Promotional Forecast View](#)
- [Promotion Parameters View](#)

In this task, you view the promotional forecast. There is also the ability to analyze the effects on the forecast if a promotion does or does not occur. This analysis is performed by turning a promotional event or attribute on or off for dates in the future plans, and/or by modifying the Simulated Effect. This task is intended for simulation and analysis purposes only. There is no Commit functionality.

Note: Changes to the Promotion Effectiveness task cannot be committed back to the master database; however, this task is useful for reporting on the performance of past promotions and simulating the effects of future promotions to support more accurate promotional modeling.

Promotion Effectiveness Workbook

To build the Promotion Effectiveness workbook, perform these steps:

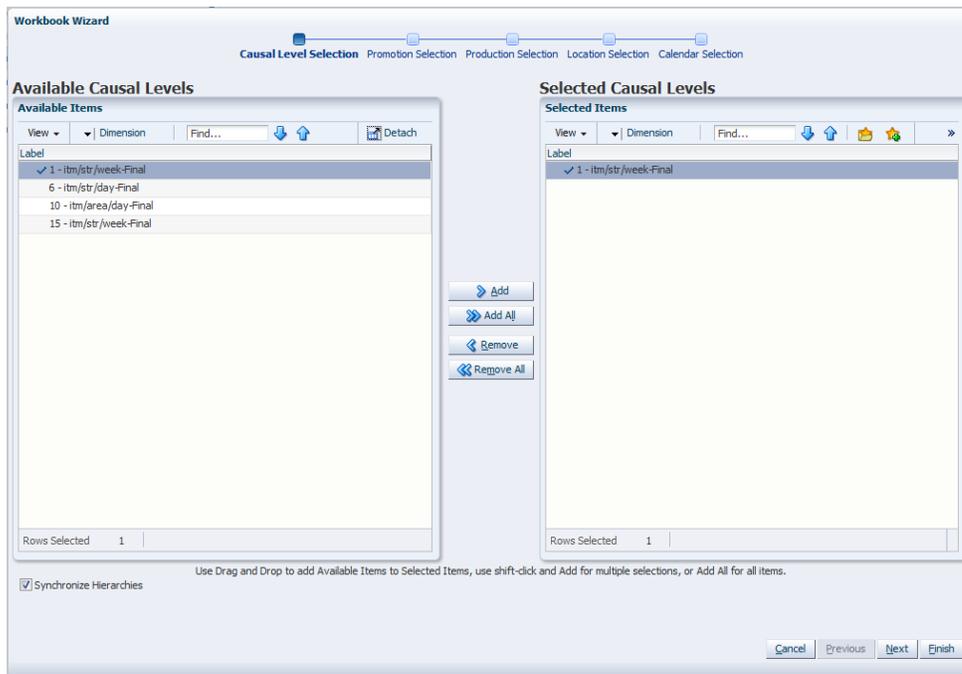
1. Click the New Workbook icon in the Promotion Effectiveness task.

Figure 13–1 Promotion Effectiveness Task



2. The Workbook wizard opens. Select the causal level and click **Next**.

Figure 13–2 Workbook Wizard: Causal Level



3. Select the promotions to include and click **Next**.

Note: If RDF detects cross-promotional effects for your selected causal level, then RDF uses those promotions and the wizard skips to "Workbook Wizard: Select Products".

Figure 13–3 Workbook Wizard: Select Promotions

Workbook Wizard

Causal Level Selection **Promotion Selection** Production Selection Location Selection Calendar Selection

Available Promotions

Available Items

View ▾ | Dimension ▾ | Find... | [Down Arrow] [Up Arrow] [Detach]

Label

- Back to School
- Christmas
- Christmas Season
- ✓ Circular
- Direct Mail
- Easter
- ✓ In-Store Display
- ✓ New Years Day
- ✓ Promotional Price
- Signage

[Add] [Add All] [Remove] [Remove All]

Rows Selected 1

Selected Promotions

Selected Items

View ▾ | Dimension ▾ | Find... | [Down Arrow] [Up Arrow] [Add] [Remove] [Detach]

Label

- ✓ Circular
- ✓ In-Store Display
- ✓ New Years Day
- ✓ Promotional Price

Rows Selected 1

Use Drag and Drop to add Available Items to Selected Items, use shift-click and Add for multiple selections, or Add All for all items.

Synchronize Hierarchies

[Cancel] [Previous] [Next] [Finish]

4. Select the products to include and click **Next**.

Figure 13–4 Workbook Wizard: Select Products

Workbook Wizard

Causal Level Selection Promotion Selection **Production Selection** Location Selection Calendar Selection

Available Product

Available Items

View ▾ | Dimension ▾ | Find... | [Down Arrow] [Up Arrow] [Detach]

Label

- ✓ Private Label Pumpkin 15 oz
- Private Label Sliced Peaches 8.5 oz
- Private Label Sliced Peaches 14.5 oz
- Private Label Tomatoes 14.5 oz
- Private Label Tomatoes 28 oz
- Private Label Tomato Paste 5 oz
- Private Label Tomato Sauce 8 oz
- Private Label Tomato Sauce 15.5 oz
- Private Label Whole Tomatoes 28 oz
- Q S Laundry Stain Remov Trig32
- Snuggles 64 oz
- Snuggles Cuddle up 70 Ct
- Snuggles Fabric Softner 35 Ld
- Snuggles Fresh Rain 64 oz
- Snuggles Shts Fresh Rain 70 Ct
- Solid Pack Pumpkin 15 oz
- Surf Powder w/BI 32 oz
- Sweet Potatoes 13 oz
- Tide Liquid Orig Scent 32 loads
- Tide Regular 40 loads
- Tide Ultra Scented Scent 30 loads

[Add] [Add All] [Remove] [Remove All]

Rows Selected 1

Selected Product

Selected Items

View ▾ | Dimension ▾ | Find... | [Down Arrow] [Up Arrow] [Add] [Remove] [Detach]

Label

- ✓ 24 X 16 oz Coke Crate
- ✓ Fruit in the Bottom LF 8 OZ:Strawberry
- ✓ Fruit in the Bottom LF 27 OZ:Banana
- ✓ Fruit in the Bottom LF 27 OZ:Blueberry
- ✓ Fruit in the Bottom LF 27 OZ:Cherry
- ✓ Fruit in the Bottom LF 27 OZ:Strawberry
- ✓ Private Label Pumpkin 15 oz

Rows Selected 1

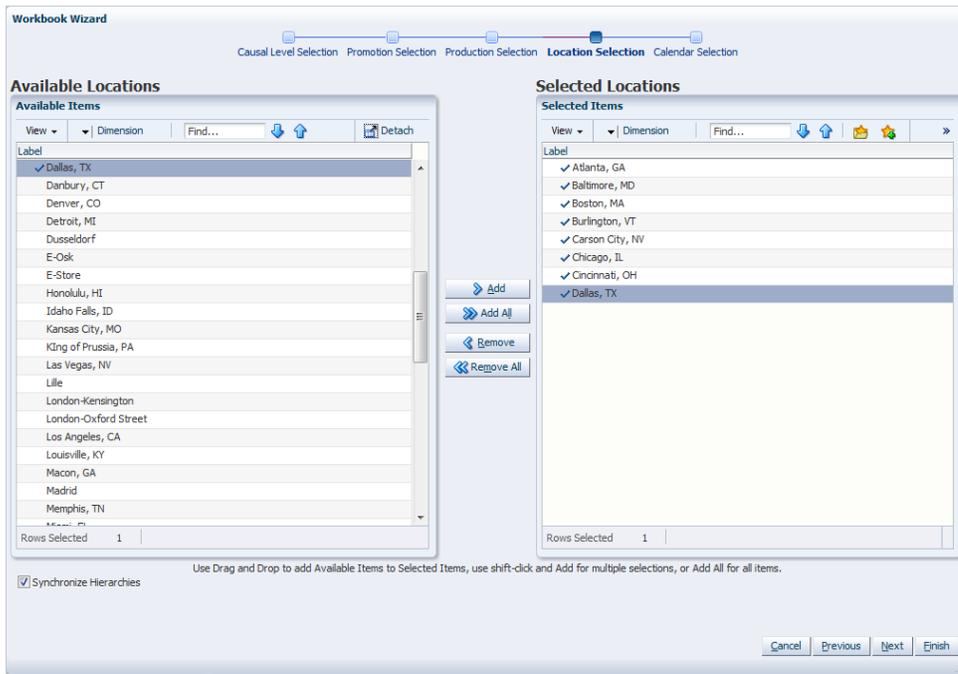
Use Drag and Drop to add Available Items to Selected Items, use shift-click and Add for multiple selections, or Add All for all items.

Synchronize Hierarchies

[Cancel] [Previous] [Next] [Finish]

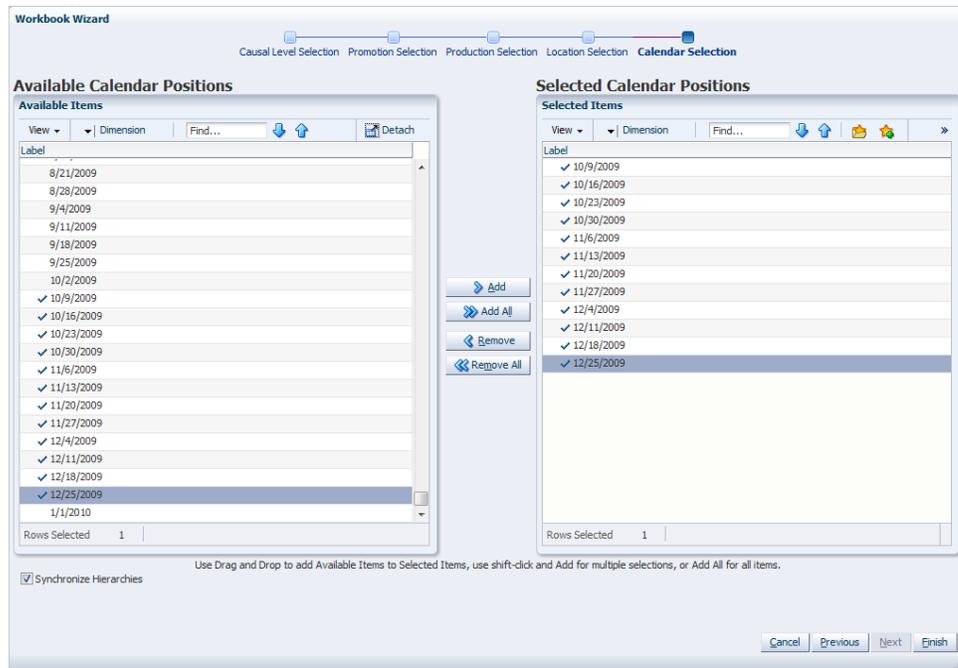
5. Select the locations and click **Next**.

Figure 13–5 Workbook Wizard: Locations



6. Select the days you want to work with and click **Finish**.

Figure 13–6 Workbook Wizard: Select Calendar Positions



The Promotion Effectiveness workbook is built.

Promotion Model Type View

The Promotion Model Type view allows you to review the model type.

Figure 13–7 Promotion Model Type View

	Circular	In-Store	New Years	Promotional
Model Type for promotional variable	Linear	Linear	Linear	Exponential

Measures: Promotion Model Type View

The Promotion Model Type view contain the following measures:

Model Type for Promotional Variable

This measure stores the model type for the selected promotion variables. The choices are exponential and linear. An exponential type is displayed if the promotion variable is nonlinear. A linear type is displayed when the promotion variables are real or Boolean.

For additional information, see "[Promotion Variables](#)" on page 13-25.

Promotions View

The Promotions view allows you to review and enable or disable promotions for item/locations for desired periods. Once the changes in the promo calendar have been made, click Calculate to update the forecast values in the Promotional Forecast view.

The promotions must be enabled for the chosen causal level. Note that multiple Promotions views are possible. Namely, there is one view for every distinct promotion intersection. Promotions with identical intersections are displayed in the same view.

Figure 13–8 Promotions View

	10/9/2009	10/16/2009	10/23/2009	10/30/2009	11/6/2009	11/13/2009	11/20/2009	11/27/2009	12
Promotion Variable Circular	<input type="checkbox"/>								
Promotion Variable In-Store	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	
Promotion Variable	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	

Measures: Promotions View

The Promotions view contain the following measure:

Promotion Variables

Promotion Variables are defined as either Boolean (by default) or Real types during the configuration process in the Promote Plug-In. Refer to the *Oracle Retail Demand Forecasting Configuration Guide* for additional details on configuring the Promote solution.

For additional information, see "[Promotion Variables](#)" on page 13-25.

Aggregated Promotion Lift Override View

The Aggregated Promotion Lift Override view allows you to review the promotional lift overrides, which can be generated at a low as well as an aggregate level in the hierarchy.

Figure 13–9 Aggregated Promotion Lift Override View

	10/9/2009	10/16/2009	10/23/2009	10/30/2009	11/6/2009	11/13/2009
Aggregated Promotion Lift	0.00	0.00	0.00	0.00	0.00	0.00
Promotional Lift Override 1 -	0.00	0.00	0.00	0.00	0.00	0.00

Measures: Aggregated Promotion Lift Override View

The Aggregated Promotion Lift Override view contains the following measures:

Aggregated Promotion Lift Override

This measure allows you to override (in units) the system-generated promotional lift. The override is at an aggregate level and spread down to the forecast level.

Promotional Lift Override

This measure allows you to override (in units) the system generated promotional lift. The override is at the same intersection as the forecast.

This measure can be equal to the [Aggregated Promotion Lift Override](#) or zero (0). It is generated based on the Aggregated Promotion Lift Override and the Apply Promotion Lift Override Indicator. If the indicator is set to **true** (which applies the override), then this value should be the same as the spread down [Aggregated Promotion Lift Override](#). If the indicator is **false** (which does not apply the override), then this value equals zero (0).

Promotional Forecast View

The Promotional Forecast view allows you visibility to all components of the forecast, like Self, Halo, and Cannibalization effects for both regular price changes, as well as promotions. This is the place the user reviews the impact of the promo-related parameters set in the Promotions and Promotion Parameters views. The process of changing parameters, re-generating the promo forecast, and reviewing the impact in this view, implements the promotional what-if functionality.

Figure 13–10 Promotional Forecast View

	10/9/2009	10/16/2009	10/23/2009	10/30/2009	11/6/2009	11/13/2009
Forecasted Baseline 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00
Promotional Forecast with Promote Lift Override 1 -	0.00	0.00	0.00	0.00	0.00	0.00
Total Promotional Lift Unit 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00
Total Regular Price Lift Unit 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00
Price Self Lift Units 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00
Price Halo Lift Units 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00
Price Cannibalization Lift Unit 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00
Promotion Self Lift Units 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00
Promotion Halo Lift Units 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00
Promotion Cannibalization Lift Unit 1 - itm/str	0.00	0.00	0.00	0.00	0.00	0.00
Promotional Lift Override 1 - itm/str/week-Final	0.00	0.00	0.00	0.00	0.00	0.00

Measures: Promotional Forecast View

The Promotional Forecast view can contain the following measures:

Forecast Baseline

The Forecast Baseline is a forecast generated on past sales data that contains no promotions (that is, normal demand given no causal effects).

Price Cannibalization Lift Unit

This measure displays the negative effect, measured in units, that a price change of certain items has on this item.

Price Halo Lift Units

This measure displays the positive effect, measured in units, that a price change of certain items has on this item.

Price Self Lift Units

This measure displays the effect that a price change has on the sales units of an item.

Promotion Cannibalization Lift Units

This measure displays the negative effect, measured in units, that the promotion of certain items has on this item.

Promotion Halo Lift Units

This measure displays the positive effect, measured in units, that the promotion of certain items has on this item.

Promotion Self Lift Units

This measure displays the effect that a promotion has on the sales units of an item.

Promotional Forecast with Promote Lift Override

This measure displays the promotional forecast where the promotional peaks are not calculated by RDF, but imported as lift overrides.

Promotional Lift Override

Override, in units, of the system-generated promotional lift. The override is at the same intersection as the forecast.

This measure can be equal to the [Aggregated Promotion Lift Override](#) or zero (0). It is generated based on the Aggregated Promotion Lift Override and the Apply Promotion Lift Override Indicator. If the indicator is set to **true** (which applies the override), then this value should be the same as the spread down [Aggregated Promotion Lift Override](#). If the indicator is **false** (which does not apply the override), then this value equals zero (0).

Total Promotional Lift Unit

This measure displays the total effect that a promotion can have on an item's sales units. It is the sum of Promotion Self Lift Units, Promotion Halo Lift Units, and Promotion Cannibalization Lift Units.

Total Regular Price Lift Unit

This measure displays the total effect that a regular price change can have on an item's sales units. It is the sum of Regular Price Self Lift, Regular Price Halo Lift, and Regular Price Cannibalization Lift.

Promotion Parameters View

The Promotion Parameters view allows you to set up rules that need to be followed when forecast is simulated. For instance, you may not want to apply regular price change effects when an item is on promotion or you may want to incorporate Halo and no Cannibalization effects in the forecast. click **Calculate** to update the forecast values in the Promotional Forecast view.

Figure 13–11 Promotion Parameters View

	24 X 16 oz	Fruit in the	Fruit in the	Fruit in the	Fruit in the	Fruit in th
System Calculated Effect 1 - itm/str/week-Final Circular	1.00	1.00	1.00	1.00	1.00	1
System Effect Override 1 - itm/str/week-Final Circular	1.00	1.00	1.00	1.00	1.00	1
Promo Effect Type 1 - itm/str/week-Final Circular	Automatic	Automatic	Automatic	Automatic	Automatic	Automatic
Apply Regular Price Lift Final 1 - itm/str/week-Final	Apply Sytem	Apply Syte				
Apply Regular Price Lift over Promotion Final 1 - itm/str/week-Final	<input type="checkbox"/>					
Apply Regular Price Self Lift Final 1 - itm/str/week-Final	<input checked="" type="checkbox"/>					
Apply Regular Price HALO Lift Final 1 - itm/str/week-Final	<input checked="" type="checkbox"/>					
Apply Regular Price Cannibalization Lift Final 1 - itm/str/week-Final	<input checked="" type="checkbox"/>					
Apply Promotion Halo Lift Final 1 - itm/str/week-Final	<input type="checkbox"/>					
Apply Promotion Cannibalization Lift Final 1 - itm/str/week-Final	<input type="checkbox"/>					
Apply Promotional Lift Override 1 - itm/str/week-Final	<input type="checkbox"/>					

Measures: Promotion Parameters View

The Promotion Parameters view contains the following measures:

Apply Promotion Cannibalization Lift

Select this parameter (set to **true**) if the application of Cannibalization Lifts is enabled for the forecast level.

Apply Promotion Halo Lift

Select this parameter (set to **true**) if the application of Halo lifts is enabled for the forecast level.

Apply Promotional Lift Override

Select this parameter (set to **true**) if you want RDF to ignore the system calculated lifts and apply the overrides.

Apply Regular Price Cannibalization Lift

Is only effective if [Apply Regular Price Lift](#) is set to Apply System or Apply Override. Select this parameter (set to **true**) if the application of Regular Price Cannibalization Lifts is enabled for the forecast level.

Apply Regular Price Halo Lift

Is only effective if [Apply Regular Price Lift](#) is set to Apply System or Apply Override. Select this parameter (set to **true**) if the application of regular price Halo lifts is enabled for the forecast level.

Apply Regular Price Lift

If regular price lifts are enabled for the forecast level, select one of the following options:

Option	Description
Do Not Apply	Does not apply regular price effects (including Self, Cannibalization, and Halo effects)
Apply System	Applies system calculated effects
Apply Override	Ignores system calculated effects and applies overrides

Apply Regular Price Lift Over Promotion

Is only effective if [Apply Regular Price Lift](#) is set to Apply System or Apply Override. Select this parameter (set to **true**) if the application of regular price lifts is enabled for time periods that are promoted. If the parameter is clear, then regular price lifts are not applied for periods with promotions.

Apply Regular Price Self Lift

Is only effective if [Apply Regular Price Lift](#) is set to Apply System or Apply Override. Select this parameter (set to **true**) if the application of regular price lifts is enabled for the forecast level.

Promo Effect Type

Promo Effect types define how Promo Effects are treated in the causal model-fitting process (which includes a call to the lower-level regression engine) and the forecast generation process where the model is used to extend the forecast over the forecast horizon.

[Table 13–2](#) lists the Promo Effect types.

System Calculated Effect

The System Calculated Effect is a read-only measure indicating the lift effect generated by the system.

System Effect Override

The user-specified lift effect. This user-entered effect is active if used in conjunction with the Override All and Override Future Only Promotion Effect Types. Otherwise, it equals 1.00.

Promotion Maintenance Task

The Promotion Maintenance task provides a view to the system-calculated and adjusted lift effects. You can edit effects at any product/location intersection and determine how these changes are factored into the promotional models. The Promotion Maintenance task contains one view. There may be multiple versions of this view, defined at various causal levels.

This task has this view:

- [Promotion Parameters View](#)

Promotion Maintenance Workbook

To build the Promotion Maintenance workbook, perform these steps:

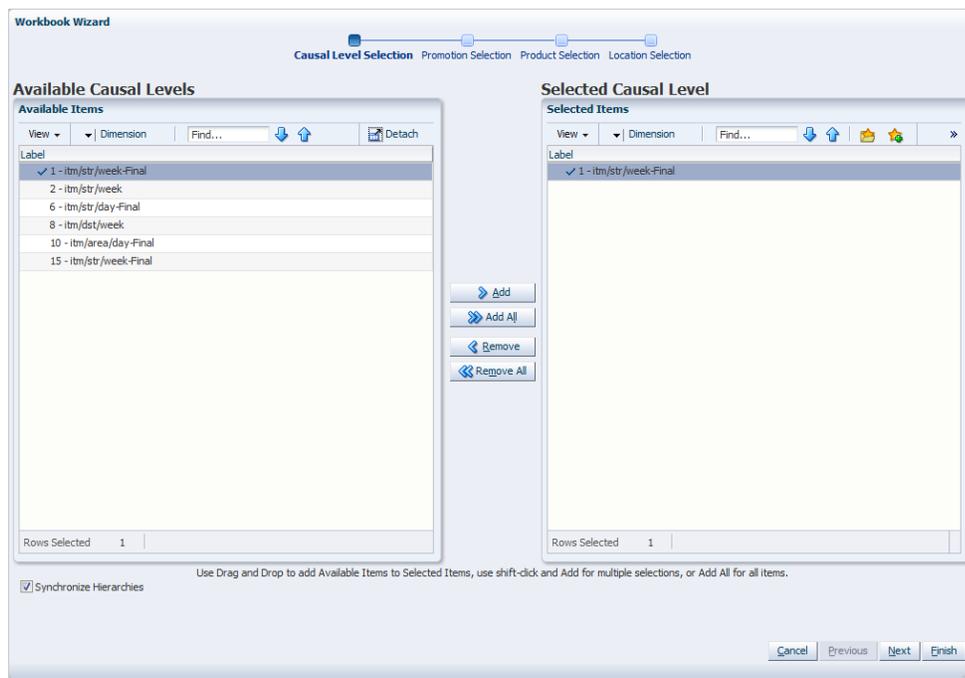
1. Click the New Workbook icon in the Promotion Maintenance task.

Figure 13–12 Promotion Maintenance Task



2. The Workbook wizard opens. Select the causal level and click **Next**.

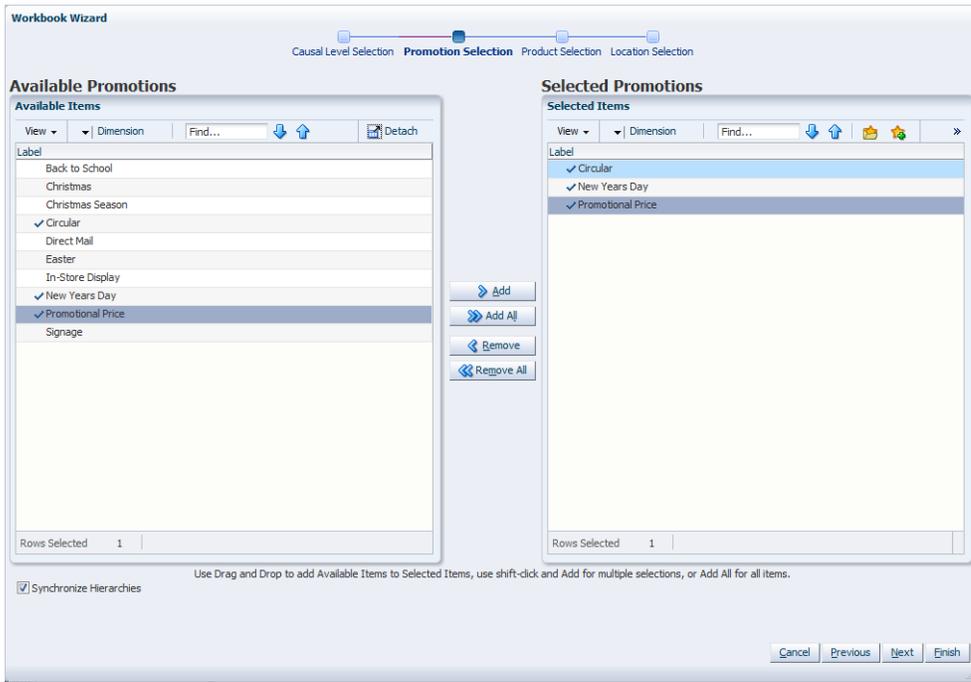
Figure 13–13 Workbook Wizard: Causal Level



3. Select the promotion events to include and click **Next**.

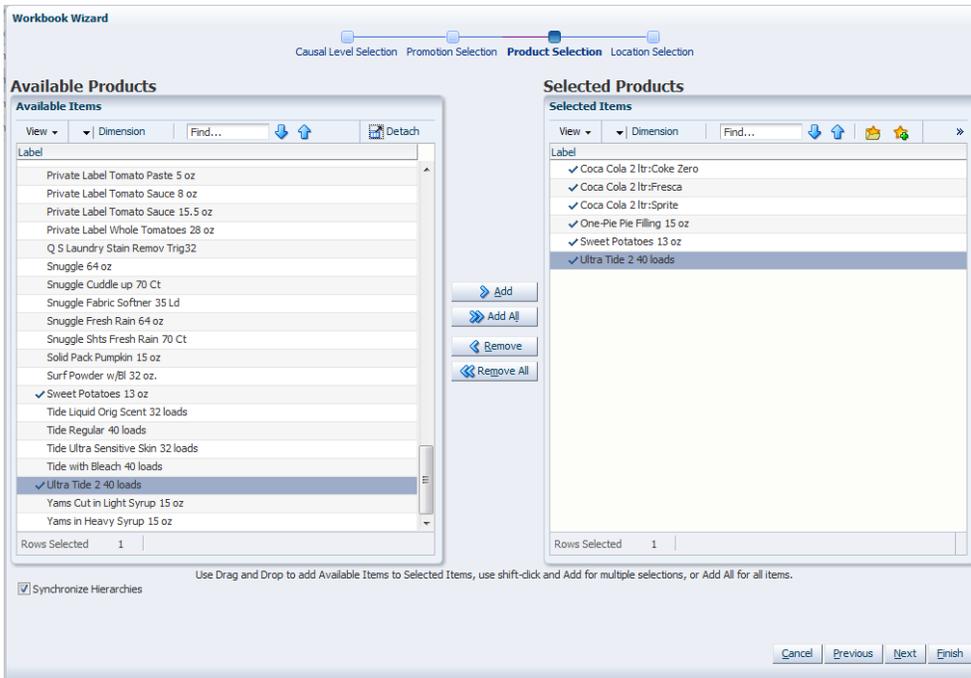
Note: If RDF detects cross-promotional effects for your selected causal level, then RDF uses those promotions and the wizard skips to "Workbook Wizard: Select Products".

Figure 13–14 Workbook Wizard: Select Promotions



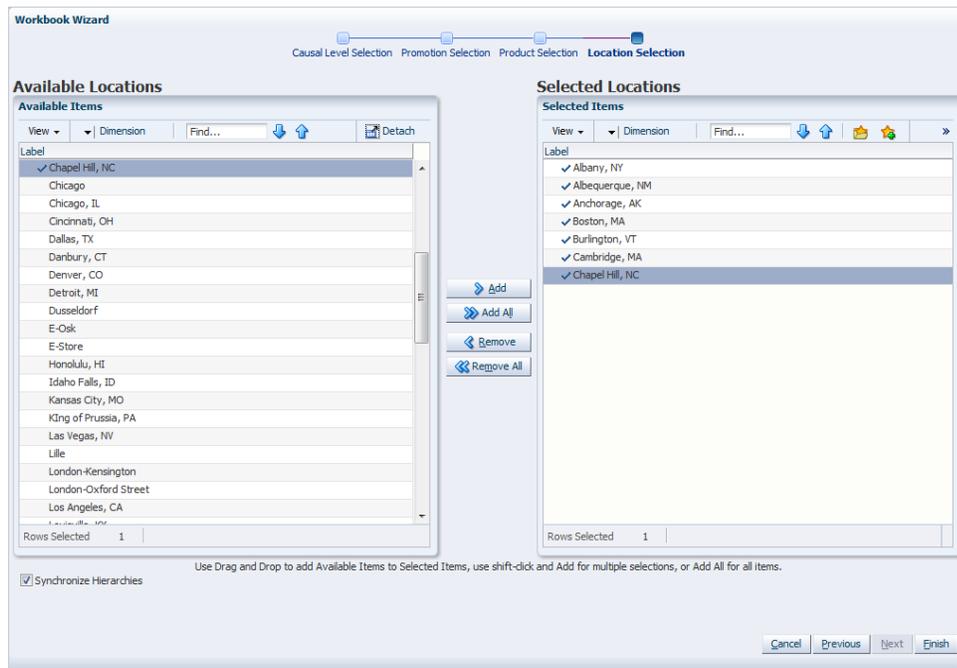
4. Select the products you want to work with and click **Next**.

Figure 13–15 Workbook Wizard: Select Products



5. Select the locations you want to work with and click **Finish**.

Figure 13–16 Workbook Wizard: Select Locations



The Promotion Maintenance workbook is built.

Promotion Parameters View

The Promotion Parameters view allows you to view and modify the system-calculated effects of a given promotion.

Note: Changes that are committed in the Promotion Maintenance task are not incorporated into the forecast until another forecast is generated in RDF batch forecast process.

Figure 13–17 Promotion Parameters View

Promotion Parameters for Level 1 - itm/str/week-Final					
Product	Location				
Coca Cola 2 ltr:Coke Zero	Albany, NY	Albuquerque, NM	Anchorage, AK	Boston, MA	Burlington, VT
Promo Effect Type 1 - itm/str/week-Final Circular	Automatic	Automatic	Automatic	Automatic	Automatic
System Effect Override 1 - itm/str/week-Final Circular	1.00	1.00	1.00	1.00	1.00
System Calculated Effect 1 - itm/str/week-Final Circular	1.00	1.00	1.00	1.00	1.00

Measures: Promotion Parameters View

The Promotion Parameters view contains the following measures:

Promo Effect Type

Promo Effect types define how promotion effects are treated in the causal model-fitting process (which includes a call to the lower-level regression engine) and the forecast generation process where the model is used to extend the forecast over the forecast horizon. [Table 13–2](#) lists the options.

System Calculated Effect

The System Calculated Effect is a read-only measure indicating the lift effect generated by the system.

System Effect Override

The user-specified lift effect. This user-entered effect is active if used in conjunction with the Override All and Override Future Only Promotion Effect Types. Otherwise, it equals 1.00.

Table 13–2 Promo Effect Types

Promo Effect Type	Description
Automatic	The inclusion of the Promo Effect is decided by regression. If the Promo Effect is found to be significant on the training set, it is included in the model. Otherwise, it is rejected. Automatic is the system default Promotion Effect Type.
Forced In	The Promo Effect is forced in to the model, thus regression is not given a choice to reject even if the effect is considered insignificant by regression. As a result, we will always return an effect even if it has a negative impact to the demand forecast.
Disabled	The variable is excluded from the model, hence no effect is returned either.
Override All	This type allows you to specify a causal effect that is used during the forecasting process. For Promo Effects specified as Override All, you also specify the corresponding causal effect in the Promotion Effect Override. The causal engine de-causalizes the training data using the user-specified effect. During forecast generation, the user-specified effect is used to determine the causal forecast. Therefore, you must change the Promotion Effect Type when this user-specified effect is no longer to be used.

Table 13–2 (Cont.) Promo Effect Types

Promo Effect Type	Description
Override Future Only	This type allows you to specify a causal effect that is used during the forecasting process. For Promo Effects specified as Override Future, you also specify the corresponding causal effect in the Promotion Effect Override. The calculated effect is not written back to the effects measure, but it is used to de-causalize the data. During forecast generation the calculated effect is ignored, and, instead, the user-specified effect is used to produce a causal forecast. Therefore, you must change the Promotion Effect Type when this user-specified effect is no longer to be used.
Override from Higher Level	This promotion type is used in conjunction with the Causal Higher Intersection set in the Forecast Administration Workbook . If the Causal Higher Intersection is not specified, no promotional effect is calculated. It allows the system to use the geometric mean of causal effects computed from product/location combination in the same group (the intersection level specified in the Causal Higher Intersection during system setup) for product/location combinations for which from sales history alone a causal effect is unable to be computed for that Promo Effect. The way the system handles this, is by having the override effects measure filled in with higher-level effects for those variables specified as Override Higher Level. During forecast generation the effect calculated as average of low level effects is used to produce a causal forecast. Note that if no effects are found for any of the product/locations in the Causal Higher intersection group, then the system uses the geometric mean of causal effects calculated from all Product/Locations as the override effect being passed in.
Automatic Boolean	If the promotion variable is always be set to 0.00 or 1.00 (meaning inactive or active, respectively), this type produces the same results as the Automatic type; however, Automatic Boolean improves the performance (speed) of the forecasting engine during the batch run of the forecast.

Promotion Management Task

The Promotion Management task allows you to enable and disable promotions for causal forecast levels and to specify whether promotions can have a negative effect.

This task has these views:

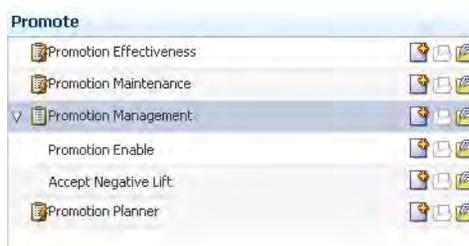
- [Promo Variable Model Type View](#)
- [Promotion Enable View](#)
- [Accept Negative Lift View](#)

Promotion Management Workbook

To build the Promotion Management workbook, perform these steps:

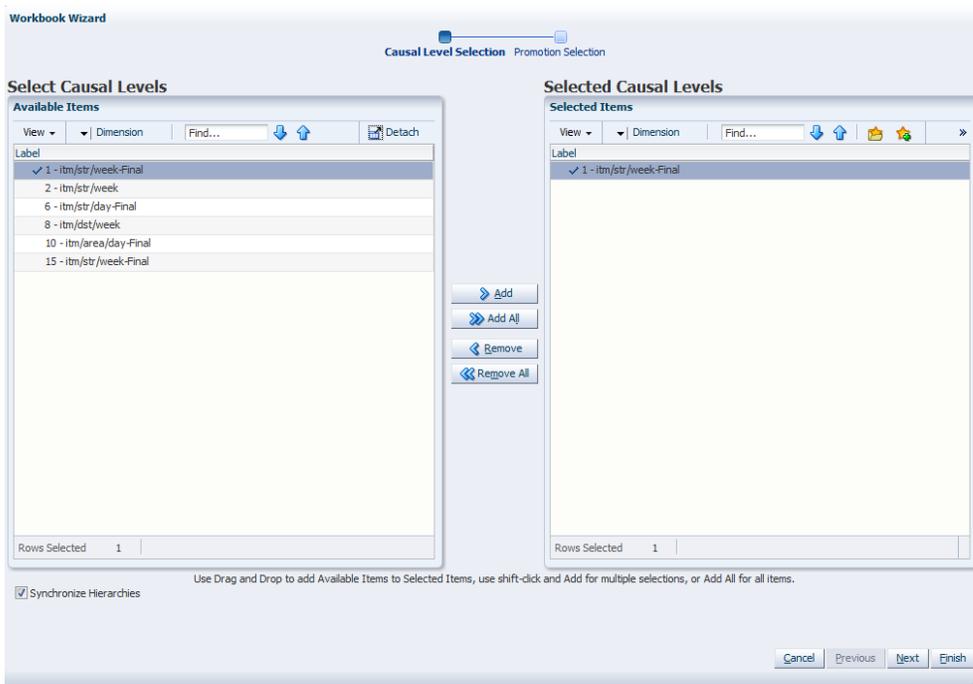
1. Click the New Workbook icon in the Promotion Management task.

Figure 13–18 Promotion Management Task



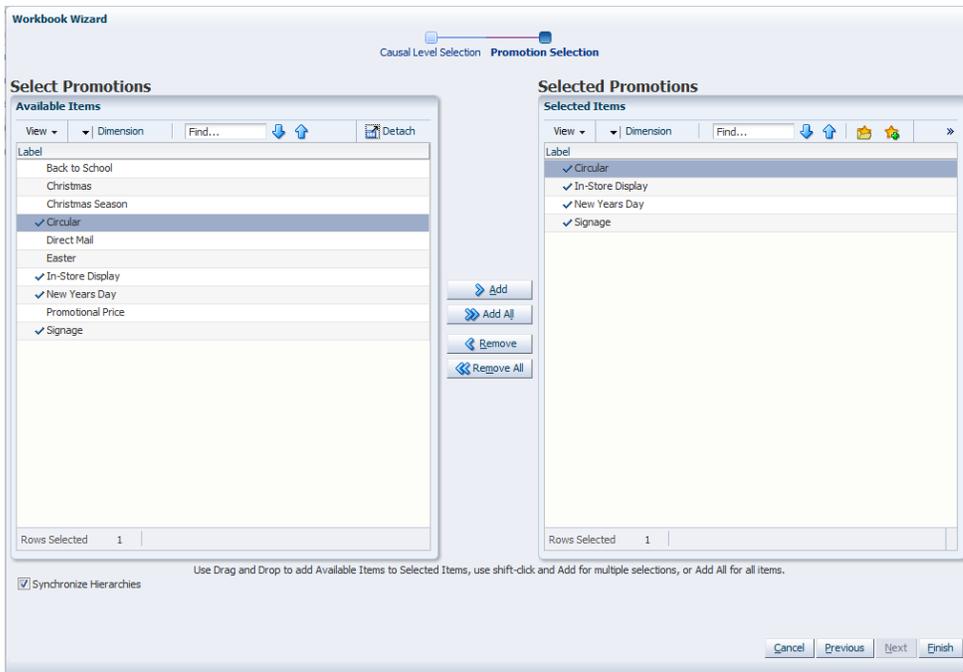
2. The Workbook wizard opens. Select the causal level and click **Next**.

Figure 13–19 Workbook Wizard: Select Causal Level



3. Select the promotions you want to work with and click **Finish**.

Figure 13–20 Workbook Wizard: Select Promotions

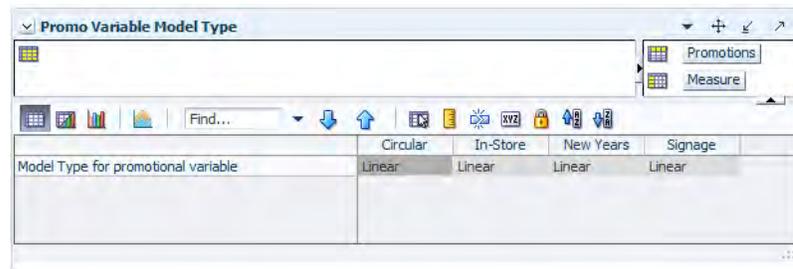


The Promotion Management workbook is built.

Promo Variable Model Type View

The Promo Variable Model Type view allows you to view the model type for the selected promotion variables.

Figure 13–21 Promo Variable Model Type View



Measure: Promo Variable Model Type View

The Promo Variable Model Type view contains the following measure:

Model Type for Promotional Variable

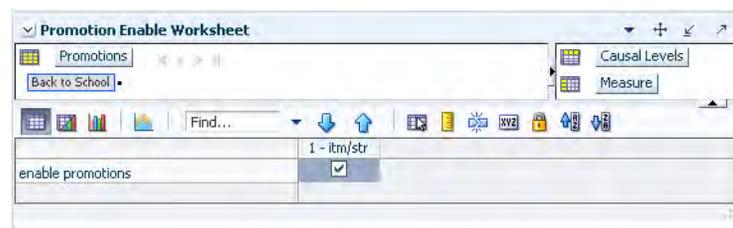
This measure stores the model type for the selected promotion variables. The choices are exponential and linear. An exponential type is displayed if the promotion variable is nonlinear. A linear type is displayed when the promotion variables are real or Boolean.

For additional information, see "[Promotion Variables](#)" on page 13-25.

Promotion Enable View

The Promotion Enable view allows you to enable or disable a promotion for the causal forecast levels. The Promotion Enable view is built at the hierarchy intersections of the promotion variables and the causal levels selected during the wizard process.

Figure 13–22 Promotion Enable View



Measures: Promotion Enable View

The Promotion Enable view contains the following measure:

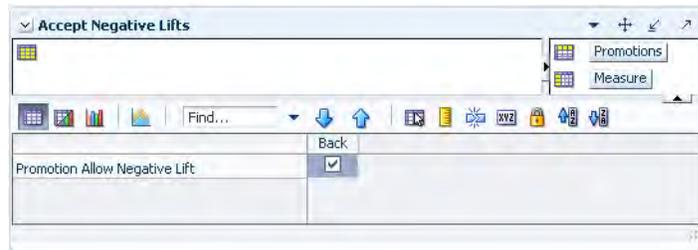
Enable Promotions

This is a read/write Boolean measure. The Enable Promotions measure allows you to enable a subset of promotions for a certain causal forecast level. It defaults to **true** for all causal forecast levels.

Accept Negative Lift View

The Accept Negative Lift view allows you to specify whether a promotion is allowed to have negative effect. The Promotion Enable view is built at the hierarchy intersections of the promotion variables selected during the wizard process.

Figure 13–23 *Accept Negative Lift View*



Measures: Accept Negative Lift View

The Accept Negative Lift view contains the following measure:

Promotion Allow Negative Lift

This is a read/write Boolean measure. The Promotion Allow Negative Lift measure allows you to specify whether a promotion is allowed to have negative effect. When this measure is set to **false** for a promotion, the forecast engine does not allow that promotion to have any negative lift (Effect < 1). It defaults to **true** for all promotions.

Promotion Planner Task

In order to correlate deviations from the seasonal forecast with the occurrence of historic promotion events, the system needs visibility as to when these events were active. The system must also be informed of dates on which the status of upcoming promotion events will again be on, so the anticipated promotion effects can be built into the forecasting model.

The Promotion Planner task allows you to indicate to the system when certain events were active in the past and when they are active in the future. All promotional events should be represented as accurately as possible so the modeling routine can more precisely detect correlations between event occurrences and changes in sales values.

The Promotion Planner task consists of as many views as are necessary to represent all unique dimensional intersections associated with the promotion events contained in the task. A separate view is constructed for each of the required intersections. For example, promotion events such as Advertisement and Gift with Purchase may be loaded at the item/store/week intersection, while an event such as Christmas is loaded at the [Day] level.

In this setup, the Advertisement and Gift with Purchase promotions would appear on one view, and Christmas would appear on another. Whenever a hierarchy is not included in the base intersection (as in the case of the Christmas promotional event) the event is assumed to apply to all positions in the undefined hierarchy. Thus, Christmas is assumed to apply to all products and all locations, but only to the Day-level calendar positions specified in the [Promotions View](#).

This task has this view:

- [Promotions View](#)

Promotion Planner Workbook

To build the Promotion Planner workbook, perform these steps:

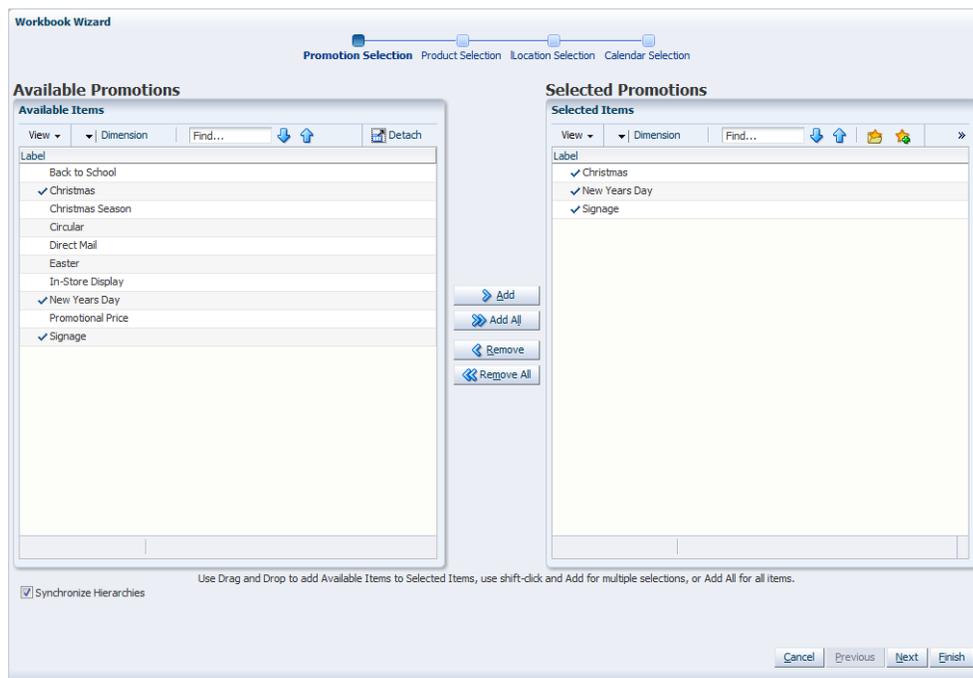
1. Click the New Workbook icon in the Promotion Planner task.

Figure 13–24 Promotion Planner Task



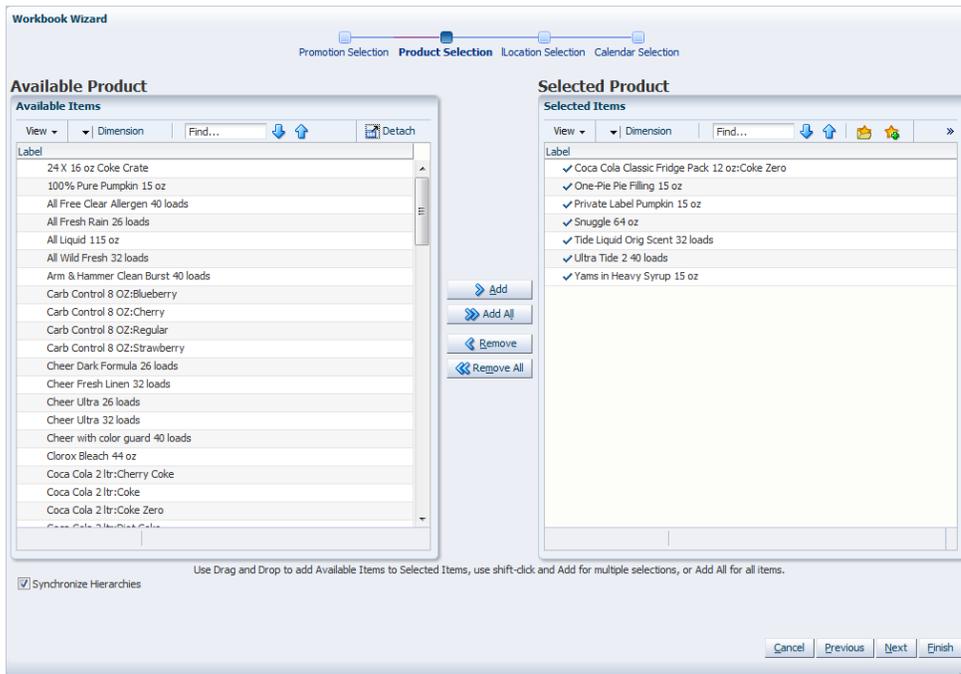
2. The Workbook wizard opens. Select the promotion events to include and click Next.

Figure 13–25 Workbook Wizard: Select Promotions



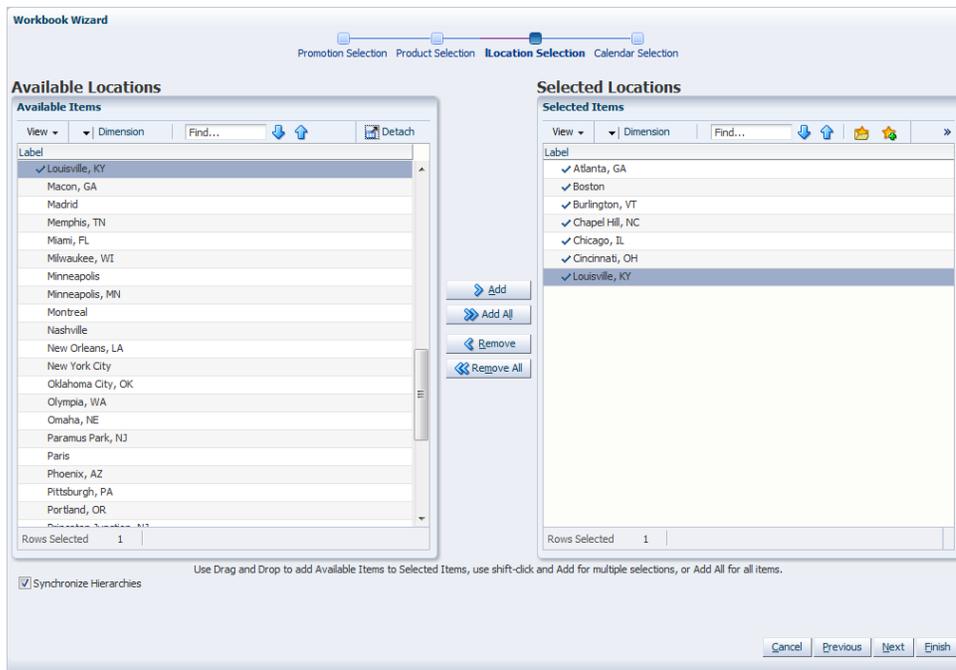
3. Select the products to include and click Next.

Figure 13–26 Workbook Wizard: Select Products



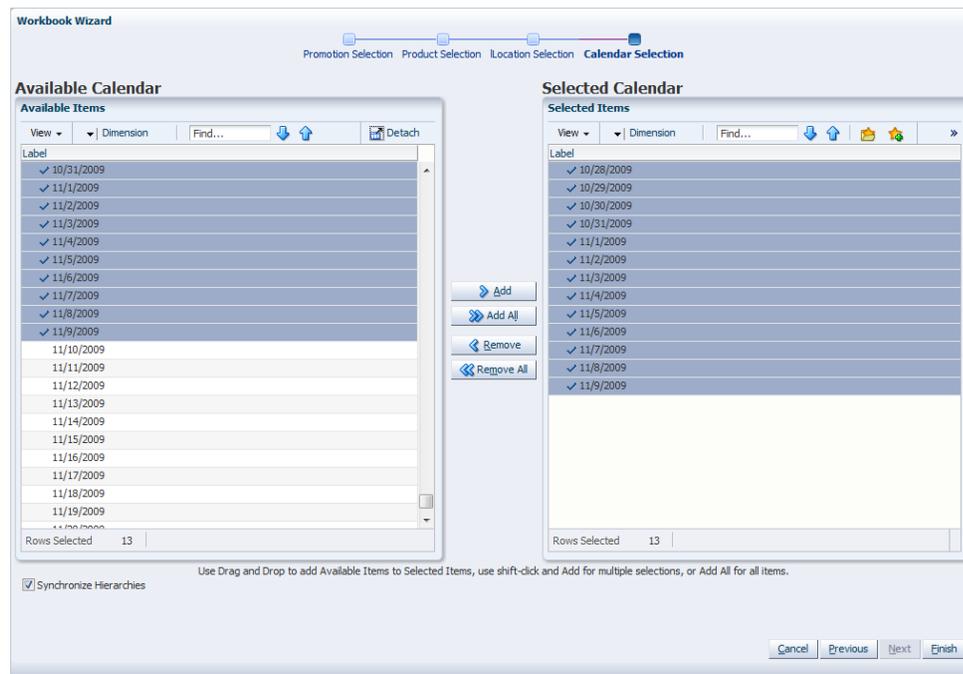
4. Select the locations to include and click **Next**.

Figure 13–27 Workbook Wizard: Select Location



5. Select the days you want to work with and click **Finish**

Figure 13–28 Workbook Wizard: Select Days



The Promotion Planner workbook is built.

Promotions View

The Promotions view allows you to view and edit the on/off information associated with each configured promotional event. This view provides an interface in which you can specify the time periods (and possibly products or locations) for which certain promotional variables are active.

Figure 13–29 Promotions View

	10/28/2009	10/29/2009	10/30/2009	10/31/2009	11/1/2009	11/2/2009	11/3/2009
Promotion Variable Christmas	<input type="checkbox"/>						

Measures: Promotions View

The Promotions view contains the following measures:

Promotion Variables

Promotion Variables are defined as either Boolean (by default) or Real or nonlinear types during the configuration process in the Promote Plug-In. Refer to the *Oracle Retail Demand Forecasting Configuration Guide* for additional details on configuring the Promote solution.

Promotion Variable Type	Description
Boolean	A check in a given cell indicates that the associated promotion event's status is on (or 100% of the lift effect applies) for that intersection. If no check is indicated, the event's status is off.
Real	A value of 1.00 in a given cell indicates that the associated promotion event's status is on (or 100% of the lift effect applies) for that intersection. A value of 0.00 indicates that the event's status is off.
Exponential	The value in a given cell can be any number. There are no restrictions as there are for Boolean or Real variables. A typical example of a metric modeled as exponential is the percent price change. Note how the percent price change can be positive (price increase) or negative (price decrease).

About Implementation

Among the ways Promo Effects can be implemented include price or discount percent. Your Oracle Retail Consultant can best determine the most accurate setup of promotion variables based upon your promotional forecasting requirements.

Procedures in Promotional Forecasting

The following list outlines the standard procedures performed in order to set up the system to run a promotional forecast:

- [Setting Up the System to Run a Promotional Forecast](#)
- Setting forecast parameters in the [Forecast Administration Workbook](#)
- Setting forecast parameters in the [Forecast Maintenance Workbook](#)
- Setting promotions to be active in the [Promotion Planner Task](#)
- Running the batch forecast
- Viewing and Edit Causal Forecast results
- Analyzing forecasts in the [Forecast Approval Workbook](#)
- Analyzing and editing causal effects in the [Promotion Maintenance Task](#)
- [Viewing a Forecast That Includes Promotion Effects](#)
- [Viewing and Editing Promotion System-Calculated Effects](#)
- [Promotion Simulation \(What-if?\) and Analysis](#)

Setting Up the System to Run a Promotional Forecast

Perform the following procedure to set up the system to run a promotional forecast.

1. Build a [Forecast Administration Workbook](#).
 - a. In the [Forecast Administration Workbook: Final Level Parameters View - Advanced Settings](#), select [Generate Baseline Forecasts](#).
 - b. Optional: **Set the Default Forecast Method to Causal** for the desired level if the level is to be use only for Promotional Forecasting.
 - c. Optional: **Set the Causal Higher Intersection** for the desired level if the Override from Higher Level promotion type is used.

- d. Optional for use with Daily Causal Forecasting: Set the values for the following parameters:
 - Causal Aggregation Profile
 - Causal Calculation Intersection
 - Causal Calculation Intersection Periodicity
 - e. Commit your changes to the master database by selecting **Commit Now** from the File menu.
2. Build a [Forecast Maintenance Workbook](#).
 - a. Set the [Forecast Method Override](#) to **Causal** for any items/locations at the desired levels that use Promotional Forecasting.
 - b. Commit your changes to the master database by selecting **Commit Now** from the File menu.
 3. Build a [Promotion Planner Workbook](#).
 - a. Set Promo Effects for items and locations historically in that the selected promotions are active.
 - b. Commit your changes to the master database by selecting **Commit Now** from the File menu.
 4. Select **Run Batch** and generate a Forecast.

Viewing a Forecast That Includes Promotion Effects

Perform this procedure to view a forecast that includes promotion effects.

1. Build a [Forecast Approval Workbook](#) and include System Baseline in your task.
2. In the [Final Level View](#), review the System Baseline and the System Forecast. The System Baseline is predicted demand given no causal effects. The System Forecast is the sum of the System Baseline and the Promotional Peak calculated during the forecast generation process based on the causal data and settings.

Viewing and Editing Promotion System-Calculated Effects

Perform this procedure to view and edit promotion system-calculated effects.

1. Build a [Promotion Maintenance Workbook](#).
 - a. In the [Promotion Parameters View](#) PromoEffects view, review the [System Calculated Effect](#).
 - b. If the [System Calculated Effect](#) is to be modeled using a method other than Automatic, edit the [Promo Effect Type](#).
 - c. If you choose to adjust the system-calculated effect, adjustments can be made to the [System Effect Override](#). You must also set the [Promo Effect Type](#) to **Override All** or **Override Future Only**.
2. Select **Run Batch** and generate a Forecast.

Promotion Simulation (What-if?) and Analysis

Follow this procedure to perform analysis on past promotions and simulate the effects of historic or future promotions:

1. Build a [Promotion Effectiveness Workbook](#).
2. In the [Promotion Parameters View](#), review the [System Calculated Effect](#) and edit the [System Effect Override](#) then select **Calculate** to recalculate the Future Forecast.

Note: Changes to the [Promotion Effectiveness Workbook](#) cannot be committed; however, this task is useful for:

- Reporting on the performance of past promotions.
 - Simulating the effects of future promotions to support more accurate promotional modeling.
-
-

Cross Promotional Effects Module (CPEM)

This chapter provides an introduction to Cross Promotional Effects Module (CPEM). CPEM is a data mining solution that determines promotional Cannibalization, or Halo relationships, or both between items or groups of items.

Refer to [Appendix B, "Appendix: RDF Workbooks \(Fusion Client\)"](#) for a complete list of RDF workbooks, and their respective tabs, views, and measures.

About Promotional Effects

We refer to Cannibalization effects when a promoted item causes a drop in sales of a similar item. An example of this behavior is when a brand of mustard is promoted, it likely experiences an increase in sales and other mustard brands' demand decreases. After the promotional period the demand pattern of the entire mustard category returns to pre-promotion levels.

Conversely, we refer to Halo effects when a promoted item, or group of items causes an increase in sales of a complementary item or items. A typical case is when hot dogs are promoted, hot dog buns are likely to see a peak in sales as well.

Cannibalization is expected to happen at low levels in the merchandise hierarchy, like style or subclass. However, Halo is detected at much higher levels, like category or department.

CPEM is a module within RDF, but it is implemented separately of a RDF forecasting environment. The nature of the problem that mining for cross promotional relationships poses, that is, search for relationships among items in different areas of business, is not compatible with the master/local domain structure of a typical RDF environment. CPEM is implemented in a single domain, such that mining for Halo effects can span multiple departments.

The inputs to CPEM are sales and promotion information for which there is formal integration with the forecasting solution.

The output of CPEM is the promotional elasticity among items, which is location specific. The intersection of the output is item/item/location.

Currently, RDF Causal is predicting promotional sales based on item Self Promo Lifts. Basically the promotion of an item is considered to affect only sales of that item. With information about cross promotional elasticities, the promotion planners have the complete picture on how a promotion affects the entire business, and can act accordingly.

CPEM Workbooks and Wizards

The CPEM wizards allow you to set up the [Effect Estimation Administration Task](#) and [Effect Estimation Review and Approval Task](#).

CPEM includes these workbooks and views:

Workbooks	Description	Views
Effect Estimation Administration Task	Allows you to set up parameters and thresholds, and their overrides, before the batch run.	Default Parameters View Cannibalization Parameter Override View Halo Parameter Override View Halo Products Mapping View
Effect Estimation Review and Approval Task	Allows you to review the cross promotional elasticities and statistical information related to the analysis. Based on the statistics, you can approve the elasticities or adjust them based on business knowledge.	Cannibalization Effects View Cannibalization Effects Drive View Cannibalization History Data View Halo Effects View Halo Effects Drive View Halo History Data View Cannibalization Statistic Report View Halo Statistic Report View

Effect Estimation Administration Task

The Effect Estimation Administration task allows you to set up parameters and thresholds, and their overrides, before the batch run.

This task has these views:

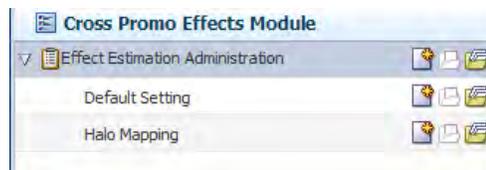
- [Default Parameters View](#)
- [Cannibalization Parameter Override View](#)
- [Halo Parameter Override View](#)
- [Halo Products Mapping View](#)

Effect Estimation Administration Workbook

To build the Effect Estimation Administration workbook, perform these steps:

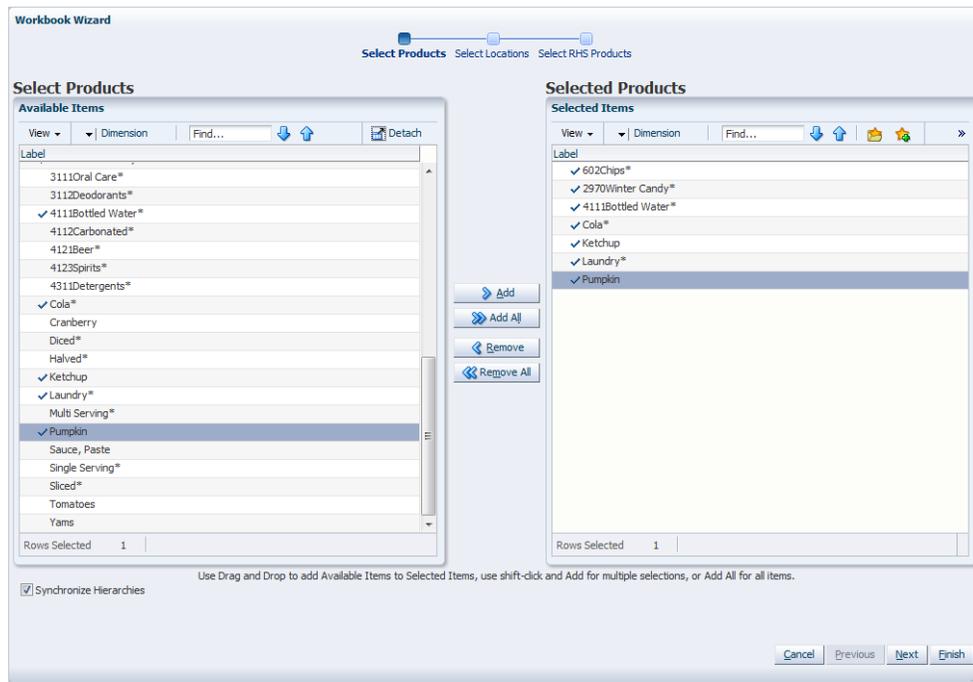
1. Click the New Workbook icon in the Effect Estimation Administration task.

Figure 14–1 *Effect Estimation Administration Task*



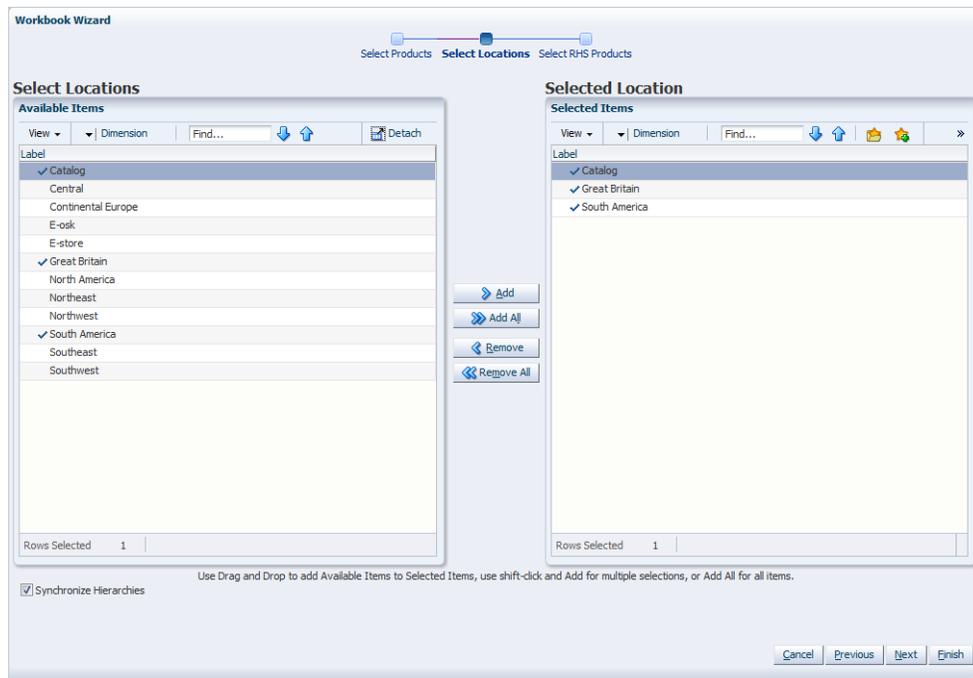
2. The Workbook wizard opens. Select the products for promotion effect estimation and click **Next**.

Figure 14–2 Workbook Wizard: Select Products



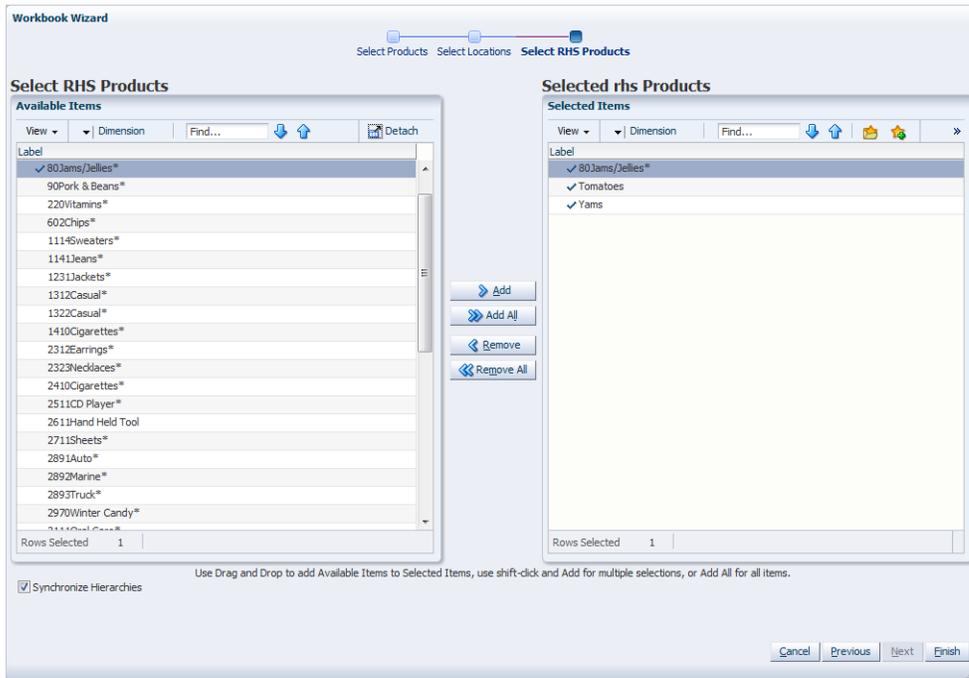
3. Select the locations for promotion effect estimation and click **Next**.

Figure 14–3 Workbook Wizard: Locations



4. Select the right-hand-side (RHS) products for promotion effect estimation and click **Finish**.

Figure 14–4 Workbook Wizard: Select RHS Products

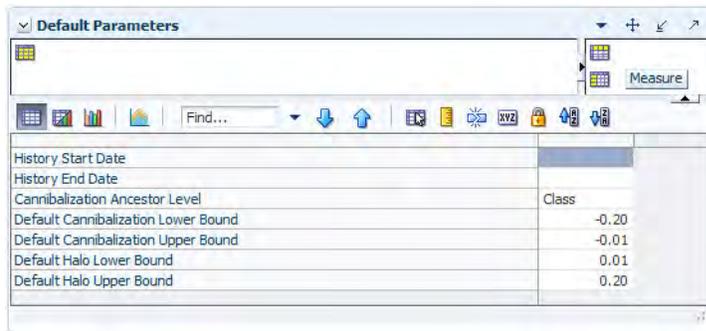


The Effect Estimation Administration workbook is built.

Default Parameters View

The Default Parameters view allows you to set up Default parameters and thresholds, and their overrides, before the batch run.

Figure 14–5 Default Parameters View



Measures: Default Parameters View

The Default Parameters view contains the following measures:

History Start Date

The first calendar date considered in the analysis. All measure data prior to this date is discarded.

History End Date

The last calendar date considered in the analysis. All measure data after this date is discarded.

Cannibalization Ancestor Level

This is the dimension in the merchandise hierarchy where Cannibalization is mined for. Only item pairs within this dimension are evaluated.

Cannibalization Lower Bound

Minimum Cannibalization value threshold. Cannibalization values that are below the threshold are discarded.

Cannibalization Upper Bound

Maximum Cannibalization value threshold. Cannibalization values that are higher than the threshold are discarded (set to zero).

Halo Lower Bound

Minimum Halo value threshold. Halo values that are below the threshold are discarded.

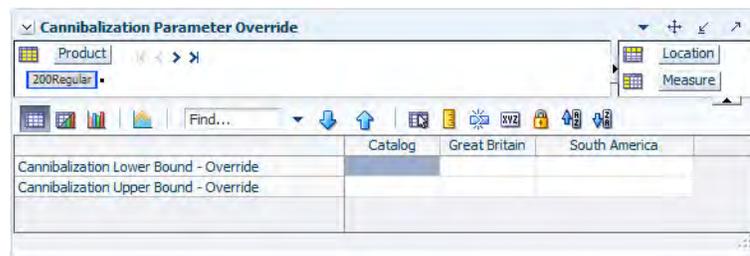
Halo Upper Bound

Maximum Halo value threshold. Halo values that are higher than the threshold are discarded (set to zero).

Cannibalization Parameter Override View

The Cannibalization Parameter Override view allows you to set up Cannibalization parameters and thresholds, and their overrides, before the batch run.

Figure 14–6 *Cannibalization Parameter Override View*

**Measures: Cannibalization Parameter Override View**

The Cannibalization Parameter Override view contains the following measures:

Cannibalization Lower Bound - Override

Override of the minimum Cannibalization value threshold. Cannibalization values that are below the threshold are discarded (set to zero). The override intersection is typically higher than item/store, for example, subclass/region.

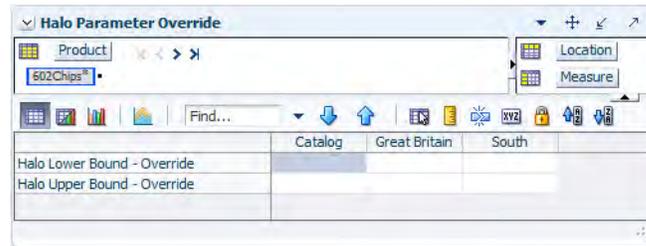
Cannibalization Upper Bound - Override

Override of the maximum Cannibalization value threshold. Cannibalization values that are higher than the threshold are discarded (set to zero). The override intersection is typically higher than item/store, for example, subclass/region.

Halo Parameter Override View

The Halo Parameter Override view allows you to set up Halo parameters and thresholds, and their overrides, before the batch run.

Figure 14–7 Halo Parameter Override View



Measures: Halo Parameter Override View

The Halo Parameter Override view contains the following measures:

Halo Lower Bound - Override

Override of the minimum Halo value threshold. Halo values that are below the threshold are discarded (set to zero). The override intersection is typically higher than item/store, for example, class/region.

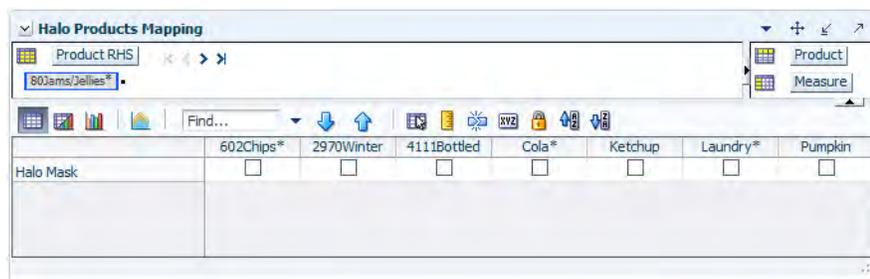
Halo Upper Bound - Override

Override of the maximum Halo value threshold. Halo values that are higher than the threshold are discarded (set to zero). The override intersection is typically higher than item/store, for example, class/region.

Halo Products Mapping View

The Halo Products Mapping view allows you to specify item pairs for which a Halo relationship is mined for. For item pairs that do not have the check box selected, the system will not attempt to look for a Halo relationship.

Figure 14–8 Halo Products Mapping View



Measures: Halo Products Mapping View

The Halo Products Mapping view contains the following measure:

Halo Mask

This measure defines item groupings that are analyzed for Halo relationship. The item groupings are typically bound to the merchandise hierarchy and can be class to class, as in [Figure 14-8](#). This is a writeable measure, and you are expected to use your business knowledge to mark pairs that could have Halo relationships.

Effect Estimation Review and Approval Task

The Effect Estimation Review and Approval workbook provides you the ability to review the cross promotional elasticities and statistical information related to the analysis. Based on the statistics, you can approve the elasticities or adjust them based on your business knowledge.

This task has these views:

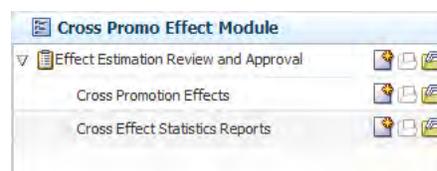
- [Cannibalization Effects View](#)
- [Cannibalization Effects Drive View](#)
- [Cannibalization History Data View](#)
- [Halo Effects View](#)
- [Halo Effects Drive Views](#)
- [Halo History Data View](#)
- [Cannibalization Statistic Report View](#)
- [Halo Statistic Report View](#)

Effect Estimation Review and Approval Workbook

To build the Effect Estimation Review and Approval workbook, perform these steps:

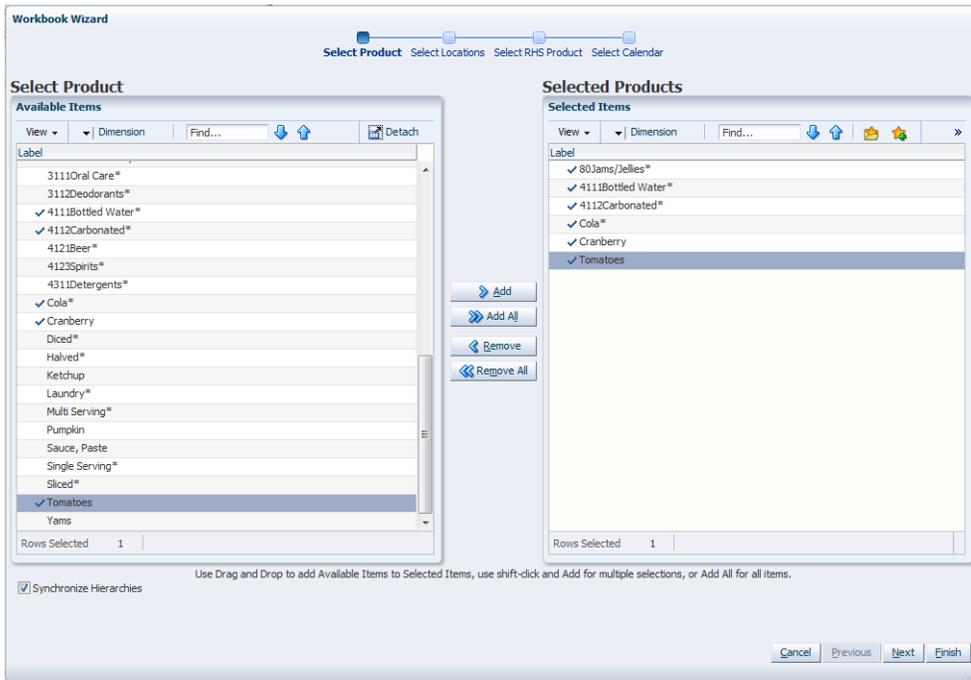
1. Click the New Workbook icon in the Effect Estimation Review and Approval task.

Figure 14-9 *Effect Estimation Review and Approval Task*



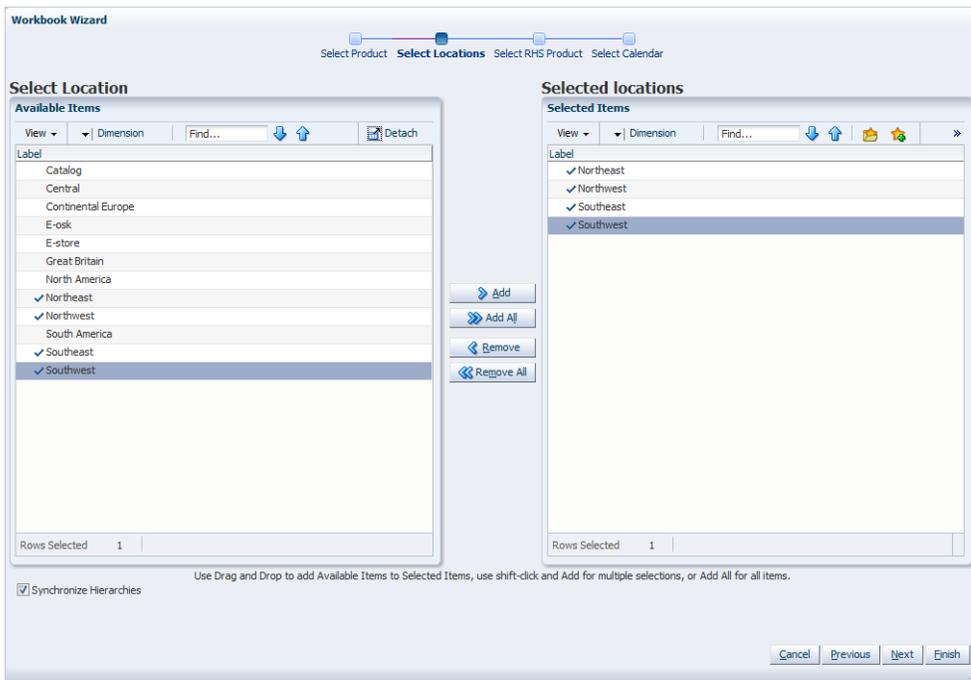
2. The Workbook wizard opens. Select the products for promotion effect estimation review and approval and click **Next**.

Figure 14–10 Workbook Wizard: Select Products



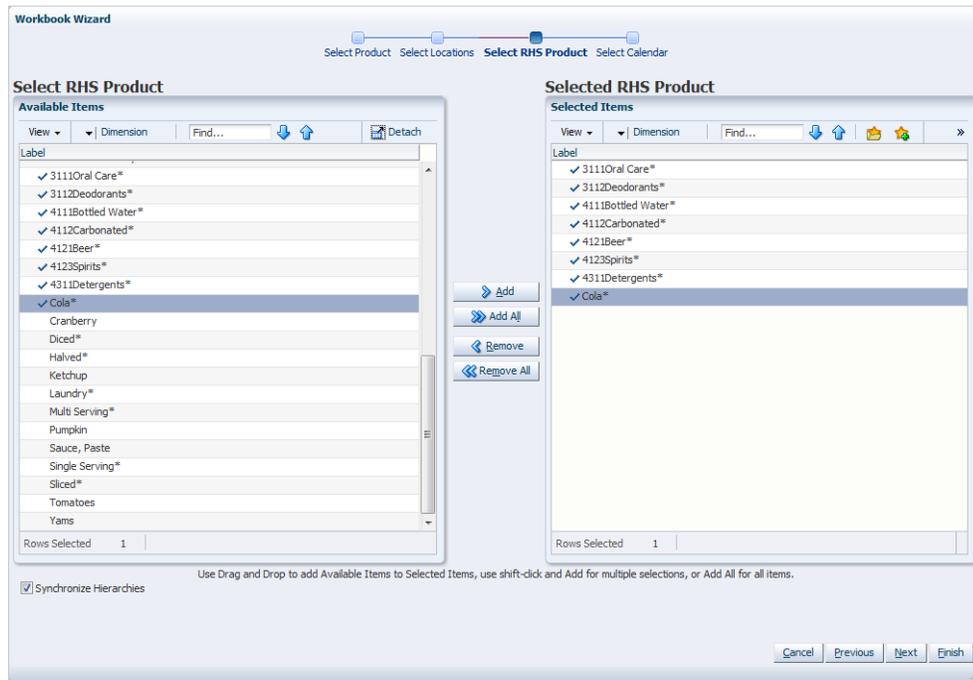
3. Select the locations for promotion effect estimation review and approval and click Next.

Figure 14–11 Workbook Wizard: Locations



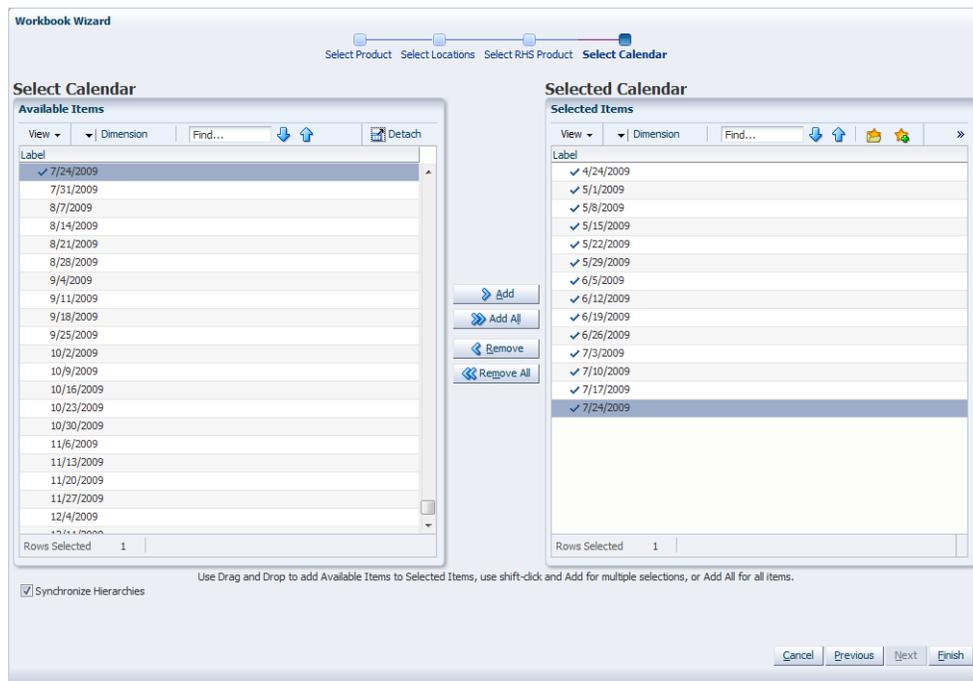
4. Select the right-hand-side (RHS) products for promotion effect estimation review and approval and click Next.

Figure 14–12 Workbook Wizard: Select RHS Products



5. Select the dates for promotion effect estimation review and approval and click **Finish**.

Figure 14–13 Workbook Wizard: Select Dates

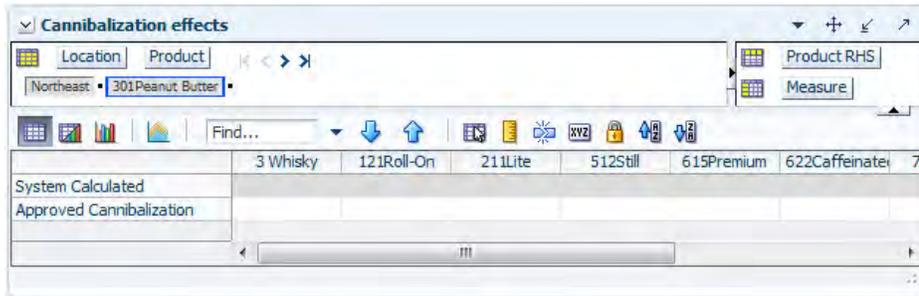


The Effect Estimation Review and Approval workbook is built.

Cannibalization Effects View

The Cannibalization Effects view allows you to review and adjust the Cannibalization cross elasticities.

Figure 14–14 Cannibalization Effects View



Measures: Cannibalization Effects View

The Cannibalization Effects view contains the following measures:

System Calculated Cannibalization Effect

This measure displays the system calculated Cannibalization effects at product/product/location level.

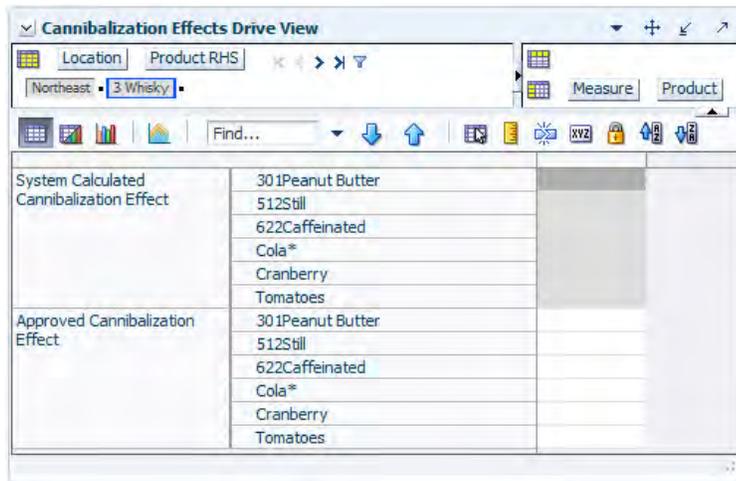
Approved Cannibalization Effect

This measure stores the Cannibalization value that you want to commit. By default it is the same value as the system calculated effect, but you may adjust it.

Cannibalization Effects Drive View

The Cannibalization Effects Drive View contains the same information as the [Cannibalization Effects View](#), but has a different layout and a position query defined.

Figure 14–15 Cannibalization Effects Drive View



Measures: Cannibalization Effects Drive View

The Cannibalization Effects Drive view contains the following measures:

System Calculated Cannibalization Effect

This measure displays the system calculated Cannibalization effects at product/product/location level.

Approved Cannibalization Effect

This measure stores the Cannibalization value that you want to commit. By default it is the same value as the system calculated effect, but you may adjust it.

Cannibalization History Data View

The Cannibalization History Data view allows you to view the historical facts used to estimate the Cannibalization effects.

Figure 14–16 *Cannibalization History Data View*

		4/24/2009	5/1/2009	5/8/2009	5/15/2009	5/22/2009	5/29/2009
Weekly Sales Baseline - Cannibalization	301Peanut Butter						
	512Still						
	622Caffeinated						
	Cola*						
	Cranberry						
	Tomatoes						
Weekly Sales history - Cannibalization	301Peanut Butter						
	512Still						
	622Caffeinated						
	Cola*						
	Cranberry						
	Tomatoes						
Weekly Normalized Price - Cannibalization	301Peanut Butter						
	512Still						
	622Caffeinated						
	Cola*						
	Cranberry						
	Tomatoes						
Weekly Promotion Variable - Cannibalization	301Peanut Butter						
	512Still						
	622Caffeinated						
	Cola*						
	Cranberry						
	Tomatoes						

Measures: Cannibalization History Data View

The Cannibalization History Data view contains the following measures:

Weekly Sales Baseline - Cannibalization

This measure represents the sales history cleansed of Promo Self Effects. It is typically at an intersection higher than the final level intersections in RDF, for example, subclass/region/week.

Weekly Sales History - Cannibalization

This measure represents the total sales history. It is typically at an intersection higher than the final level intersections in RDF, for example, subclass/region/week.

Weekly Normalized Price - Cannibalization

This measure stores the normalized price. It is typically at an intersection higher than the final level intersections in RDF, for example, subclass/region/week. There is no guarantee that all items in a subclass (or region) have the same price, hence the need to normalize them.

Weekly Promotion Variable - Cannibalization

This measure stores the information if any promotion was active for a given time period. It is typically at an intersection higher than the final level intersections in RDF, for example, subclass/region/week.

Halo Effects View

The Halo Effects view allows you to review and adjust the Halo cross elasticities.

Figure 14–17 Halo Effects View



Measures: Halo Effects View

The Halo Effects view contains the following measures:

System Calculated Halo Effect

This measure displays the system calculated Halo effects at product/product/location level.

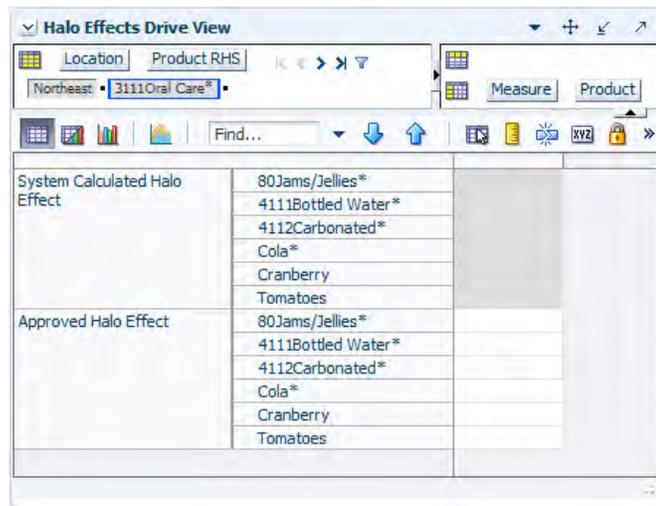
Approved Halo Effect

This measure stores the Halo value that you want to commit. By default it is the same value as the system calculated effect, but you may adjust it.

Halo Effects Drive View

The Halo Effects Drive View contains the same information as the [Halo Effects View](#), but has a different layout and a position query defined.

Figure 14–18 Halo Effects Drive View



Measures: Halo Effects Drive View

The Halo Effects Drive View contains the following measures:

System Calculated Halo Effect

This measure displays the system calculated Halo effects at product/product/location level.

Approved Halo Effect

This measure stores the Halo value that you want to commit. By default it is the same value as the system calculated effect, but you may adjust it.

Halo History Data View

The Halo History Data view allows you to view the historical facts used to estimate the Halo effects.

Figure 14–19 Halo History Data View

		4/24/2009	5/1/2009	5/8/2009	5/15/2009	5/22/2009
Weekly Sales Baseline - Halo	80Jams/Jellies*					
	4111Bottled Water*					
	4112Carbonated*					
	Cola*					
	Cranberry					
	Tomatoes					
Weekly Sales history - Halo	80Jams/Jellies*					
	4111Bottled Water*					
	4112Carbonated*					
	Cola*					
	Cranberry					
	Tomatoes					
Weekly Normalized Price - Halo	80Jams/Jellies*					
	4111Bottled Water*					
	4112Carbonated*					
	Cola*					
	Cranberry					
	Tomatoes					
Weekly Promotion Variable - Halo	80Jams/Jellies*					
	4111Bottled Water*					
	4112Carbonated*					
	Cola*					
	Cranberry					
	Tomatoes					

Measures: Halo History Data View

The Halo History Data view contains the following measures:

Weekly Sales Baseline - Halo

This measure represents the sales history cleansed of Promo Self Effects. It is typically at an intersection higher than the final level intersections in RDF, for example, class/region/week.

Weekly Sales History - Halo

This measure represents the total sales history. It is typically at an intersection higher than the final level intersections in RDF, for example, class/region/week.

Weekly Normalized Price - Halo

This measure stores the normalized price. It is typically at an intersection higher than the final level intersections in RDF, for example, class/region/week. There is no guarantee that all items in a subclass (or region) have the same price, hence the need to normalize them.

Weekly Promotion Variable - Halo

This measure stores the information if any promotion was active for a given time period. It is typically at an intersection higher than the final level intersections in RDF, for example, class/region/week.

Cannibalization Statistic Report View

The Cannibalization Statistic Report view allows you to view Cannibalization statistics on the cross effects.

Figure 14–20 *Cannibalization Statistic Report View*

	301Peanut	512Still	622Caffeinate	Cola*	Cranberry	Tomatoes
Number of Effects Found						
Minimum Effect Found						
Maximum Effect Found						

Measures: Cannibalization Statistic Report View

The Cannibalization Statistic Report view contains the following measures:

Number of Effects Found

This measure displays the number of Cannibalization cross effects found for a given intersection. A typical example for an intersection would be subclass/region.

Minimum Effect Found

This measure displays the minimum value of the Cannibalization effects found for a given intersection. A typical example for an intersection would be subclass/region.

Maximum Effect Found

This measure displays the maximum value of the Cannibalization effects found for a given intersection. A typical example for an intersection would be subclass/region.

Halo Statistic Report View

The Halo Statistic Report view allows you to view Halo statistics on the cross effects.

Figure 14–21 *Halo Statistic Report View*

	80Jams/Jellies	4111Bottled	4112Carbonat	Cola*	Cranberry	Tomatoes
Number of Effects Found						
Minimum Effect Found						
Maximum Effect Found						

Measures: Halo Statistic Report View

The Halo Statistic Report view contains the following measures:

Number of Effects Found

This measure displays the number of Halo cross effects found for a given intersection. A typical example for an intersection would be class/region.

Minimum Effect Found

This measure displays the minimum value of the Halo effects found for a given intersection. A typical example for an intersection would be class/region.

Maximum Effect Found

This measure displays the maximum value of the Halo effects found for a given intersection. A typical example for an intersection would be class/region.

Curve is an optional automated predictive solution that can generate ratio arrays from historical data at user-specified intersections. The profiles generated by Curve can be used for various purposes; for example, to convert the organization-level assortment plans into base-level weekly sales forecasts and to generate seasonal forecasts, daily forecasts, or new product forecasts using lifecycle profiles.

Curve meets the need of operational systems (such as Oracle Retail Demand Forecasting (RDF) and Oracle Retail Merchandising System (RMS)) to have sales unit predictions at a more detailed level than those provided by planning programs. The planning process attempts to establish the correct balance between different products in order to maximize sales opportunities in the available sales space. The planning process is supported by the generation of an assortment plan, which provides details of your anticipated sales volumes and stock requirements at aggregated levels. However, operational systems like RDF require data to be at the lowest level of execution (that is, item/store/week or item/store/day) because these systems are responsible for ensuring that the right quantity of each product is in the right store at the right time.

In the most basic sense, a profile represents the ratio of an aggregate dimension to the dimension for execution. For example, you may have a forecast generated at the item/store/week level, but for execution purposes the data must be spread down to the item/store/day level. It is the point of aggregation (source-level) and the desired destination intersection (final profile) that are the unique identifiers of each profile. Using this example, the point of aggregation of the data (where the data equals 100%) is item/store/week, and the desired destination intersection (where all data ratios sum to 100%) is item/store/day.

There are several parameters within RDF that may take a Curve-generated profile as an input. These are: Causal Aggregation Profile, Causal Spread Profile, Seasonal Profile, and Spreading Profile. The most common input from Curve that is required by RDF is the Spreading Profile. This profile can either be manually generated and approved by you or dynamically generated as part of the RDF batch forecast process. For more information on the different profile parameters in RDF, refer to the section, "[Measures: Causal Parameters View - Advanced Settings](#)" on page 5-11.

Dynamic Profiles

Dynamic Profiles are generated as part of the RDF forecast batch process to produce ratios used to spread source-level forecasts to a final-forecast level. To specify a profile to be generated in such a manner, the profile name (the level number of the profile, 01 for example) is entered in Spreading Profile field in RDF. When the batch forecast is executed, RDF produces an interim forecast at the final-forecast level. This interim forecast becomes the Profile Data Source in Curve and is used to generate the final

profile. The final profile is then passed back to RDF to determine the final forecast results. For information on profile configuration for integration with RDF, refer to both the *Oracle Retail Curve Configuration Guide* and *Oracle Retail Demand Forecasting Configuration Guide*.

Curve Tasks and Wizards

The Curve tasks and wizards allow you to manage the promotion events used in the system's promotional forecasting processes and view/edit the system's analysis of the effects of these events on demand.

The Curve tasks include:

Task Workbooks	Description
Profile Administration Task	Allows you to set default parameters for profile generation, which is the first step in profile generation.
Profile Approval Task	Allows you to view, analyze, revise, and approve generated profiles at historic levels.
Profile Maintenance Task	Allows you to set values for any subset of positions in Profile Administration.
Run Batch Profile Task	Allows you to generate batch profiles for all hierarchy positions set in the Profile Intersection.

Profile Administration Task

The Profile Administration Task allows you to set default parameters for profile generation, which is the first step in profile generation. These parameters are typically set during system implementation and are configured based on your business practices and needs. This configuration can be updated if you need to change certain parameters over time. However, it is not practical to change the configuration on a regular basis. The Profile Administration Task gives you the flexibility to change profiling parameters as the need arises to improve both forecasting accuracy and computational efficiency.

Profile Administration Workbook

To build the Profile Administration workbook, perform these steps:

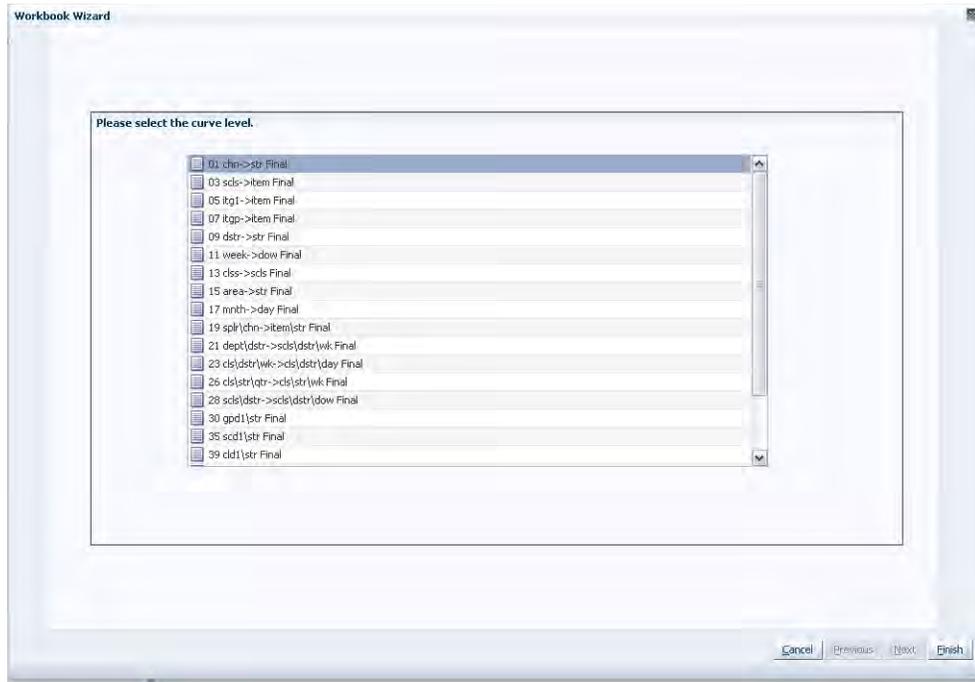
1. Click the New Workbook icon in the Profile Administration task.

Figure 15–1 Profile Administration Task



2. The Workbook wizard opens. Select the Curve level and click **Finish**.

Figure 15–2 Workbook Wizard: Curve Level

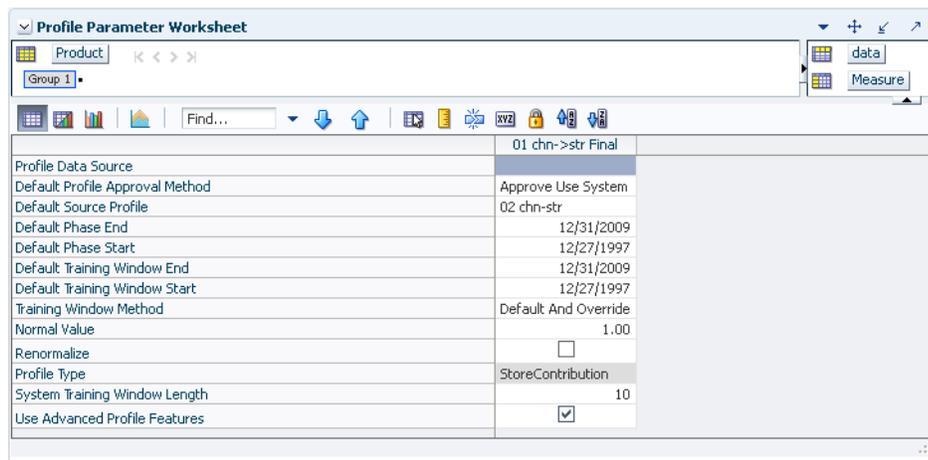


The Profile Administration workbook is built.

Profile Parameter View

The Profile Parameter View allows you to specify default values for parameters affecting profile generation. Figure 15–3 provides an example of the Profile Parameters view in a Master Domain with three partitions/Local Domains, partitioned on Group.

Figure 15–3 Profile Parameter View



Measures: Profile Parameter View

The following is a description of the measures that are contained in the Profile Parameter View:

Default Phase End

The Default Phase End defines the end date of the period in which profile results will be applied. Click the pop-up calendar to change the value in this field. If phase definitions are unavailable, the default phase end date will be used. When calculating time profiles, default dates are used for intermediate computations. For computational efficiency, use the most common phase definition as the default value.

Default Phase Start

The Default Phase Start defines the first date of the period in which profile results will be applied. Click the pop-up calendar to change the value in this field. If phase definitions are unavailable, the default phase start date will be used. When calculating time profiles, default dates are used for intermediate computations. For computational efficiency, use the most common phase definition as the default value.

Default Profile Approval Method

The Default Profile Approval Method displays the primary approval policy to be used for the profile. Select *Approve Use System* if the profile results are to be automatically approved during the batch profile generation. Select *Do Not Approve* if the profile results are to be manually approved by you. If a profile is being generated dynamically (to support Source Level Forecasting) as part of the RDF batch forecast process, the Default Profile Approval Method should be set to Approve Use System.

Default Source Profile

The Default Source Profile determines the primary source-level that will be used to generate the profile. When only a single source is configured for a profile, all profile results will be calculated using the same intersections. When multiple sources are defined for a profile, some profile results will be calculated using different intersections. The Profile Maintenance Task may be used to define exceptions to the Default Source Profile.

Default Training Window End

The Default Training Window End defines the last date in history that will be used to calculate profile results. Click the pop-up calendar to change the value in this field. The default date will be used only if the training window method is set to Defaults and Overrides.

Default Training Window Start

The Default Training Window Start defines the first date in history that will be used to calculate profile results. Click the pop-up calendar to change the value in this field. The default date will only be used if the training window method is set to Defaults and Overrides.

Normal Value

If the profile is aggregated to the Aggregation Intersection, this will be the value in all the cells (or zero outside of the phase window). The desired setting for Normal Value is usually 1.00 (100%); however, there may be instances where it is desired to have the profile normalize to a value different than 1.00. For example, when generating seasonal profiles, such as by Week of Year, the Normal Value needs to be set to 52.

Profile Data Source

The Profile Data Source displays the name of the measure that contains the data to be used to generate the profile.

Note: For Dynamic profiles used by RDF, the profile Data Source can be left blank. Curve will use the Interim Forecast generated by RDF as the Data Source for generating the profile, and the value specified here will be ignored.

Profile Type

The Profile Type is a read-only measure that displays the value for this measure set in the Curve Plug-In. The Profile Type is used to determine the profile algorithm and validation required by the profile level. Profile Types are represented with pre-defined configuration information.

The following Profile Types share the same profile algorithm. The rationale for providing different types that have the same behavior is strictly to remind you of the intent of the profile while using this task:

- **Store Contribution Profile** — The Store Contribution Profile is used to determine the data relationship between stores to aggregate dimensions in the location hierarchy.
- **Daily Profile** — The Daily Profile is used to determine the data relationship between a given day to the week in which it belongs.
- **Hourly Profile** — The Hourly Profile is used to determine the spreading ratios from aggregate dimensions to the hour, hour of day, or hour of week dimensions.
- **Product Profile** — The Product Profile is used to determine the data relationship between any two dimensions along the product hierarchy.
- **Size Profile** — Size profiles are used to determine the data relationship between any size in the product hierarchy and another dimension in the product hierarchy. Note that size can be configured as a dimension along an alternate roll-up along the product hierarchy, such as itemParent-size.
- **User Defined Profile** — The User Defined Profile may be used to support any basic profile configuration.

The following Profile Types have unique behavior:

- **Diff Profile** — Diff Profiles are used to determine spreading ratios from aggregate dimensions in the Product hierarchy to diff dimensions. Used to support the spreading of data in RMS Allocation, Diff Profiles exhibit the same behavior as the previous profile types. However, unique to Diff Profiles is special validation of the relationship between the defined diff dimensions to dimensions along the main branch of the Product hierarchy. Refer to the *Oracle Retail Demand Forecasting Implementation Guide* for more information on validation criteria.
- **Daily Seasonal Profile** — The Daily Seasonal Profile is used to determine the data relationship between a given day of the week to aggregate dimensions in the calendar hierarchy. This profile type uses training window data to compute the profile. The resulting profile is then clipped to fit within the defined phase window.
- **Life Cycle Profile** — The Life Cycle Profile uses data along a user-defined training window and then stretches or shrinks data to fit a user-defined phase window.

Renormalize

Renormalize is a Boolean measure. When set to **true**, it will automatically renormalize the calculated profile result at the corresponding final-level. Normally, the renormalization is not necessary. For example, if you have a source profile at week of

season and its final profile is at day of season, you would need to renormalize the final-level because going from week to day will do replication. At day level, the profile will sum up to greater than one (1) for a season. Since it was a week to day, it will probably sum up to seven. The renormalize will force the final profile to sum to 1.00 (100%).

System Training Window Length

The System Training Window Length is necessary when Use Training Window is set as the Training Window Method for the profile. This field specifies the number of weeks of the most recent data to use as the training window for calculating the profile. The System Training Window Length defaults to 10 weeks.

Training Window Method

The Training Window Method is used to determine the default method that is used to define the training window. The options are:

- **Default and Overrides** — Uses the default dates as set in the Training Window Start Date and Training Window End Date measures.
- **Phase Definitions and Overrides** — Calculates the Training Window Start Date and Training Window End Date based on the Phase Start Date and Phase End Date measures.
- **Use Training Window** — Used with System Training Window Length to specify the number of weeks of the most recent data to use for calculating the profile.

Use Advanced Profile Features

This parameter is selected by default; however, if dynamic profiles are being generated during the RDF batch process, setting this parameter to **false** (unchecked) improves batch performance by skipping unnecessary processes in the creation of the dynamic profiles.

When this is cleared, RDF automatically roll the Training window and Phase windows forward with the forecasts for Dynamic profiles. When selected, you must manually manage the Phase Date windows for which the profile is generated.

Note: When the Use Advanced Profile Features is cleared, then the approved profile is just a direct copy of the system generated profile.

Profile and Source Level Intersection View

The Profile and Source Level Intersection View is a read-only view to the different intersections defined for the Profile and Source level configured in the Curve Plug-in. Refer to the *Oracle Retail Predictive Application Server Solution Extension Configuration Guide* for more information on defining these intersections during configuration of the Curve Solution.

Figure 15–4 Profile and Source Level Intersection View

	01 chn->str	02 chn-str	weekitemchn_
Profile Agg Intersection	weekitemchn_	weekitemchn_	
Profile Approval Intersection	weekitemchn_		
Profile Intersection	weekitemstr_	weekitemstr_	
Stored Intersection	weekitemstr_	weekitemstr_	

Measures: Profile and Source Level Intersection View

The following is a description of the measures that are contained in the Profile and Source Level Intersection View:

Profile Agg Intersection

The Profile Agg Intersection is the intersection where the profile will sum to one (or 100%). If the profile is being used as the Spreading Profile in RDF, this Aggregation Intersection should be the same as the Source Forecast Level.

Profile Approval Intersection

Assigned only at the Final Profile, the Approval Intersection is the intersection where the profile is approved. Approval Intersection should be higher than or equal to the Aggregation Intersection. If the profile is being used as the Spreading Profile in RDF, this Approval Intersection should be the same as the Aggregation Intersection.

Profile Intersection

The Profile Intersection is the intersection where an intermediate profile is calculated. This intermediate profile is then replicated down or aggregated up to the Stored Intersection. If the Store Intersection is the same as the Profile Intersection, the values in intermediate profile are copied to the Stored Intersection. The Profile Intersection must be lower than the Aggregation Intersection.

If the profile is being used as the Spreading Profile in RDF, this Profile Intersection should be the same as the Final Forecast Level.

Stored Intersection

The Stored Intersection is the destination intersection of the profile. The intermediate profile produce at the Profile Intersection is either replicated down to or aggregated up to the Stored Intersection. If the Store Intersection is the same as the Profile Intersection, the values in intermediate profile are copied to the Stored Intersection. The Stored Intersection should not be greater than the Aggregation Intersection. If the profile is being used as the Spreading Profile in RDF, this Stored Intersection should be the same as the Profile Intersection.

Profile Approval Task

The profiles generated at the historic levels must be viewed, analyzed, revised, and approved using the Profile Approval Task. In the approval process, you select the appropriate source-level for each product/location combination. After you make any necessary changes to the profiles and commit the task, the profiles are normalized to preserve the appropriate ratios. At this time, Curve automatically spreads the

source-level profiles to the final-level and combines them. After you commit your changes, you can refresh the data in your task to display the newly generated final-level profiles.

Use the Profile Approval Task to view, analyze, revise, and approve the profiles generated at the historic levels. This task contains three tabs for each view displayed:

- [Final Profile View](#)
- [Source Profile View](#)
- [Profile Approval View](#)

Profile Approval Workbook

To build the Profile Approval workbook, perform these steps:

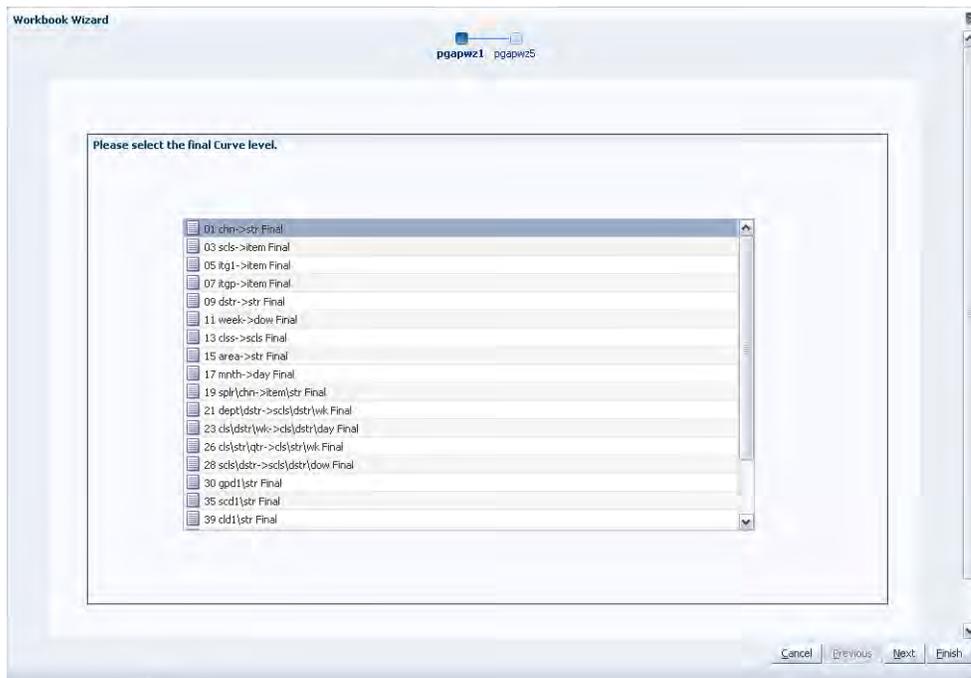
1. Click the New Workbook icon in the Profile Approval task.

Figure 15–5 Profile Approval Task



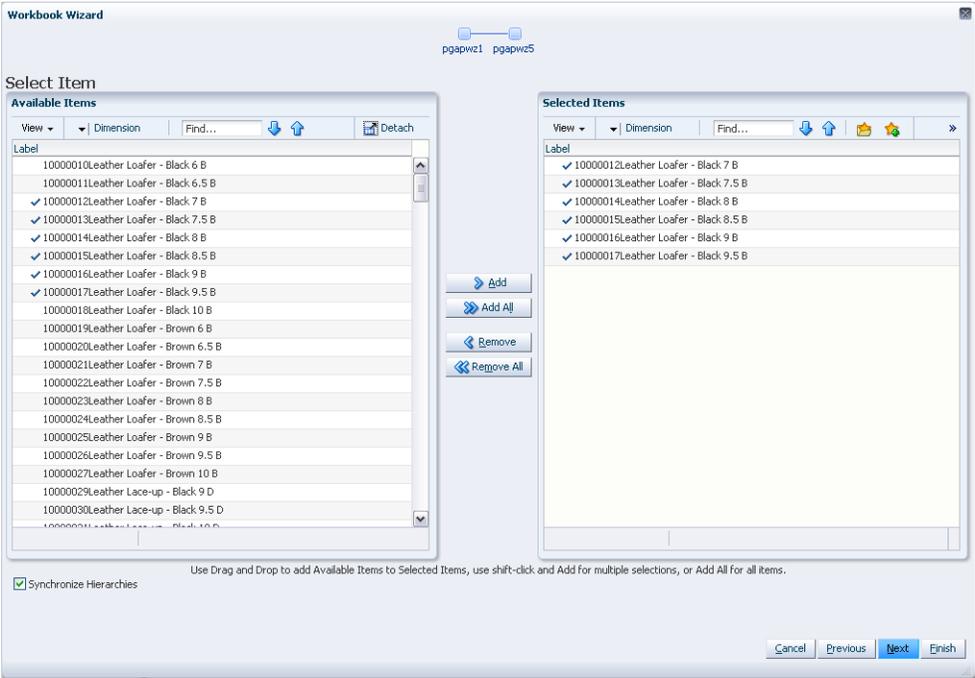
2. The Workbook wizard opens. Select the Curve level and click **Next**.

Figure 15–6 Workbook Wizard: Curve Level



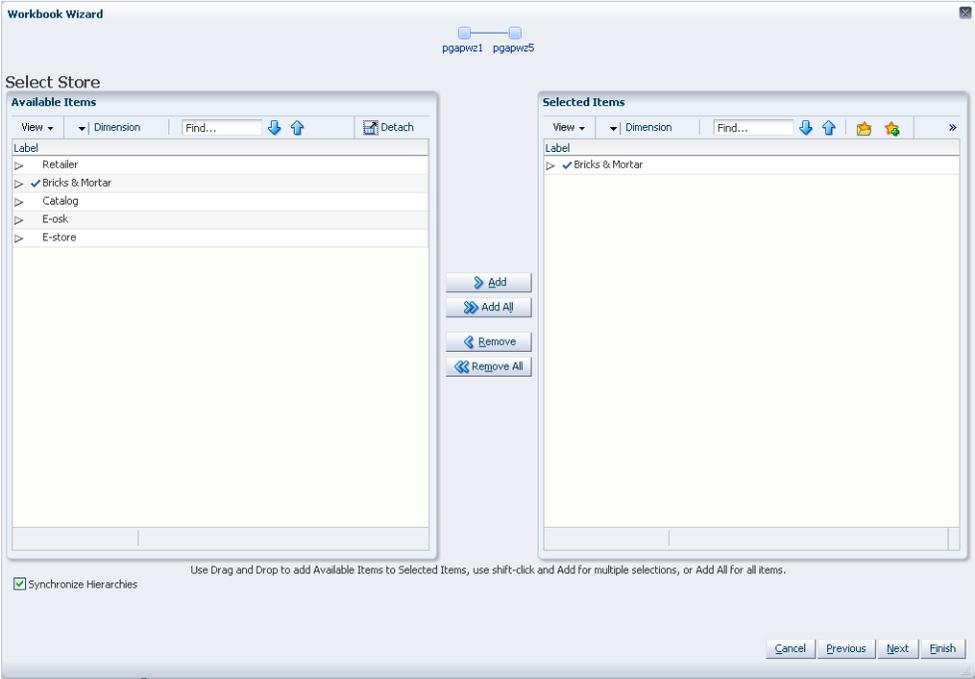
3. Select the items you want to work with and click **Next**.

Figure 15-7 Workbook Wizard: Select Item



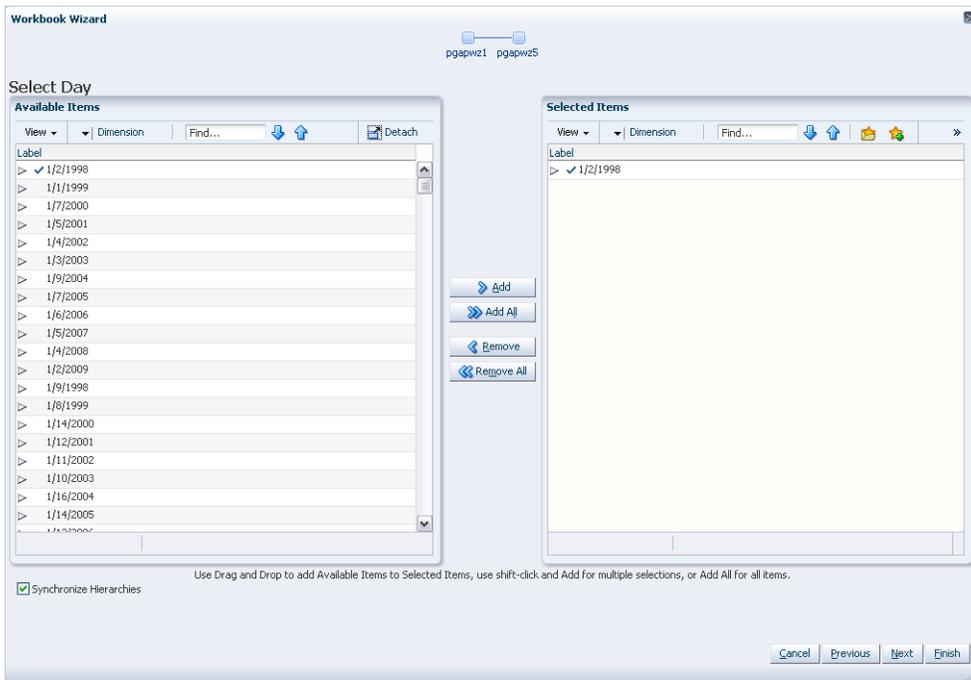
4. Select the stores you want to work with and click Next.

Figure 15-8 Workbook Wizard: Select Stores



5. Select the days you want to work with and click Finish.

Figure 15–9 Workbook Wizard: Select Day

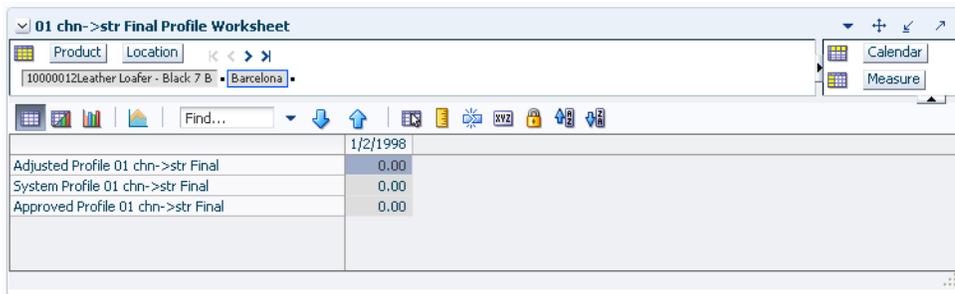


The Profile Approval workbook is built.

Final Profile View

Through this view, you can view the system-calculated final profile and make adjustments to this profile. [Figure 15–10](#) is an example of the Final Profile View.

Figure 15–10 Final Profile View



Measures: Final Profile View

The Final Profile view contains the following measures:

Adjusted Profile

This is the user-adjusted profile. If edits are necessary to the Adjusted Profile, it is first required to lock the Adjust Profile at the Aggregation Intersection. This will prevent the Normal Value from recalculating to a value different than 1.00 (100%) when the adjustments are made. To determine the Aggregation Intersection, view the intersection displayed on the Approval view.

Once adjustments are made to the profile, you must go to the Final Profile Approval View and set *Do Not Approve* to **true** for the adjusted intersection. The Approved Profile measure will update with these changes.

Approved Profile

The Approve Profile displays the approved profile values. If a profile intersection is set to Do Not Approve, no value will be displayed in this field. The system will automatically approve all profile intersections set to Approve Use System. If changes are made to the Adjusted Profile, the values in this measure will update once the Manually Approved flag is set to **true** the Final Profile Approval View.

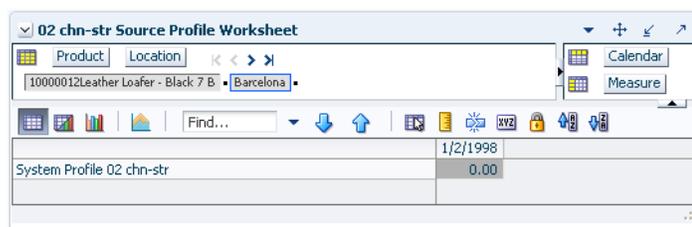
System Profile

A read-only measure that displays the system-generated profiles calculated at the final profile's profile intersection.

Source Profile View

The Source Profile View displays the profiles generated at the source-level for all product/location/calendar combinations selected to appear in the task. This view displays the source system profiles (that is, the profiles that are calculated by the system during the profile generation process).

Figure 15–11 Source Profile View



Measures: Source Profile View

The following measure is contained in the Source Level View:

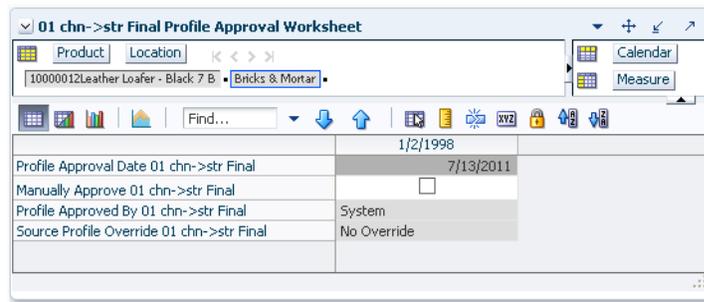
System Profile

A read-only measure that displays the system-generated profiles calculated at the source profile's profile intersection for each product/location combination displayed.

Profile Approval View

The Profile Approval View allows you to review and approve final profiles.

Figure 15–12 Profile Approval View



Measures: Profile Approval View

The Profile Approval View contains the following measures.

Profile Approved By

This measure displays who approved the profile for a given product/location combination. For all profile intersections with an Approval Method set to Approve Use System and no adjustment occurs, the Approved By measure will contain System.

Profile Approval Date

Displays the date on which a profile was approved, whether it is automatically approved by the system or manually approved by you.

Manually Approve

A Boolean measure that must be activated (checked) for all profile intersections that are set to *Do Not Approve* or for intersections in which you make changes to the Adjusted Profile on the Final Profile view.

Source Profile Override

This field displays the source-level that was used to generate the profile. If the Default Source Level was used, this measure will display No Override.

Profile Maintenance Task

After setting default parameters for profile generation in the [Profile Administration Task](#), the next step in profile generation is to select any subset of positions for which the values set in Profile Administration differ from the defaults. This step is necessary in those situations where it is not efficient to use the same parameters for all positions in the hierarchy data.

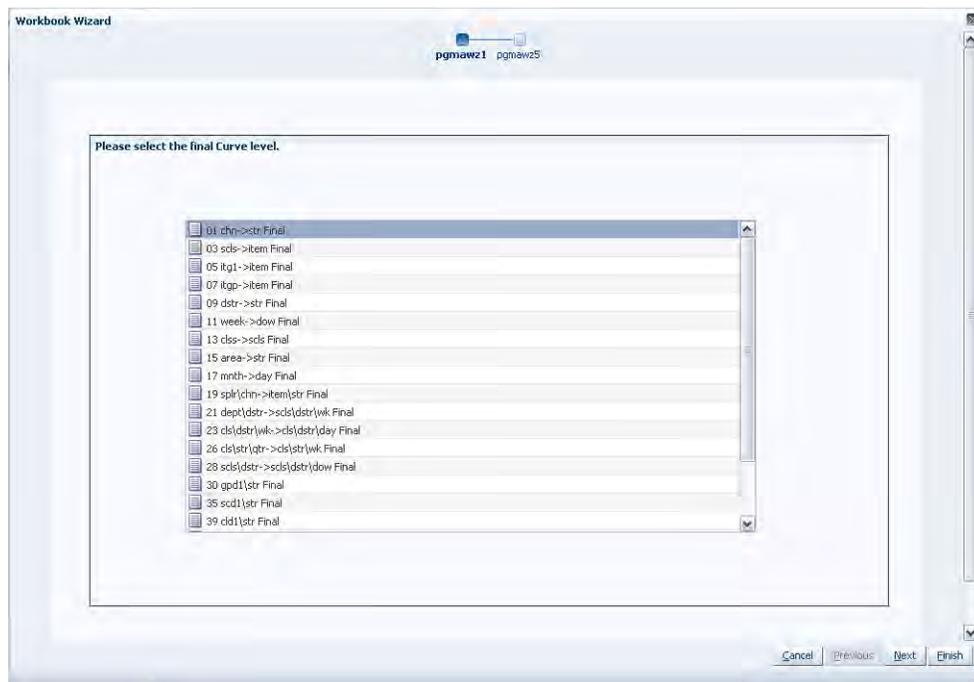
Profile Maintenance Workbook

To build the Profile Maintenance workbook, perform these steps:

1. Click the New Workbook icon in the Profile Maintenance task.

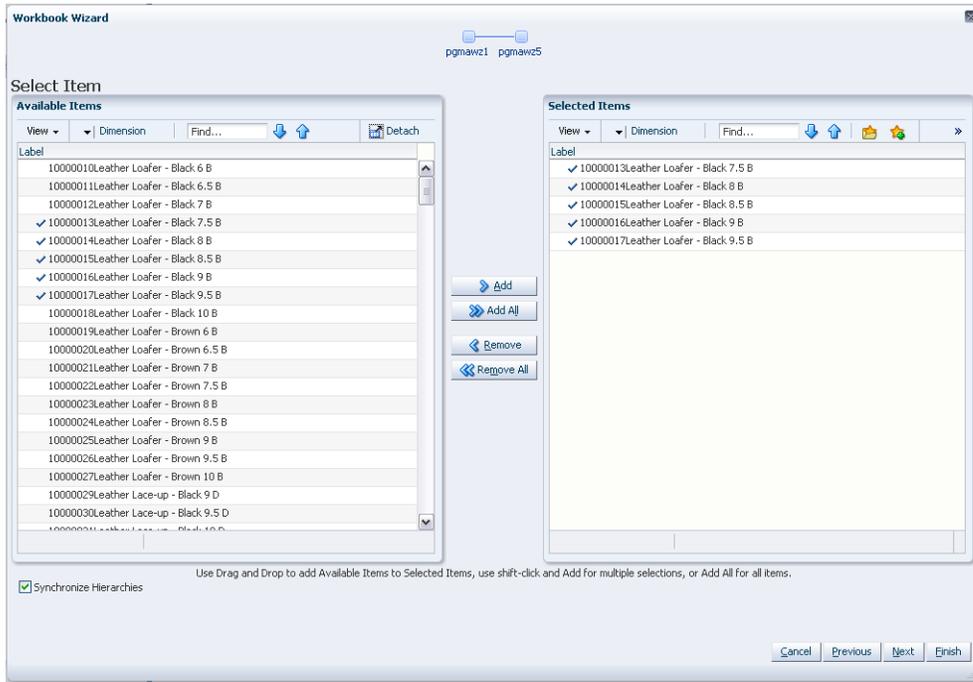
Figure 15–13 Profile Maintenance Task

2. The Workbook wizard opens. Select the Curve level and click **Next**.

Figure 15–14 Workbook Wizard: Curve Level

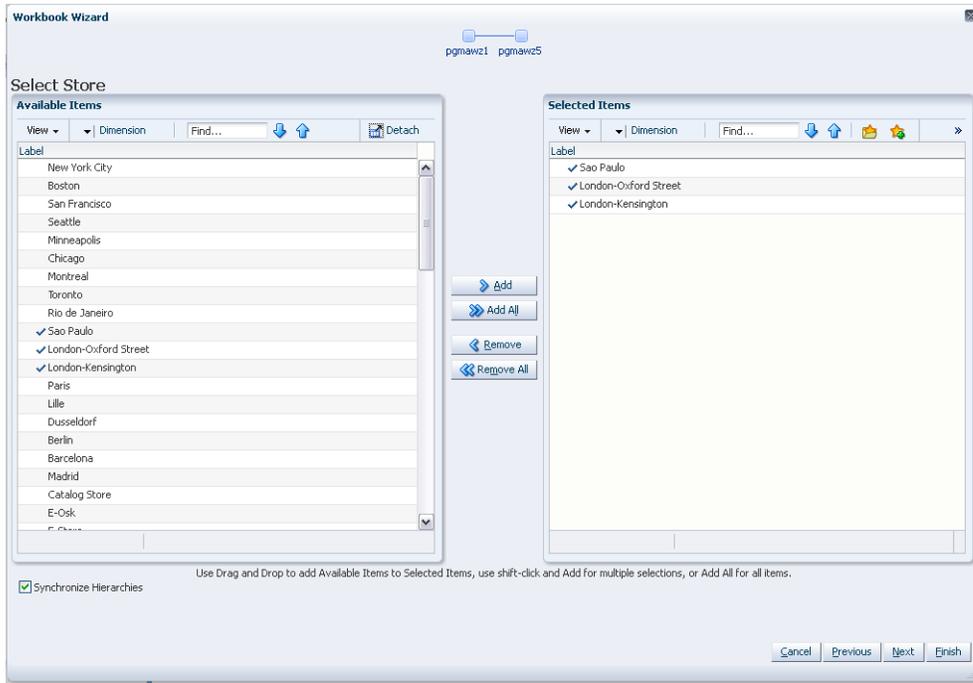
3. Select the products you want to work with and click **Next**.

Figure 15–15 Workbook Wizard: Select Items



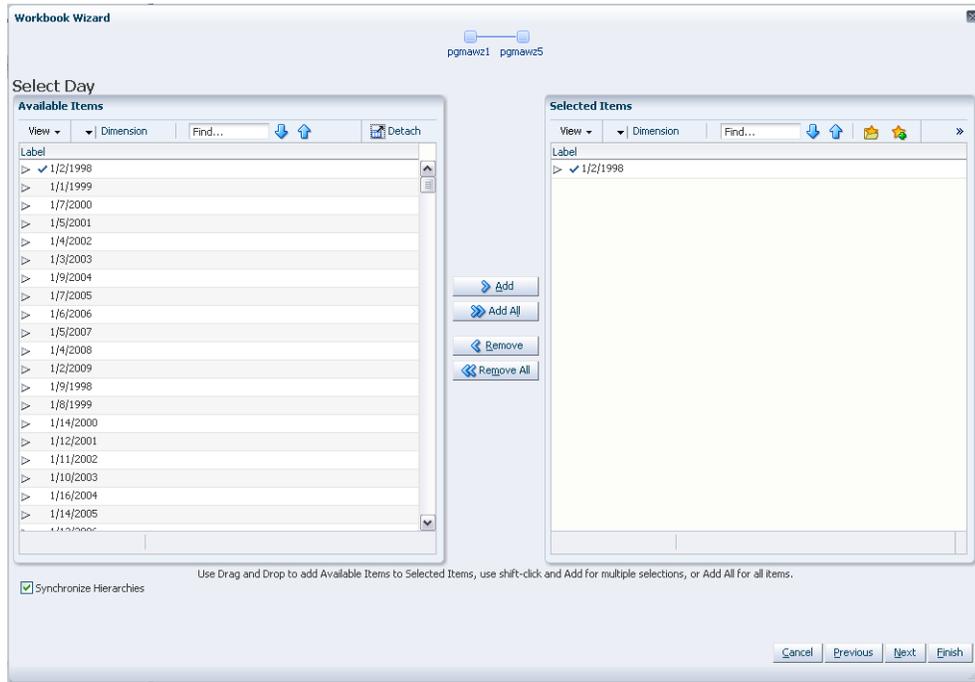
4. Select the locations you want to work with and click **Next**.

Figure 15–16 Workbook Wizard: Select Stores



5. Select the days you want to work with and click **Finish**.

Figure 15–17 Workbook Wizard: Select Day



The Profile Maintenance workbook is built.

Profile Maintenance: Final Approval and Sourcing View

The Final Approval and Sourcing view allows edits to intersections that vary from the default values set in the Profile Administration Task.

Figure 15–18 Final Approval and Sourcing View

	10000013Leather	10000014Leather	10000015Leather	10000016Leather	10000017Leather
Profile Approval Method 01 chn->str Final	No Override				
Source Profile Override 01 chn->str Final	No Override				

Measures: Final Approval and Sourcing View

Profile Approval Method

The Profile Approval Method displays the primary approval policy that is used for the profile. No Override is displayed in this field if the Default Approval Method is used. *Select Approve Use System* if profile results are to be automatically approved during the batch profile generation. Select *Do Not Approve* if profile results are to be manually approved by you. If a profile is being generated dynamically (to support Source Level Forecasting) as part of the RDF batch forecast process, the Default Profile Approval Method should be set to Approve Use System.

Source Profile Override

Make edits to the Source Profile Override if the source-level for an intersection varies from the Default Source Profile. No Override is displayed in this field if the value set in the Default Source Profile is to be used. When only a single source is configured for a profile, all profile results are calculated using the same intersections. Therefore, edits to this parameter are not required. When multiple sources are defined for a profile, some profile results are calculated using different intersections. Select the appropriate source-level for an intersection.

Profile Maintenance: Final Training Window View

Figure 15–19 Final Training Window View

	10000013Leather	10000014Leather	10000015Leather	10000016Leather	10000017Leather
System Training Window Begin 01 chn->str Final					
System Training Window End 01 chn->str Final					

Measures: Final Training Window

The following is a description of the measures contained in the Final Approval and Sourcing View:

System Training Window Begin

Select a date in this field if this date is different from the default value set in the Training Window Start date in Profile Administration. Click the pop-up calendar to change the value in this field.

System Training Window End

Select a date in this field if this date is different from the default value set in the Training Window End date in Profile Administration. Click the pop-up calendar to change the value in this field.

Run Batch Profile Task

The batch profile generation process creates profile results for all hierarchy positions set in the Profile Intersection. Profiles may be run from the back-end of the domain using the *curvebatch* executable or run manually using the Run Batch Profile workbook. For more information on *curvebatch*, refer to the *Oracle Retail Predictive Application Server Administrators Guide*.

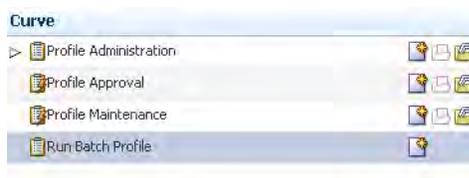
Run Batch Profile Workbook

The Run Batch Profile wizard automatically executes *curvebatch* within the Simple or Local Domain. If **Next** is selected from the last wizard window, the wizard will not advance to the completion message until the profiles have been generated. Depending on the data set, this process may take a several minutes before the system advances to the final window.

To run the Run Batch Profile wizard, perform these steps:

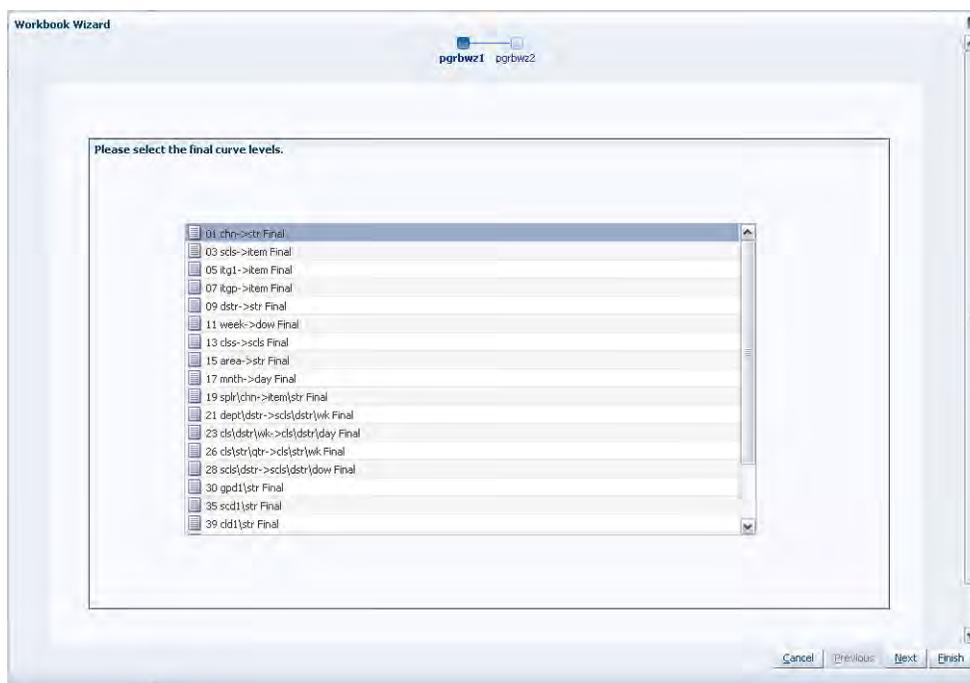
1. Click the New Workbook icon in the Run Batch Profile task.

Figure 15–20 Run Batch Profile Task



2. The Workbook wizard opens. Select the Curve level and click **Next**.

Figure 15–21 Workbook Wizard: Curve Level



3. Verify the Curve level and click **Finish**.

Figure 15–22 Workbook Wizard: Verify Curve level



The Run Batch Profile workbook is built.

Grade is a clustering tool that provides insight into how various parts of a retailer's operations can be grouped together. Typically, a retailer may cluster stores over item sales to create logical groupings of stores based upon sales of particular products. This provides increased visibility to where products are selling, and it allows the retailer to make more accurate decisions in merchandising.

Beyond this traditional use of clusters, Grade is flexible enough to cluster any business measure based on products, locations, time, promotions, customers, or any hierarchy configured in the solution.

Key Grade functionality includes:

- Two methods of creating Grades/Clusters:
 - Breakpoints—the sorting of data points into groups based on user-defined indexes
 - Clustering (or the BaNG) Algorithm—the optimization of data points into clusters based on the user-defined number of clusters
- *Group By* capabilities: support the segmentation of clusters for more detailed and focused cluster generation
- Clustering statistics: provide insight into the relationship of members within a cluster and how all clusters relate to one another
- Cluster What-if: allows user changes to members assigned to clusters and the review of recalculated clustering statistics

Regardless of the method employed to create clusters, Grade is designed to support the decision-making process necessary to create effective and actionable groupings of data. This chapter describes the process to generate grades/clusters and analyze results. All of Grade's functionality exists in the:

- [Breakpoints Administration Workbook](#)
- [Delete Cluster Run Wizard](#)
- [Cluster Review Workbook](#)
- [Generate Breakpoint Grades Wizard](#)
- [Generate Clusters Wizard](#)

Cluster Methods in Grade

Grade supports two cluster methods: Breakpoint and BaNG. The following sections describe these two methods.

Breakpoint

Breakpoint method clusters, or groups, data points based on user-defined thresholds, or breakpoints. Breakpoints are properly defined by entering a decreasing sequence of real numbers that terminate in a value of zero. Breakpoints are entered in the Breakpoints Administration Workbook. Breakpoints are then used to cluster data points based on the index to average of each data point.

Example 16–1 Breakpoint

While clustering stores, if the user has entered the breakpoints 2.0, 1.5, 1.0, 0.5, and 0.0, the system will generate five grades. Stores that sell more than 2.0 times the average sales (over all stores) will be assigned to the highest grade, and stores that sell less than 0.5 times the average will be assigned to the lowest grade.

Index to Average

Index to Average is the average sales of the stores in each grade divided by the average sales of all stores. This value provides a relative indication of how well a grade performed compared to the total store average. A value of 1.00 indicates that average sales in the grade were the same as average sales across the chain. A value of 3.74, for example, indicates that average sales in the grade were 3.74 times the chain average.

Batch Neural Gas Algorithm (BaNG)

The BaNG algorithm automatically generates optimal clusters based on user-specified number of clusters and clustering criteria. The algorithm provides a means for clustering data based on data distributions. For example, while clustering on weekly store sales data, the BaNG algorithm considers the Euclidean distance of the individual Store/week level data points from a cluster center to determine the clusters. This is different from the Breakpoint method, where clustering is performed based on average sales.

The BaNG algorithm iteratively updates cluster centers while considering the distance of each data vector from the cluster centers and its contribution to each cluster center. For every data point, cluster centers are ranked based on their distance from the data point within each iteration.

Additionally, the cluster centers are guided, using a control parameter, to gradually spread from the center of the distribution to their optimal locations.

The BaNG algorithm is a non-trivial extension of the K-means clustering approach. It is usually faster than the K-means and is guaranteed to converge.

BaNG versus Breakpoint

The BaNG algorithm generates statistically optimal clusters based on the number of clusters specified by the user. Breakpoint generates clusters based on user input breakpoints, and the number of clusters generated depends on the breakpoints.

In order to generate store clusters that vary by Dept, users need to specify a Group By option of Dept. Breakpoint will cluster stores based on the ratio of total store sales versus the average total store sales within each Group By intersection. BaNG can consider an additional dimension for generating the clusters. For example, BaNG can cluster based on weekly sales of each stores within the Dept. Here, weekly sales are the coordinates over which the clustering is performed.

Grade in a Global Domain Environment

When implemented in a global domain environment, the following workbooks are available to be accessed from the Master domain:

- [Breakpoints Administration Workbook](#)

Note: The Breakpoint Administration workbook is accessible from the Master domain to allow for the centralized administration of breakpoints.

The remaining Grade workbooks can only be accessed from local domains. These include:

- [Delete Cluster Run Wizard](#)
- [Cluster Review Workbook](#)
- [Generate Breakpoint Grades Wizard](#)
- [Generate Clusters Wizard](#)

Grade Tasks and Wizards

The Grade tasks and wizards allow you to generate grades/clusters and analyze results.

The Grade tasks and wizards include:

Task Workbooks	Description
Breakpoints Administration Task	Allows you to set the index to average for each breakpoint.
Delete Cluster Run Task	Allows you to delete clusters from the system based on Cluster Run Label and generation date (birth date).
Cluster Review Task	Allows you to support Grade/cluster What-if functionality; the re-assignment of members to grades/clusters on the fly and the recalculation of clustering statistics.
Generate Breakpoint Grades Task	Allows you to range the input data and hierarchies that will be used to produce grades based on the Breakpoints method.
Generate Clusters Task	Allows you to range the input data and hierarchies that will be used to produce clusters based on the Clustering (BaNG) method.

Breakpoints Administration Task

The Breakpoints Administration task is used with the breakpoints method of grading. In this task, you set the index to average for each breakpoint. This includes the ability to set multiple breakpoint configurations to allow for Grades to be produced and compared using different breakpoint settings.

Breakpoints Administration Workbook

To build the Breakpoints Administration workbook, perform these steps:

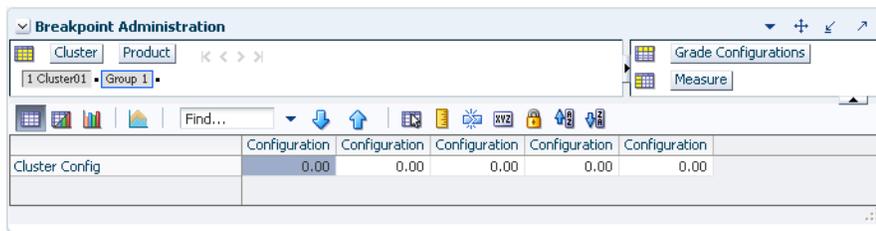
1. Click the New Workbook icon in the Breakpoints Administration task.

Figure 16–1 Breakpoints Administration Task

The [Breakpoint Administration View](#) opens.

Breakpoint Administration View

If the breakpoints method is to be used to generate grades, you must go through the process of setting of the Index to Average for each grade's breakpoint range. [Figure 16–2](#) displays an example of the Breakpoint Administration view:

Figure 16–2 Breakpoint Administration View

On the Breakpoint Administration view, for a Configuration/Cluster intersection, set the Index to Average to be used by the Breakpoints algorithm for sorting data points into grades. The number of configurations available in this view is based on the Maximum Number of Clusters configured in the Grade Plug-In. Only one configuration is required for use with the breakpoints method.

Breakpoints should be set from high Index to Average to low Index to Average, starting with the first Cluster ordered in the list of available clusters.

Note: The Cluster Labels may vary based on the configuration.

Example: If the Index to Average is set for Configuration 01 as:

- 1 Cluster01: 1.25
- 2 Cluster02: 0.75
- 3 Cluster03: 0.50
- 4 Cluster04: 0.25
- 5 Cluster05: 0.00

Using this example, the Breakpoints algorithm groups data based on the following:

Data	Description
1 Cluster01	Sort all data with an Index to Average at or above 1.25 into 1 Cluster01.
2 Cluster02	Sort all data with an Index to Average from 0.75 to 1.24 into 2 Cluster02.
3 Cluster03	Sort all data with an Index to Average from 0.50 to 0.74 into 3 Cluster03.
4 Cluster04	Sort all data with an Index to Average from 0.25 to 0.49 into 4 Cluster04.
5 Cluster05	Sort all data with an Index to Average from 0.00 to 0.24 into 5 Cluster05.
Junk Cluster	Using the previous example, all data points with an Index to Average that is less than 0.00 will be sorted into the Junk Cluster.
No Cluster	Any data points with null values in history (no loaded history) will not be graded. Refer to the Cluster Review Task for more information on Cluster Membership results.

Delete Cluster Run Task

The Delete Cluster Run wizard allows you to delete clusters from the system-based on Cluster Run Label and generation date (birth date).

Delete Cluster Run Wizard

To run the Delete Cluster Run wizard, perform these steps:

1. Click the New Workbook icon in the Delete Cluster Run task.

Figure 16–3 Delete Cluster Run Task



2. This wizard window is a single select pick-list that includes the Cluster Run Label and the date/time stamp (birth date) for all cluster runs currently stored in the system. Select the Cluster Run Label to be deleted.
3. Select **Next** or **Finish**.

If **Next** is selected, Delete Cluster Run Succeeded is displayed once the cluster deletion process is completed.

Cluster Review Task

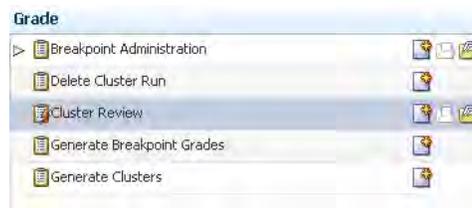
The Cluster Review workbook is a view of grade or cluster results and statistics. This workbook supports Grade/cluster What-if, the re-assignment of members to Grades/clusters on the fly, and the recalculation of clustering statistics. In addition, measures from different cluster runs may be inserted into the workbook to compare results.

Cluster Review Workbook

To run the Cluster Review wizard, perform these steps:

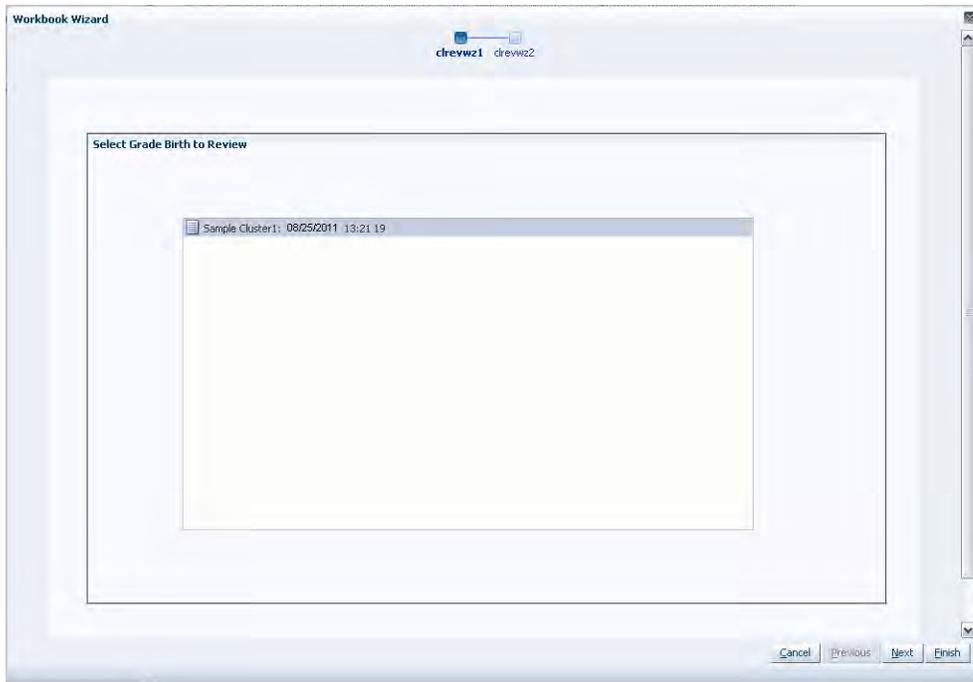
1. Click the New Workbook icon in the Cluster Review task.

Figure 16–4 Cluster Review Task



2. The Workbook wizard opens. Select the Grade Birth you want to review and click **Finish**.

Figure 16–5 Workbook Wizard: Select Grade Birth



The Cluster Review workbook is built.

Cluster Input Summary View

The Cluster Input Summary view allows you to see the settings used to generate the grades/clusters being reviewed in the workbook.

Figure 16–6 Cluster Input Summary View

Property	Value
Number Of Clusters	3
Dimension to Cluster Over	DSTR
Cluster Run Name	Sample
Cluster Method	BANG
Measure to Cluster	DPOS
Dimension to Cluster	WEEK
Cluster By Group Intersection	MNTH

Measures: Cluster Input Summary View

Number of Clusters

Displays the number of clusters used during the cluster generation process.

Dimension to Cluster Over

Displays the Cluster Over dimension selected by the user during the cluster generation process.

The algorithm uses the positions in this dimension as the coordinates when clustering. Internally, data will be aggregated to this level before performing the clustering.

For example, when using the SKU/Store/week level sales data to generate store clusters, the user has the option of specifying a dimension, such as Week, that will be used to aggregate the sales data to. In this case, the algorithm considers Store/Week positions and clusters stores based on their weekly sales.

A value will only be displayed in this field if the BaNG method was used.

Cluster Run Name

Displays the name assigned by the user to the cluster run at the time the clusters were generated.

Measure to Cluster

Displays the measure selected as the clustering Data Source during the cluster generation process.

Dimension to Cluster

Displays the Cluster dimension selected by the user during the cluster generation process.

Cluster By Group Intersection

Displays Group By dimensions if selected during the cluster generation process. For example, while clustering stores by time period, such as quarter, the Cluster By Group is Quarter (Qtr) within the Calendar hierarchy. This means that clusters are generated one quarter at a time.

Generate Breakpoint Grades Task

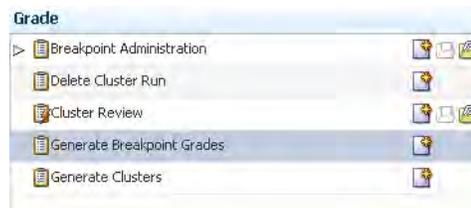
The Generate Breakpoint Grades wizard allows you to range the input data and hierarchies that will be used to produce grades based on the Breakpoints method.

Generate Breakpoint Grades Wizard

To run the Generate Breakpoint Grades wizard, perform these steps:

1. Click the New Workbook icon in the Generate Breakpoint Grades task.

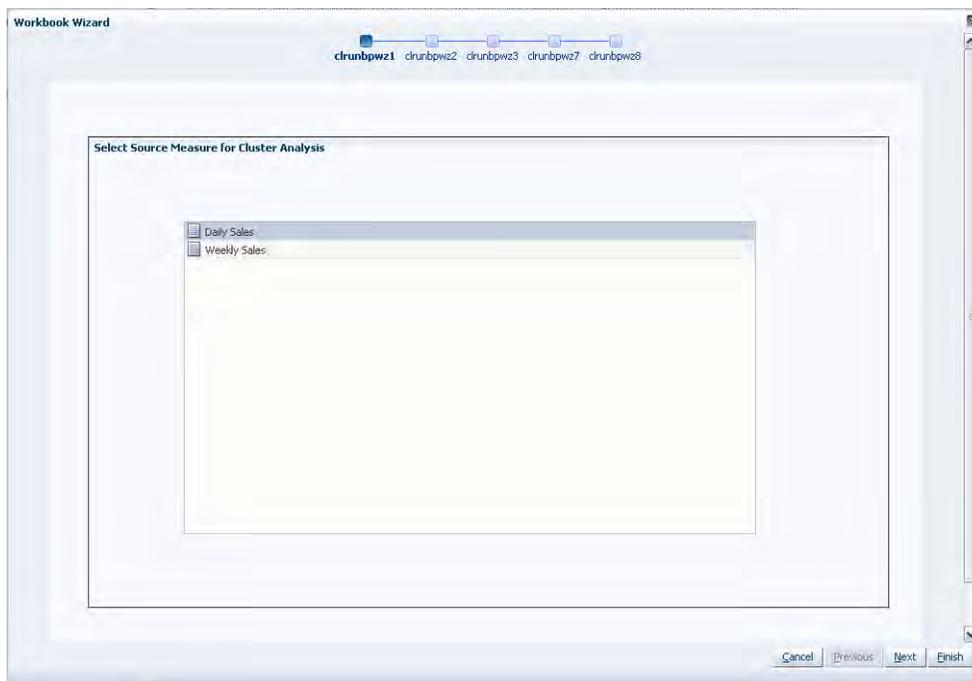
Figure 16–7 Generate Breakpoint Grades Task



2. The Workbook wizard opens. Select the source measure for Cluster Analysis and click **Next**.

The options include all source measures defined at the time of the configuration.

Figure 16–8 Workbook Wizard: Source Measure for Cluster Analysis

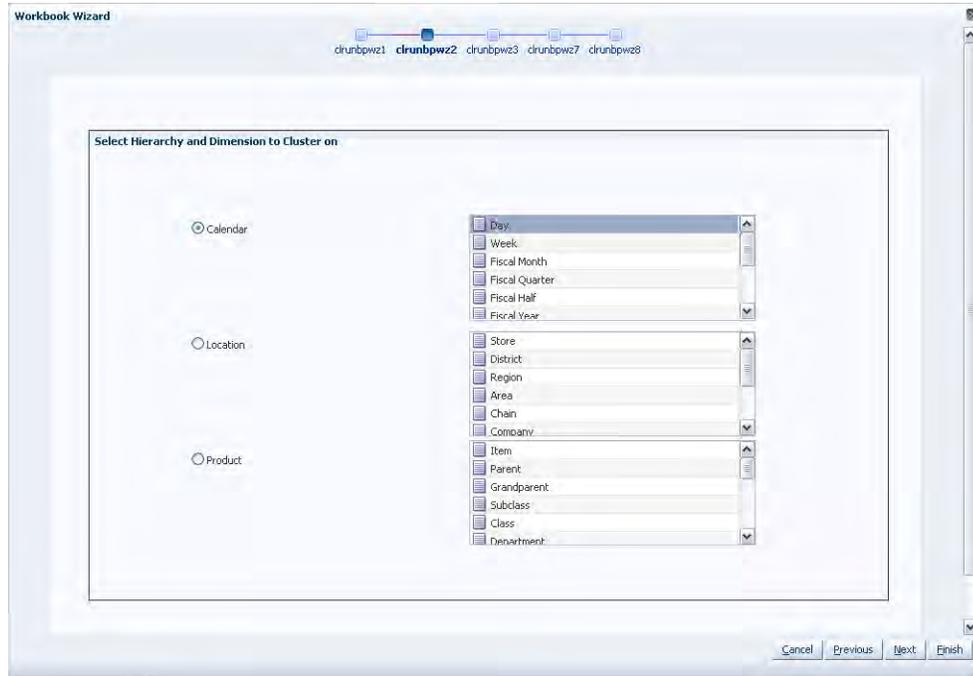


3. Select the Hierarchy and Dimension to Cluster On and click **Next**.

The Cluster On dimension is the dimension that will be clustered. For example, while generating store clusters, store would be the Cluster On dimension. All positions in this dimension are sorted into grades based on the Breakpoints Configuration that will be selected later in this wizard process.

Note: The dimension chosen should be at or higher than the base intersection of the selected Source Measure.

Figure 16–9 Workbook Wizard: Hierarchy and Dimension to Cluster On



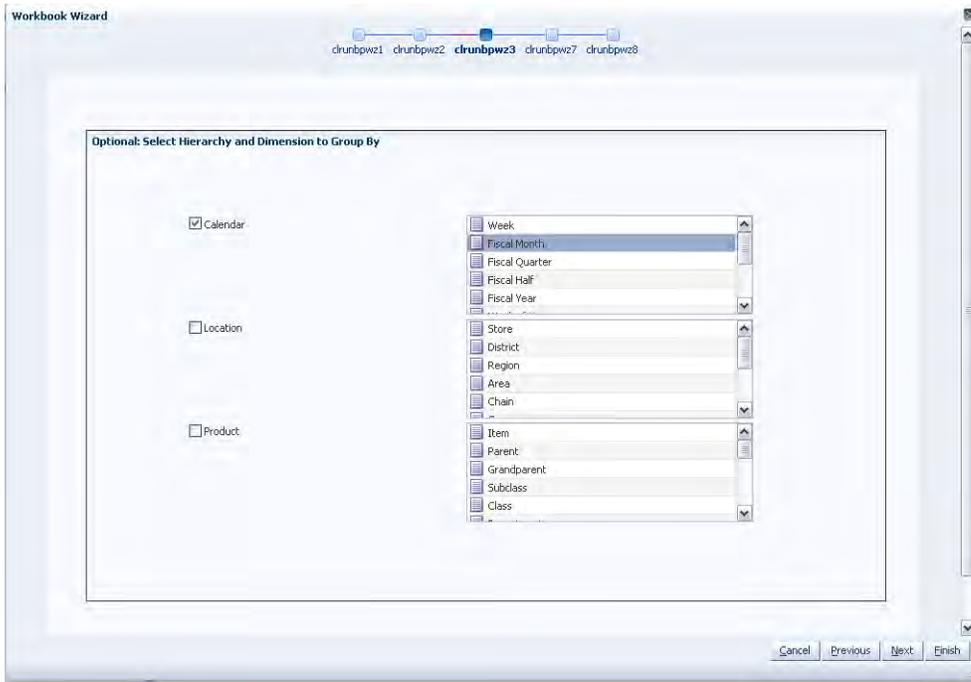
4. Optional Step

Select the Hierarchy and Dimension to Group By and click **Next**.

By Group is an optional setting that is used to partition data along a hierarchy dimension. This functionality allows for grades to be further segmented along multiple hierarchy dimensions. Grades are generated independently for each position within the selected dimensions. The dimensions displayed are equal to or higher than the Cluster On dimension.

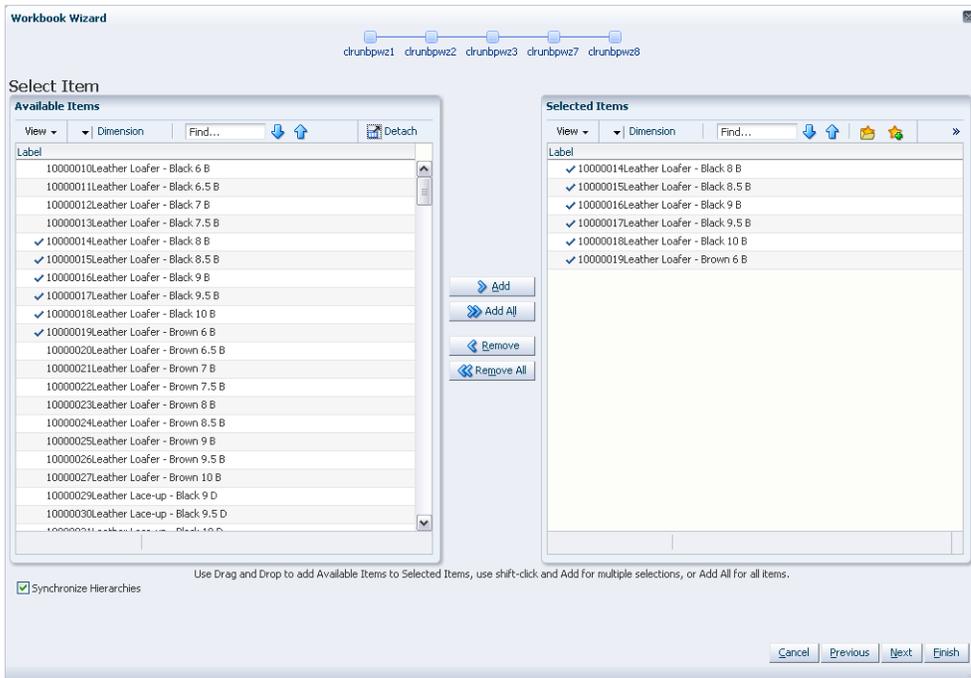
For example, while clustering stores by Dept, that is, cluster stores by a Department at a time, the Cluster By Group is Department.

Figure 16–10 Workbook Wizard: Select Hierarchy and Dimension to Group By



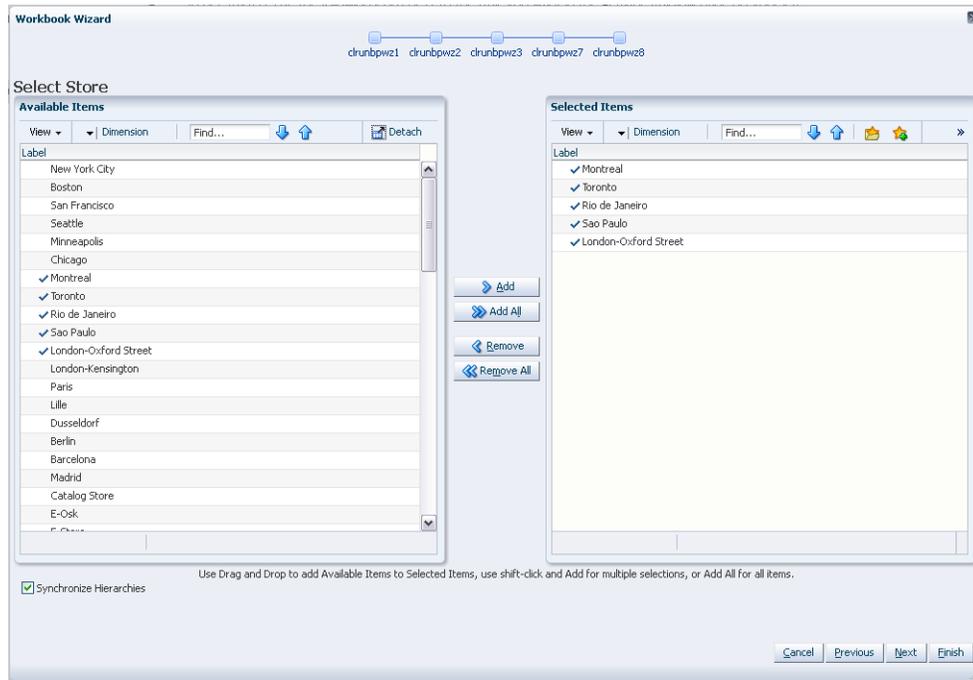
5. Select the products for which you want to generate clusters and click **Next**.

Figure 16–11 Workbook Wizard: Select Items



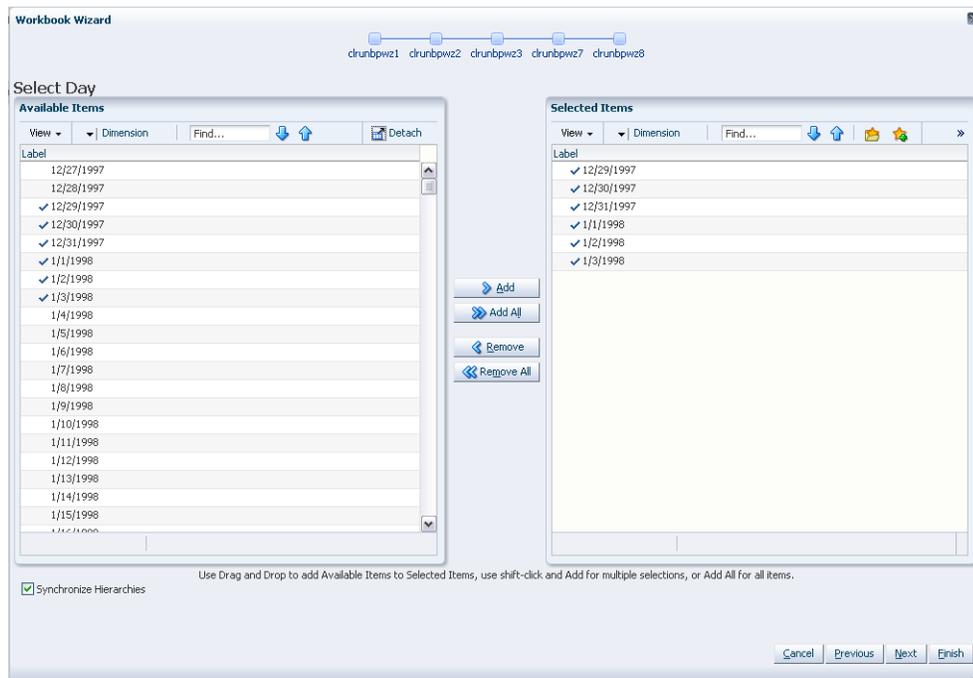
6. Select the locations for which you want to generate clusters and click **Next**.

Figure 16–12 Workbook Wizard: Select Stores



7. Select the time periods that you want to be considered when generating clusters and click Next.

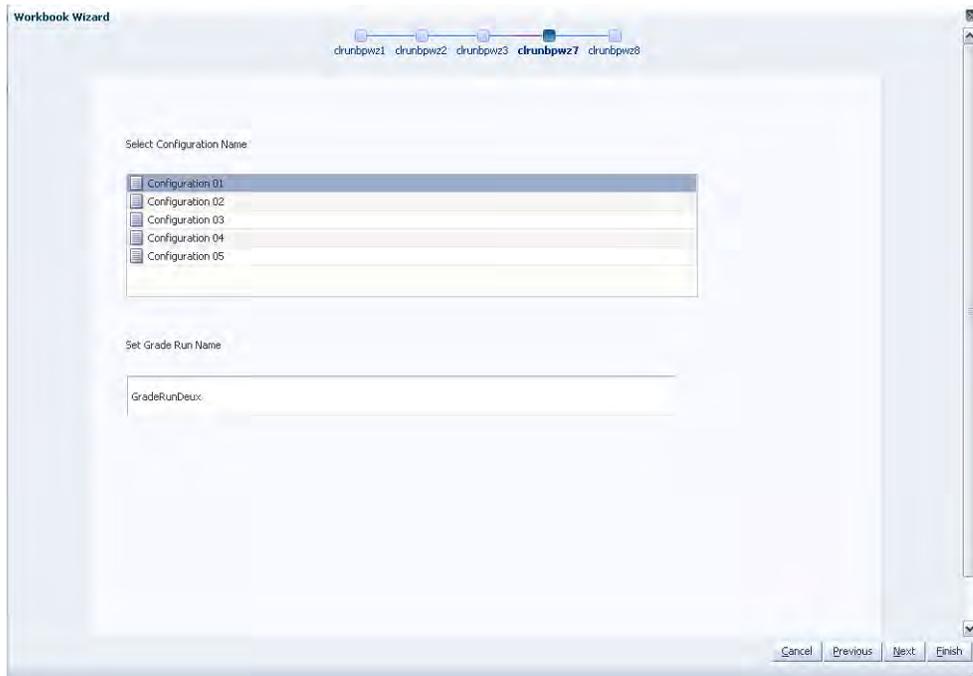
Figure 16–13 Workbook Wizard: Select Day



8. Select the Breakpoint Configuration that will be used to produce the grades. These configurations must be set prior to generating grades using the Breakpoint Administration task.

Click **Next**.

Figure 16–14 Workbook Wizard: Select Configuration Name



9. Set the Grade Run Name. Assign a label that will be used to identify the clustering run.
10. Select **Next** or **Finish**.

If **Next** is selected, Break Point Run Succeeded is displayed once the grading process is completed. If **Finish** is selected this message will be skipped.

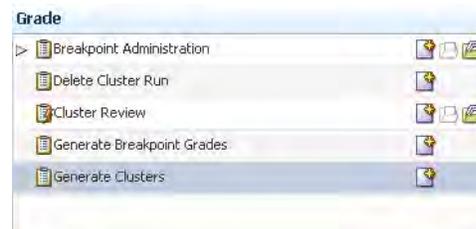
Generate Clusters Task

The Generate Clusters wizard allows you to range the input data and hierarchies that will be used to produce clusters based on the Clustering (BaNG) method.

Generate Clusters Wizard

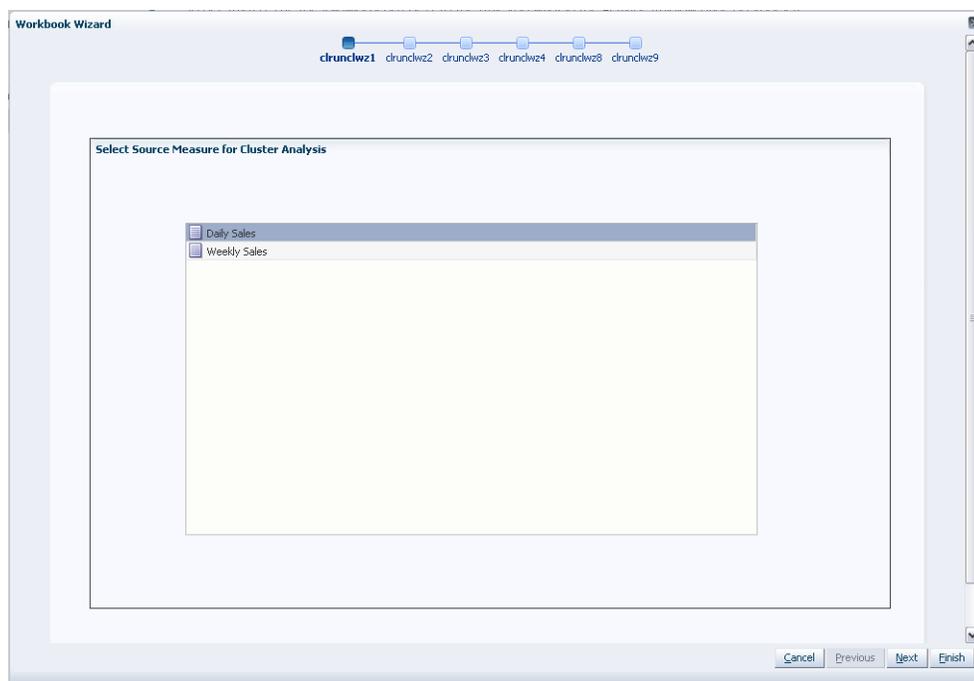
To run the Generate Clusters wizard, perform these steps:

1. Click the New Workbook icon in the Generate Clusters task.

Figure 16–15 Generate Clusters Task

2. The Workbook wizard opens. Select the source measure for Cluster Analysis and click **Next**.

The options include all source measures defined at the time of the configuration.

Figure 16–16 Workbook Wizard: Select Source Measure for Cluster Analysis

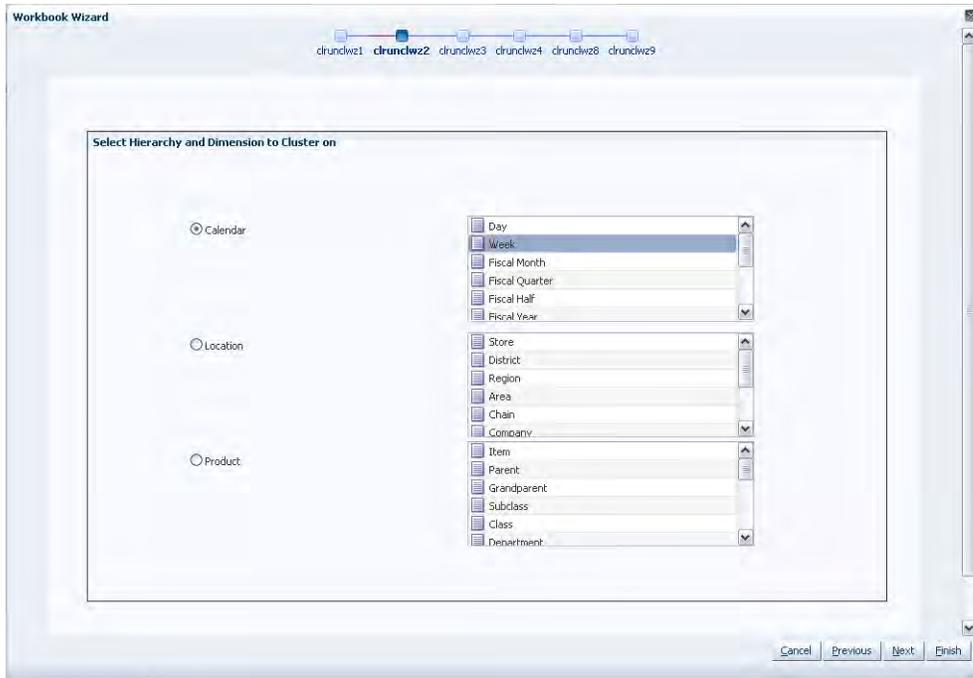
3. Select the Hierarchy and Dimension to Cluster On and click **Next**.

The Cluster Over dimension allows you to define the dimension that will be used for clustering. The algorithm uses the positions in this dimension as the co-ordinates when clustering. Internally, data will be aggregated to this level before performing the clustering. For example, when using SKU/Store/week level sales data to generate store clusters, the user has the option of specifying a dimension, such as Week, to which sales data will be aggregated. In this case, the algorithm considers Store/Week positions and clusters stores based on their weekly sales.

Note: The hierarchies presented in this wizard window include all hierarchies associated with the base intersection of the source measure except the hierarchy and dimension chosen to Cluster On.

Note: The dimension chosen should be at or higher than the base intersection of the selected Source Measure.

Figure 16–17 Workbook Wizard: Hierarchy and Dimension



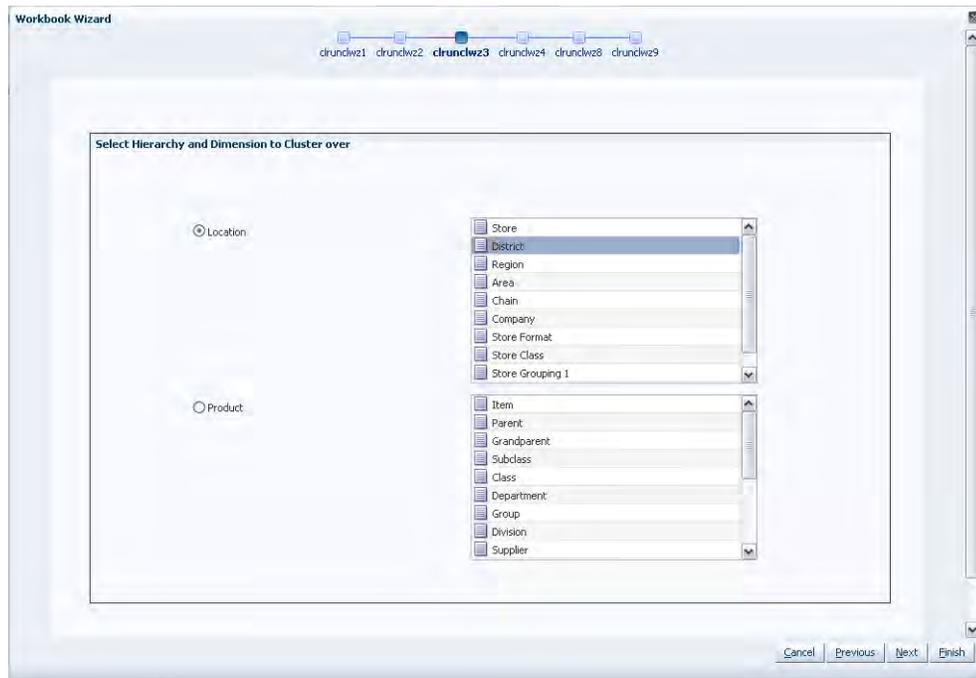
4. Optional Step

Select the Select Hierarchy and Dimension to Group By and click **Next**.

By Group is an optional setting that is used to partition data along a hierarchy dimension. This functionality allows for clusters to be further segmented along multiple hierarchy dimensions. Clusters are generated independently for each position within the selected dimensions. The dimensions displayed are equal to or higher than the Cluster On and Cluster Over dimensions.

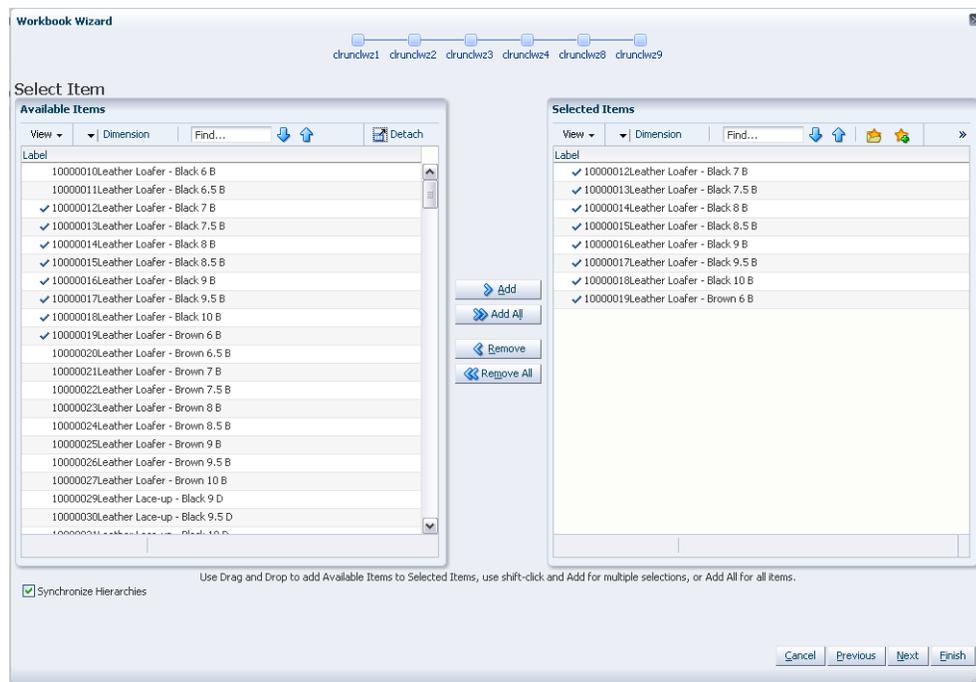
For example, while clustering stores by Dept, that is, cluster stores by a Department at a time, the Cluster By Group is Department.

Figure 16–18 Workbook Wizard: Select Hierarchy and Dimension to Group By



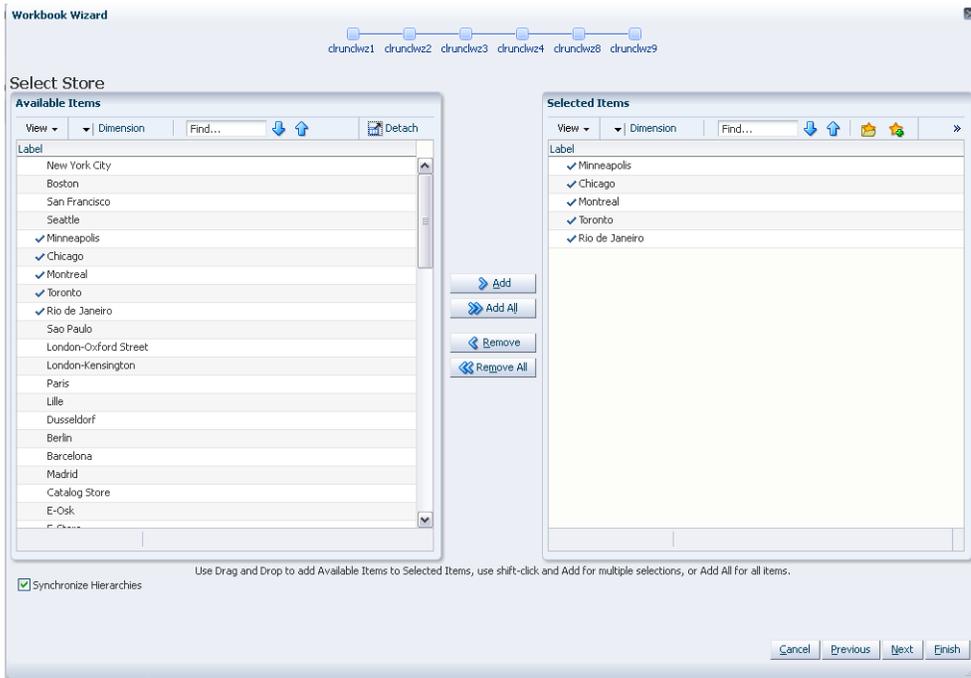
5. Select the products for which you want to generate clusters and click **Next**.

Figure 16–19 Workbook Wizard: Select Items



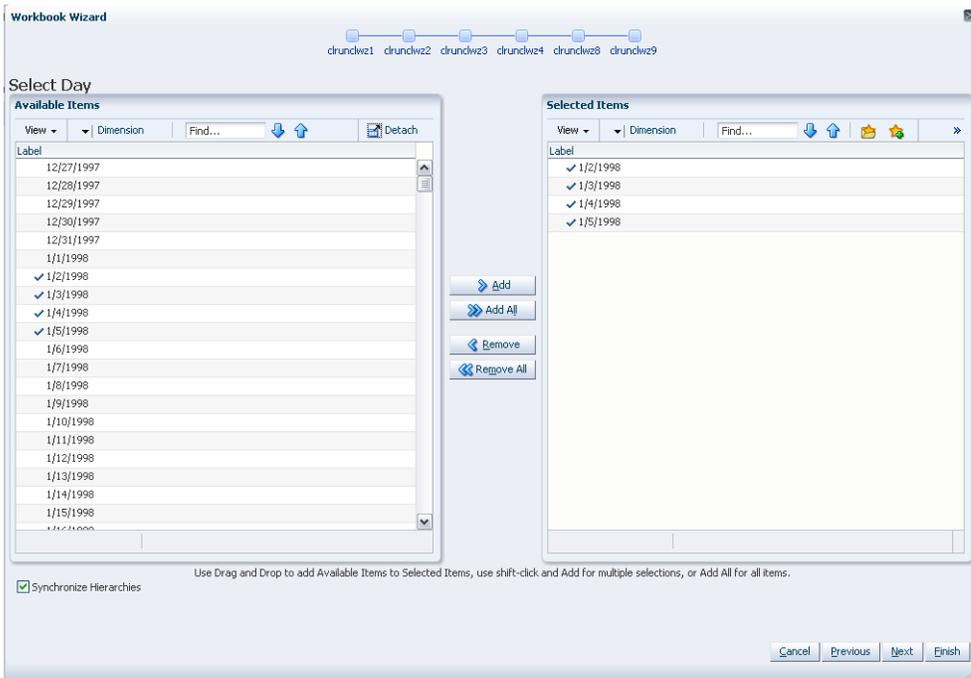
6. Select the locations for which you want to generate clusters and click **Next**.

Figure 16–20 Workbook Wizard: Select Stores

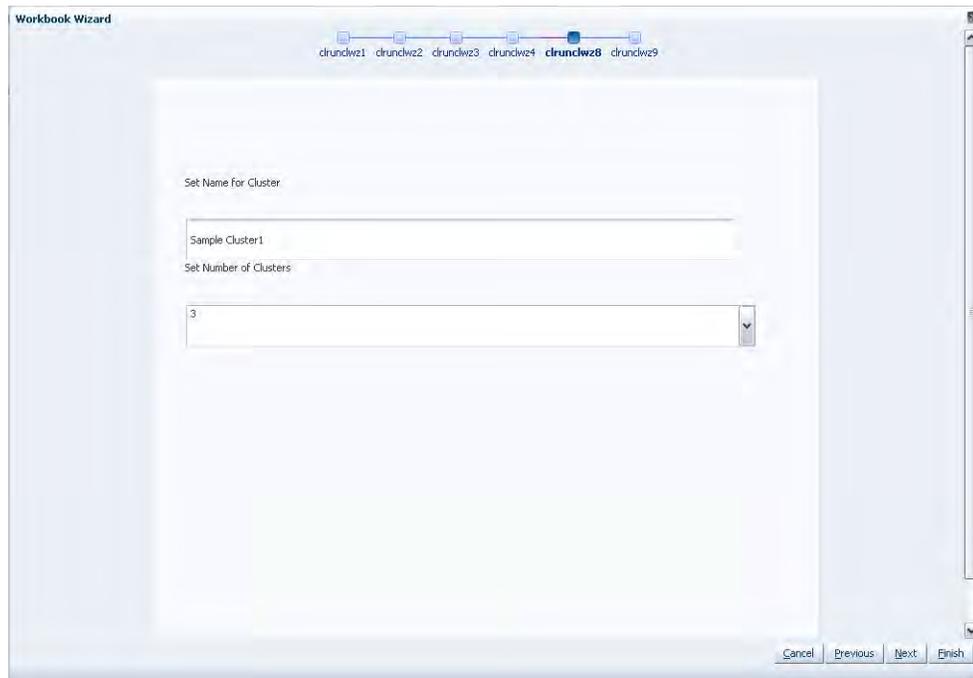


7. Select the time periods that you want to be considered when generating clusters and click Next.

Figure 16–21 Workbook Wizard: Select Day



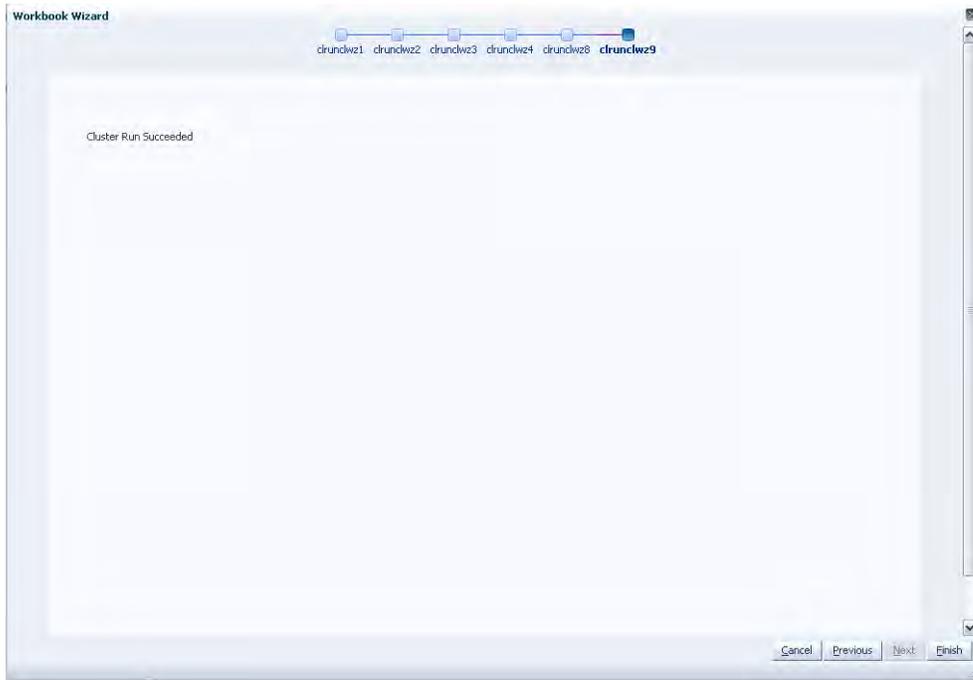
8. Assign a label that will be used to identify the clustering run. Click Next.

Figure 16–22 Workbook Wizard: Set Name and Number of Clusters

9. Set the name and number of clusters:
 - a. Set the name. Assign a label that will be used to identify the clustering run.
 - b. Set the number of clusters. Define the number of clusters that are generated during the cluster generation process. If a Group By dimension was selected, this is the number of clusters that will be generated for each data partition within the dimension. The number of clusters that may be generated is based on the maximum number of clusters configured in the Grade configuration. refer to the *Oracle Retail Grade Configuration Guide* for more information on the Grade solution configuration.
10. Select **Next** or **Finish**.

If **Next** is selected, Cluster Run Succeeded is displayed once the grading process is completed as shown in [Figure 16–23](#). If **Finish** is selected this message will be skipped.

Figure 16–23 Workbook Wizard: Cluster Run Succeeded



Appendix: Preprocessing

Preprocessing is a filtering module that automatically adjusts historical data to correct data points that do not represent general demand pattern. Essentially, it smoothes out spikes and dips in historical sales data, replacing stock-out data and data from short term events, such as promotions and temporary price changes, with data points that more accurately represent typical sales for that period. By adjusting the historical sales, Preprocessing can provide smarter data to the RDF Causal Engine, thus creating a smarter baseline forecast.

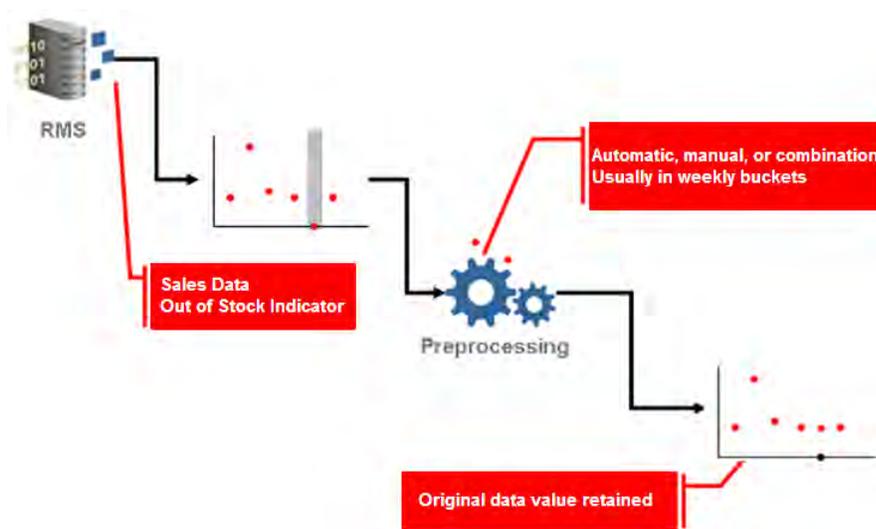
Note: There are no workbooks associated with Preprocessing - it is available as a configuration option.

Common Preprocessing corrections are:

- Out of stock — Interfaced from RMS, weekly or daily
- Outliers — Indicator not required, depends on method
- Short term events — Promotions, temporary price changes

For example, [Figure A-1](#) illustrates how Preprocessing adjusts for stock-outs.

Figure A-1 *Preprocessing for Stock-outs*



In [Figure A-1](#), RMS sends historical sales data to the Preprocessing module of RDF. In that sales data, RMS has flagged out-of-stock instances with indicators (the gray portion of the first data set). Preprocessing takes note of that out-of-stock indicator and adjusts the sales for that time period to reflect a more typical sales quantity, taking into account trending and seasonality. Note in [Figure A-1](#) that Preprocessing has removed the dip in sales in the second data set and has replaced it with a new data point.

Preprocessing Methods

Preprocessing uses several methods to massage historical data. The following sections detail these methods:

- [Standard Median](#)
- [Retail Median](#)
- [Standard Exponential Smoothing](#)
- [Lost Sales Standard Exponential Smoothing](#)
- [Override](#)
- [Increment](#)
- [Forecast Sigma](#)
- [Forecast Sigma Event](#)
- [Clear](#)
- [No Filtering](#)

Standard Median

Standard Median calculates baselines on long time ranges.

Input: none

Optional parameter: window length

Figure A-2 *Standard Median Formula*

$$LSOVER(t) = \text{median value over} \left(t - \frac{\text{window}}{2}, t + \frac{\text{window}}{2} \right)$$

Figure A-3 *Standard Median Example*



When data points for the full window are not available, Preprocessing pads the beginning and end of the time series with the first and the last data points, respectively, so that there are values for the full window.

Retail Median

Retail Median calculates baselines on long time ranges and improves side effects by making five standard median filter passes.

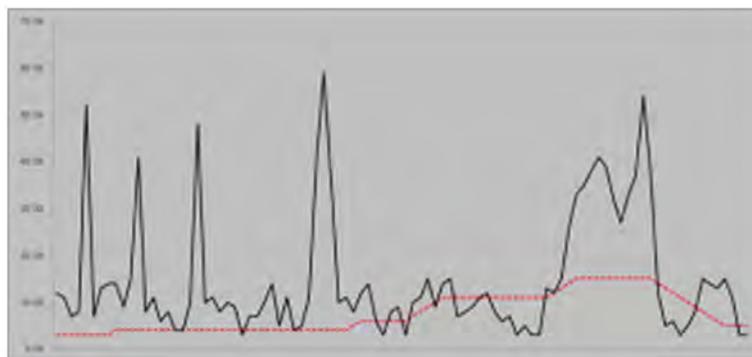
Input: none

Optional parameter: window length

Figure A–4 Retail Median Formula

$$\begin{aligned} \text{Median1} &= \text{StdMedian}(\text{Src}) \\ \text{Median2} &= \text{StdMedian}(\text{Median1}) \\ \text{Diff1}(t) &= \text{Median2}(t) - \text{Median2}(t-1) \\ \text{Median_Diff1} &= \text{StdMedian}(\text{Diff1}) \\ \text{Smooth1}(t) &= \text{Smooth1}(t-1) + \text{Median_Diff1}(t) \\ \text{Diff2}(t) &= \text{Smooth1}(t) - \text{Smooth1}(t-1) \\ \text{Avg_Diff2} &= \text{StdMedian}(\text{Diff2}) \\ \text{Smooth2}(t) &= \text{Smooth2}(t-1) + \text{Avg_Diff2}(t) \end{aligned}$$

Figure A–5 Retail Median Example



Standard Exponential Smoothing

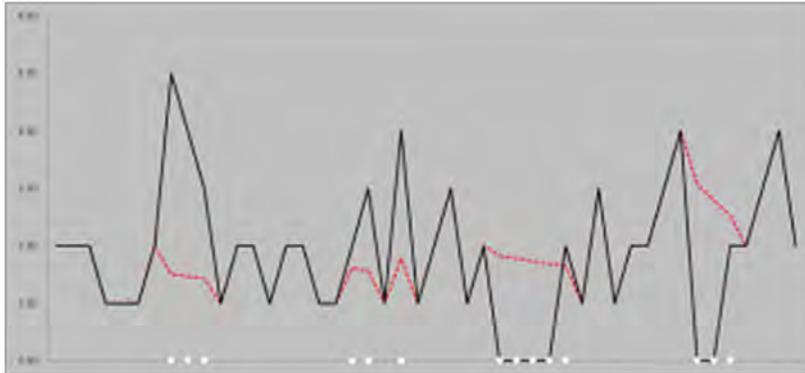
Standard Exponential Smoothing removes spikes (such as promotional promo, temporary price changes, and so on), as well as filling the gaps (out of stock, unusual events such as a fire or hurricane).

Input: event indicator

Optional parameters: ES (Exponential Smoothing) Parameter, number of future and past periods used to calculate the future/past velocities and an event flag.

Figure A-6 Standard Exponential Smoothing Formula

$$\begin{aligned}
 \text{past velocity} &= \frac{\sum_{i=1}^{np} (1-\alpha)^{i-1} * \text{src}(t_f - i)}{\sum_{i=1}^{np} (1-\alpha)^{i-1}} \\
 \text{future velocity} &= \frac{\sum_{i=1}^{nf} (1-\alpha)^{i-1} * \text{src}(t_l + i)}{\sum_{i=1}^{nf} (1-\alpha)^{i-1}} \\
 \text{LSOVER}(t) &= \text{past velocity} + \frac{\text{future velocity} - \text{past velocity}}{t_l - t_f + 2} * (t - t_f + 1), \text{ where } t \in [t_f, t_l]
 \end{aligned}$$

Figure A-7 Standard Exponential Smoothing Example

When event flags exist within the future and past velocity windows, rather than consider the entire window, Preprocessing only considers unflagged data points after the last event flag in the history window to compute the past velocity. It does a similar process for the future window by using the unflagged data points prior to the first event flag in the future window to compute the future velocity. Consecutive events are smoothed using the same velocities. A data point becomes flagged, and hence not part of the future/past velocity calculation, if either the event indicator or the optional event flag are on.

If future velocities cannot be calculated, then the past velocities, if they exist, are used as future and past velocities, and vice versa. When neither of the velocities can be calculated, there is no adjustment.

If the velocity window contains all zero values, then the calculated velocity is zero. A velocity of zero is a legitimate value if it occurs within the selling window. A velocity of zero is not acceptable if it is calculated based on values outside of the selling window.

Lost Sales Standard Exponential Smoothing

Lost Sales Standard Exponential Smoothing calculates baselines on long time ranges. Lost Sales Standard Exponential Smoothing makes positive adjustments to the flagged periods and to the period immediately following the flagged period.

Input: out-of-stock indicator

Optional parameters: ES (Exponential Smoothing) parameter and a number of future and past periods used to calculate the future/past velocities, an event flag, and a scalar parameter indicating if the period immediately following an out-of-stock period should be adjusted.

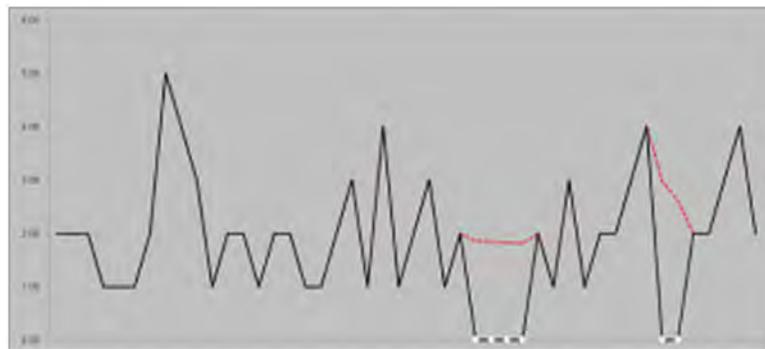
Figure A-8 Lost Sales Standard Exponential Smoothing Formula

$$\text{past velocity} = \frac{\sum_{i=1}^{np} (1-\alpha)^{i-1} * src(t_f - i)}{\sum_{i=1}^{np} (1-\alpha)^{i-1}}$$

$$\text{future velocity} = \frac{\sum_{i=1}^{nf} (1-\alpha)^{i-1} * src(t_l + i)}{\sum_{i=1}^{nf} (1-\alpha)^{i-1}}$$

$$LSOVER(t) = \text{past velocity} + \frac{\text{future velocity} - \text{past velocity}}{t_l - t_f + 2} * (t - t_f + 1), \text{ where } t \in [t_f, t_l]$$

Figure A-9 Lost Sales Standard Exponential Smoothing Example



When event flags exist within the future and past velocity windows, rather than consider the entire window, Preprocessing only considers unflagged data points after the last event flag in the history window to compute the past velocity. It does a similar process for the future window by using the unflagged data points prior to the first event flag in the future window to compute the future velocity. Consecutive events are smoothed using the same velocities. A data point becomes flagged, and hence not part of the future/past velocity calculation, if either the event indicator or the optional event flag is on.

If future velocities cannot be calculated, then the past velocities, if they exist, are used as future and past velocities, and vice versa. When neither of the velocities can be calculated, there is no adjustment.

If the velocity window contains all zero values, then the calculated velocity is zero. A velocity of zero is a legitimate value if it occurs within the selling window. A velocity of zero is not acceptable if it is calculated based on values outside of the selling window. Note that by default, the periods being adjusted are the periods flagged by an out-of-stock indicator and the period immediately following any such period. If the optional scalar parameter POA (Partial Outage Allowed) is set to **false**, then this extra period will not be adjusted, and only the out-of-stock periods will be adjusted.

Override

Override fills gaps in data when a reference measure exists.

Input: reference measure (R(t)) to copy data from

Optional parameter: outage/mask (M(i)), adjustment ratio (a)

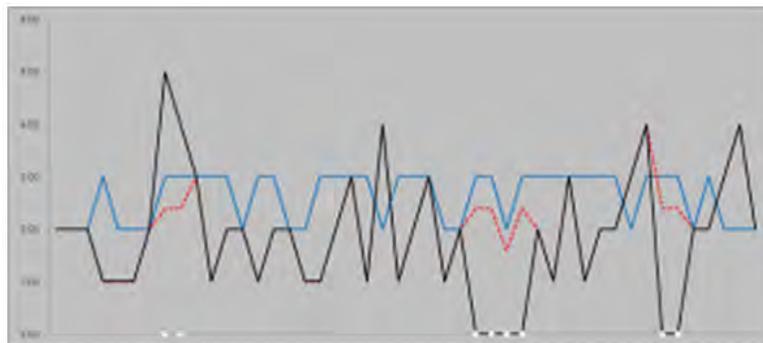
Formula: Overrides *LSOVER* with the Src adjusted by the adjustment ratio according to the mask:

Figure A–10 Override Formula

$$LSOVER(t) = a * R(t) \quad \text{if } M(i) \text{ is true}$$

$$LSOVER(t) = Src(t) \quad \text{if } M(i) \text{ is false}$$

Figure A–11 Override Example



Increment

Increment updates gaps or outliers in data when a reference measure exists.

Input: reference measure (R(t)) to copy data from

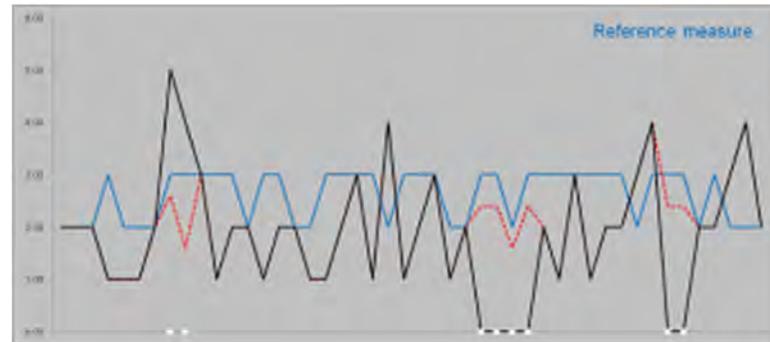
Optional parameter: outage/mask (M(i)), adjustment ratio (a)

Increments the Src with the reference adjusted by the adjustment ratio according to the mask:

Figure A–12 Increment Formula

$$LSOVER(t) = SRC(t) + a * R(t) \quad \text{if } M(i) \text{ is true}$$

$$LSOVER(t) = Src(t) \quad \text{if } M(i) \text{ is false}$$

Figure A–13 Increment Example

Forecast Sigma

Forecast Sigma removes spikes in recent data when no indicators are available.

Inputs: forecast and confidence intervals

Optional parameters:

- number of stddev for upper bound
- number of stddev for lower bound
- forecast lower bound
- minimum history required for filtering

Formula: If the difference in the sales and forecast is larger than a threshold, the override value is brought within some bounds of the forecast.

Forecast Sigma Event

Like Forecast Sigma, Forecast Sigma Event removes spikes in recent data when no indicators are available but also takes outage as input.

Inputs: outage, forecast, confidence intervals

Optional parameters:

- number of stddev for upper bound
- number of stddev for lower bound
- forecast lower bound
- minimum history required for filtering

Formula:

If outage is on:

$$LSOVER = forecast$$

Otherwise, if the difference in the sales and forecast is larger than a threshold, the override value is brought within some bounds of the forecast.

Clear

This Preprocessing method clears the Preprocessing adjustments from previous runs.

No Filtering

This Preprocessing method does not filter the source data.

Preprocessing for Stock-outs

When using Preprocessing to correct for stock-outs, the system expects out-of-stock indicators from a merchandising system like RMS. The system can be set up for automatic adjustment of sales history to correct for stock-outs as well as for manual user overrides under exception cases.

When set to automatically adjust sales history to correct stock-outs, Preprocessing takes into account trending and seasonality and adjusts the sales that were flagged by the out-of-stock indicator to reflect a more typical sales quantity.

Preprocessing for Promotional Forecasting

Preprocessing adjusts promotional data in a similar way that it does stock-outs. Typically, historical data shows a higher rate of sales during promotional periods. Were these spikes in sales to be left in historical sales data and loaded in the RDF Causal Engine, the baseline forecast created from this data would reflect similar spikes in future sales.

Appendix: RDF Workbooks (Fusion Client)

This appendix provides a convenient reference of all the RDF workbooks on the RPAS Fusion Client and their respective views and measures by task:

- [Forecast Product/Location Cloning Administration Task](#)
- [Forecast Administration Task](#)
- [Forecast Approval Task](#)
- [Forecast Delete Task](#)
- [Forecast Scorecard Task](#)
- [Interactive Forecasting Task](#)
- [Forecast Like-Item, Sister-Store Task](#)
- [Forecast Maintenance Task](#)
- [Run Batch Forecast Task](#)
- [Promote \(Promotional Forecasting\)](#)
 - [Promotion Effectiveness Task](#)
 - [Promotion Maintenance Task](#)
 - [Promotion Management Task](#)
 - [Promotion Planner Task](#)
- [Cross Promotional Effects Module \(CPEM\) Task](#)

Forecast Product/Location Cloning Administration Task

This section lists the workbook found in [Chapter 4, "Forecast Product/Location Cloning Administration Task,"](#) and its respective views and measures.

Product/Location Cloning Administration Workbook

[Table B-1](#) lists the views and measures for the [Product/Location Cloning Administration Workbook](#). For additional information, click any link to go to that section of the *Oracle Retail Demand Forecasting User Guide for the RPAS Fusion Client*.

Table B–1 Product/Location Cloning Administration Workbook

View	Measures
Adjustment Parameters View	Cloned History Adjustment - Alpha (Range 0-1) Cloned History Adjustment - Recent Sales Threshold Cloned History Adjustment - Calculation Window Length
Location Cloning View	Location Clones Location Clone Contributions Clone Location Adjustment Ratio
Product Cloning View	Product Clones Product Clone Contributions Clone Product Adjustment Ratio
Clone Mask View	Clone Mask

Forecast Administration Task

This section lists the workbook found in [Chapter 5, "Forecast Administration Task,"](#) and its respective views and measures.

Forecast Administration Workbook

[Table B–2](#) lists the settings, views, and measures for the [Forecast Administration Workbook](#). For additional information, click any link to go to that section of the *Oracle Retail Demand Forecasting User Guide for the RPAS Fusion Client*.

Table B-2 Forecast Administration Workbook

Settings	View	Measures	
Basic Settings	Final Level View - Basic Settings	Default Approval Method Default Source Level Forecast Cycle Forecast Start Date	Forecast Length Next Run Date
	Final and Source Level Parameters View - Basic Settings	Forecast Data Source Default History Start Date Data Plan Default Forecast Method	Seasonal Profile Spreading Profile Source Level Forecast Data Source Source Level History Start Date
Advanced Settings	Causal Parameters View - Advanced Settings	Causal - Disable Seasonal Regression Causal - Effects Only Causal - Lift Estimate Causal - Seasonal Baseline Selection Criteria Causal - Smoothing Method Causal Aggregation Profile Causal Baseline Minimum Winters History Length Causal Beta Scale Causal Calculation Intersection Causal Calculation Intersection Periodicity Causal Capping Number	Causal Capping Ratio Causal Data Source Causal External Baseline Causal Higher Intersection Causal Short Seasonal Threshold Causal SLC - ROS Causal Spread Profile Disable Causal Minimum Causal Capping History Minimum Causal Capping History (Periods) Use Causal Capping

Table B-2 (Cont.) Forecast Administration Workbook

Settings	View	Measures	
Advanced Settings	Final Level Parameters View - Advanced Settings	Components - Baseline Components - Baseline Interval Components - Regular Price Lifts Components - Promotional Lifts Days to Keep Forecasts Default Keep Last Changes Generate Baseline Forecasts Generate Cumulative Interval Generate Intervals	Generate Methods Generate Parameters Interval Cap Lower Ratio Interval Cap Upper Ratio Item End Date Action Like TS Duration (Periods) Store Interim Forecast Updating Last Week Forecast Updating Last Week Forecast Number of Weeks
	Final and Source Level Parameters View - Advanced Settings	Bayesian Alpha (range (0, infinity)) Bayesian Cap Ratio Croston's Min Gaps DD Duration (weeks) Default Moving Average Window Length Fallback Method Holt Min Hist (Periods) Max Alpha (Profile) (range [0,1]) Max Alpha (Simple, Holt) (range [0,1])	Max Alpha (Winters) (range [0,1]) Max Gamma (Holt) (range [0,1]) Max Gamma (Winters) (range [0,1]) Seasonal Smooth Index Trend Damping Factor (range [0,1]) Winters Min Hist (Periods) Winters Mode Winters Mode Impact on AutoES and Causal Deseasonalized Demand Array
	Regular Price Parameters View-Advanced Settings	Apply Regular Price Lift Apply Regular Price Lift over Promotion Regular Price Decaying Factor Apply Regular Price Self Lift	Apply Regular Price Halo Lift Apply Regular Price Cannibalization Lift Regular Price Effective Periods Price Elasticity Matrix Data Source
	Promo Affinity Parameters View	Apply Promotion Cannibalization Lift Apply Promotion Halo Lift	Promo Cannibalized Matrix Data Source Promo Halo Matrix Data Source

Forecast Approval Task

This section lists the workbook found in [Chapter 6, "Forecast Approval Task,"](#) and its respective views and measures.

Forecast Approval Workbook

[Table B-3](#) lists the views and measures for the [Forecast Approval Workbook](#). For additional information, click any link to go to that section of the *Oracle Retail Demand Forecasting User Guide for the RPAS Fusion Client*.

Table B-3 Forecast Approval Workbook

View	Measures	
Advanced Final Level View	Promotion Cannibalization Lift Units Promotion Halo Lift Units Promotion Indicator Promotion Self Lift Units Regular Price	Regular Price Cannibalization Lift Regular Price Halo Lift Regular Price Self Lift Total Promotion Lift Units Total Regular Price Lift Units
Final Level View	Adjusted Cumulative Interval Adjusted Forecast Approved Cumulative Interval Approved Forecast Approved System Forecast Forecast Cumulative Interval Forecast Interval History Data Interim Forecast Last Approved Forecast	System Baseline System Forecast Total Promotion Lift Units Total Regular Price Lift Units
Approval View	Approval Comment Approval Date	Approved By Manually Approved

Table B–3 (Cont.) Forecast Approval Workbook

View	Measures	
Valid Forecast Run View	Final Level Run Flag	Final Level Run Flag
Source Level View	History Data System Baseline	System Forecast
Final or Source System Parameters View	Forecast Method Picked System Generated Alpha System Generated Level	System Generated Gamma System Generated Trend

Forecast Delete Task

This section describes the wizard found in [Chapter 7, "Forecast Delete Task,"](#).

Wizard

The [Forecast Delete Wizard](#) is used in a global domain environment. Deleting a forecast in the Master domain deletes the selected birth date within all domains. Deleting a forecast in the Local domain deletes the selected birth date only within the Local domain.

Forecast Scorecard Task

This section lists the workbook found in [Chapter 8, "Forecast Scorecard Task,"](#) and its respective views and measures.

Forecast Scorecard Workbook

[Table B–4](#) lists the views and measures for the [Forecast Scorecard Workbook](#). For additional information, click any link to go to that section of the *Oracle Retail Demand Forecasting User Guide for the RPAS Fusion Client*.

Table B-4 Forecast Scorecard Workbook

View	Measures
Error Measure View	History Data for Forecast Window Forecast Measures Mean Error Mean Absolute Error Root Mean Squared Error Mean Absolute Percentage Error Percentage Absolute Error
Actuals versus Forecasts View	Actuals History Data Forecast Measures

Interactive Forecasting Task

This section lists the workbook found in [Chapter 9, "Interactive Forecasting Task,"](#) and its respective views and measures.

Interactive Forecasting Workbook

[Table B-5](#) lists the views and measures for the [Interactive Forecasting Workbook](#). For additional information, click any link to go to that section of the *Oracle Retail Demand Forecasting User Guide for the RPAS Fusion Client*.

Table B-5 Interactive Forecasting Workbook

View	Measures
Forecasting Parameter View	Forecast Method Forecast Picked Method History Start Date Forecast Start Date Forecast End Date
Interactive Forecasting View	History Data System Forecast

Forecast Like-Item, Sister-Store Task

This section lists the workbook found in [Chapter 10, "Forecast Like-item, Sister-Store Task,"](#) and its respective views and measures.

Forecast Like-Item, Sister-Store Workbook

[Table B-6](#) lists the views and measures for the [Forecast Like-Item, Sister-Store Workbook](#). For additional information, click any link to go to that section of the *Oracle Retail Demand Forecasting User Guide for the RPAS Fusion Client*.

Table B-6 Forecast Like-Item, Sister-Store Workbook

View	Measures
Like-Item View	Like SKU
Sister-Store View	Sister Store
Advanced Parameter View	Adjustment Ratio Forecast Start Date Override Forecast End Date Override History Start Date Override Substitute Methods

Forecast Maintenance Task

This section lists the workbook found in [Chapter 11, "Forecast Maintenance Task,"](#) and its respective views and measures.

Forecast Maintenance Workbook

[Table B-7](#) lists the settings, views, and measures for the [Forecast Maintenance Workbook](#). For additional information, click any link to go to that section of the *Oracle Retail Demand Forecasting User Guide for the RPAS Fusion Client*.

Table B-7 Forecast Maintenance Workbook

Settings	View	Measures	Measures
Basic Settings	Final and Source Level Views - Basic Settings	Approval Method Override Bayesian Alpha Override Forecast Method Override Keep Last Changes Override Moving Average Window Length Override	Optimal Source Levels Pick Optimal Level Source Level Override Winters Mode Override
Advanced Settings	Advanced Final Parameter View	Apply Promotion Cannibalization Lift Override Apply Promotion Halo Lift Override Apply Promotional Lift Override Apply Regular Price Cannibalization Lift Override Apply Regular Price Halo Lift Override Apply Regular Price Lift Over Promotion Override Apply Regular Price Lift Override Apply Regular Price Self Lift Override	Forecast End Date Override Forecast Start Date Override History Start Date Override Regular Price Decaying Factor Override Regular Price Effective Periods Override

Run Batch Forecast Task

This section describes the wizard found in [Chapter 12, "Run Batch Forecast Task,"](#).

Wizard

The [Run Batch Forecast Workbook](#) is a wizard that automatically executes PreGenerateForecast and Generate within the Simple or Local Domain.

Promote (Promotional Forecasting)

This section lists the tasks and workbooks found in [Chapter 13, "Promote \(Promotional Forecasting\),"](#) and their respective views and measures.

Promotion Effectiveness Task

This section lists the workbook found in the [Promotion Effectiveness Task](#) and its respective views and measures.

Promotion Effectiveness Workbook

Table B–8 lists the views and measures for the [Promotion Effectiveness Workbook](#). For additional information, click any link to go to that section of the *Oracle Retail Demand Forecasting User Guide for the RPAS Fusion Client*.

Table B–8 Promotion Effectiveness Workbook

View	Measures	
Promotion Model Type View	Model Type for Promotional Variable	
Promotions View	Promotion Variables	
Aggregated Promotion Lift Override View	Aggregated Promotion Lift Override Promotional Lift Override	
Promotional Forecast View	Forecast Baseline Price Cannibalization Lift Unit Price Halo Lift Units Price Self Lift Units Promotion Cannibalization Lift Units Promotion Halo Lift Units	Promotion Self Lift Units Promotional Forecast with Promote Lift Override Promotional Lift Override Total Promotional Lift Unit Total Regular Price Lift Unit
Promotion Parameters View	Apply Promotion Cannibalization Lift Apply Promotion Halo Lift Apply Promotional Lift Override Apply Regular Price Cannibalization Lift Apply Regular Price Halo Lift Apply Regular Price Lift	Apply Regular Price Lift Over Promotion Apply Regular Price Self Lift Promo Effect Type System Calculated Effect System Effect Override

Promotion Maintenance Task

This section lists the workbook found in the [Promotion Maintenance Task](#) and its respective views and measures.

Promotion Maintenance Workbook

Table B–9 lists the views and measures for the [Promotion Maintenance Workbook](#). For additional information, click any link to go to that section of the *Oracle Retail Demand Forecasting User Guide for the RPAS Fusion Client*.

Table B–9 Promotion Maintenance Workbook

View	Measures
Promotion Parameters View	Promo Effect Type System Calculated Effect System Effect Override

Promotion Management Task

This section lists the workbook found in the [Promotion Management Task](#) and its respective views and measures.

Promotion Management Workbook

[Table B–10](#) lists the views and measures for the [Promotion Management Workbook](#). For additional information, click any link to go to that section of the *Oracle Retail Demand Forecasting User Guide for the RPAS Fusion Client*.

Table B–10 Promotion Management Workbook

View	Measures
Promo Variable Model Type View	Model Type for Promotional Variable
Promotion Enable View	Enable Promotions
Accept Negative Lift View	Promotion Allow Negative Lift

Promotion Planner Task

This section lists the workbook found in the [Promotion Planner Task](#) and its respective views and measures.

Promotion Planner Workbook

[Table B–11](#) lists the views and measures for the [Promotion Planner Workbook](#). For additional information, click any link to go to that section of the *Oracle Retail Demand Forecasting User Guide for the RPAS Fusion Client*.

Table B–11 Promotion Planner Workbook

View	Measures
Promotions View	Promotion Variables

Cross Promotional Effects Module (CPEM) Task

This section lists the workbooks found in [Chapter 14, "Cross Promotional Effects Module \(CPEM\),"](#) and their respective tabs, views, and measures.

Effect Estimation Administration Workbook

[Table B–12](#) lists the tabs, views, and measures for the [Effect Estimation Administration Workbook](#). For additional information, click any link to go to that section of the *Oracle Retail Demand Forecasting User Guide for the RPAS Fusion Client*.

Table B–12 *Effect Estimation Administration Workbook*

Tabs	Views	Measures
Effect Estimation Administration	Default Parameters View	History Start Date History End Date Cannibalization Ancestor Level Cannibalization Lower Bound Cannibalization Upper Bound Halo Lower Bound Halo Upper Bound
	Cannibalization Parameter Override View	Cannibalization Lower Bound - Override Cannibalization Upper Bound - Override
	Halo Parameter Override View	Halo Lower Bound - Override Halo Upper Bound - Override
	Halo Products Mapping View	Halo Mask

Effect Estimation Review and Approval Workbook

[Table B–13](#) lists the tabs, views, and measures for the [Effect Estimation Review and Approval Workbook](#). For additional information, click any link to go to that section of the *Oracle Retail Demand Forecasting User Guide for the RPAS Fusion Client*.

Table B-13 *Effect Estimation Review and Approval Workbook*

Tabs	Views	Measures
Cross Promo Effect	Cannibalization Effects View	System Calculated Cannibalization Effect Approved Cannibalization Effect
	Cannibalization Effects Drive View	System Calculated Cannibalization Effect Approved Cannibalization Effect
	Cannibalization History Data View	Weekly Sales Baseline - Cannibalization Weekly Sales History - Cannibalization Weekly Normalized Price - Cannibalization Weekly Promotion Variable - Cannibalization
	Halo Effects View	System Calculated Halo Effect Approved Halo Effect
	Halo Effects Drive View	System Calculated Halo Effect Approved Halo Effect
	Halo History Data View	Weekly Sales Baseline - Halo Weekly Sales History - Halo Weekly Normalized Price - Halo Weekly Promotion Variable - Halo
Cross Effect Statistics Reports	Cannibalization Statistic Report View	Number of Effects Found Minimum Effect Found Maximum Effect Found
	Halo Statistic Report View	Number of Effects Found Minimum Effect Found Maximum Effect Found

Appendix: Curve Workbooks (Fusion Client)

This appendix provides a convenient reference of all the Curve workbooks on the RPAS Fusion Client and their respective views and measures by task.

The following sections list the workbooks, views, and measures found in [Chapter 15, "Curve"](#):

- [Profile Administration Task](#)
- [Profile Approval Task](#)
- [Profile Maintenance Task](#)
- [Run Batch Profile Task](#)

Profile Administration Task

This section lists the workbook found in the [Profile Administration Task](#) and its respective views and measures.

Profile Administration Workbook

[Table C-1](#) lists the views and measures for the [Profile Administration Workbook](#). For additional information, click any link to go to that section of the *Oracle Retail Demand Forecasting User Guide for the RPAS Fusion Client*.

Table C-1 Profile Administration Workbook

View	Measures	
Profile Parameter View	Default Phase End Default Phase Start Default Profile Approval Method Default Source Profile Default Training Window End Default Training Window Start Normal Value	Profile Data Source Profile Type Renormalize System Training Window Length Training Window Method Use Advanced Profile Features
Profile and Source Level Intersection View	Profile Agg Intersection Profile Approval Intersection	Profile Intersection Stored Intersection

Profile Approval Task

This section lists the workbook found in the [Profile Approval Task](#) and its respective views and measures.

Profile Approval Workbook

[Table C–2](#) lists the views and measures for the [Profile Approval Workbook](#). For additional information, click any link to go to that section of the *Oracle Retail Demand Forecasting User Guide for the RPAS Fusion Client*.

Table C–2 Profile Approval Workbook

View	Measures
Final Profile View	Adjusted Profile Approved Profile System Profile
Source Profile View	System Profile
Profile Approval View	Profile Approved By Profile Approval Date Manually Approve Source Profile Override

Profile Maintenance Task

This section lists the workbook found in the [Profile Maintenance Task](#) and its respective views and measures.

Profile Maintenance Workbook

[Table C–3](#) lists the views and measures for the [Profile Maintenance Workbook](#). For additional information, click any link to go to that section of the *Oracle Retail Demand Forecasting User Guide for the RPAS Fusion Client*.

Table C–3 Profile Maintenance Workbook

View	Measures
Profile Maintenance: Final Approval and Sourcing View	Profile Approval Method Source Profile Override
Profile Maintenance: Final Training Window View	System Training Window Begin System Training Window End

Run Batch Profile Task

The [Run Batch Profile Task](#) allows you to run the [Run Batch Profile Workbook](#); a wizard that automatically executes curvebatch within the Simple or Local Domain.

Appendix: Grade Workbooks (Fusion Client)

This appendix provides a convenient reference of all the Grade tasks, wizards, and workbooks on the RPAS Fusion Client.

The following sections list the tasks found in [Chapter 16, "Grade"](#):

- [Breakpoints Administration Task](#)
- [Delete Cluster Run Task](#)
- [Cluster Review Task](#)
- [Breakpoint Grades Task](#)
- [Generate Clusters Task](#)

Breakpoints Administration Task

Use the [Breakpoints Administration Task](#) to generate the [Breakpoints Administration Workbook](#) that allows you to set multiple breakpoint configurations to allow for grades to be produced and compared using different breakpoint settings.

Delete Cluster Run Task

The [Delete Cluster Run Wizard](#) allows you to delete clusters from the system based on Cluster Run Label and generation date (birth date).

Cluster Review Task

The [Cluster Review Task](#) is a wizard that generates the [Cluster Review Workbook](#), which allows you to view grade or cluster results and statistics.

[Table D-1](#) lists the views and measures for the [Cluster Review Workbook](#). For additional information, click any link to go to that section of the *Oracle Retail Demand Forecasting User Guide for the RPAS Classic Client*.

Table D-1 Cluster Review Workbook

View	Measures	
Cluster Input Summary View	Number of Clusters	Measure to Cluster
	Dimension to Cluster Over	Dimension to Cluster
	Cluster Run Name	Cluster By Group Intersection

Breakpoint Grades Task

The [Generate Breakpoint Grades Task](#) is a wizard that allows you to range the input data and hierarchies, which is used to produce grades based on the Breakpoints method.

Generate Clusters Task

The [Generate Clusters Task](#) is a wizard that allows you to range the input data and hierarchies which is used to produce clusters based on the Clustering (BaNG) method.

Glossary

Note: With a few exceptions, this glossary contains definitions of terms specific to RDF. For further definitions of terms and concepts relating to the RPAS user interface, refer to the *Oracle Retail Predictive Application Server Online Help* or *Oracle Retail Predictive Application Server User Guide*.

Additive Seasonal Method

Also referred to as Additive Winters Model, this model is similar to the Multiplicative Winters model but it is used when zeros are present in the data. This model adjusts the un-seasonalized values by adding the seasonal index for the forecast horizon.

Alert

A notice displayed to system users that a forecast value is above or below user-defined limits (an exception).

Alert Manager Window

A window that displays the alerts assigned to you. This dialog provides a list of all identified instances in which a monitored measure's values fall outside a set of defined limits. You may pick an alert from this list and have RCS automatically build a workbook containing the measure values that triggered the alert.

AutoES Method or Automatic Exponential Smoothing Method

RDF fits the sales data to a variety of exponential smoothing (time series) models of forecasting, and the best model is chosen for the final forecast. The candidate methods considered by AutoES are:

- SimpleES
- IntermittentES
- TrendES
- Multiplicative Seasonal
- Additive Seasonal and SeasonalES

The final selection between the models is made according to a performance criterion (Bayesian Information Criterion) that involves a trade-off between the model's fit over the historic data and its complexity.

Bayesian Method

Useful for short lifecycle forecasting and for new products with little or no historic sales data. The Bayesian method requires a product's known sales plan (created externally to RDF) and considers a plan's shape (the selling profile or lifecycle) and scale (magnitude of sales based on Actuals). The initial forecast is equal to the sales plan, but as sales information comes in, the model generates a forecast by merging the sales plan with the sales data. The forecast is adjusted so that the sales magnitude is a weighted average between the original plan's scale and the scale reflected by known history.

Causal Method

Causal is a forecasting method used for promotional forecasting and can only be selected if Promote is implemented. Typically, the Causal method is used at the Final Levels (that is, item/week/week). Causal uses a Stepwise Regression sub-routine to determine the promotional variables that are relevant to the time series and their lift effect on the series. AutoES utilizes the time series data and the future promotional calendar to generate future baseline forecasts. By combining the future baseline forecast and each promotion's effect on sales (lift), a final promotional forecast is computed.

Croston's Model of Exponential Smoothing

See [IntermittentES](#) or [Intermittent Exponential Smoothing](#).

Curve

An optional automated predictive solution that transforms organization-level assortment plans into base-level weekly sales forecasts.

Exception

A forecast value that is greater than or less than a user-defined limit.

Exponential Smoothing

A form of a weighted moving average. Its weight declines in data exponentially. Most recent data is weighted more heavily. It requires the smoothing constant. It ranges from 0 to 1 and is subjectively chosen.

Final Forecast Level

A low level in a hierarchy from which a forecast is generated and at which approvals and data exports can be performed. Often, data from forecasts at a low level is insufficient to generate reliable forecasts without first aggregating the data to a higher level and then spreading the data back to the low level.

Forecast

In RDF, Forecast refers to RDF's statistical forecasting capabilities.

Forecast-Driven Planning

Planning that keys off of forecasts fed directly into a planning system. Connection to RDF is built directly into the business process supported by Oracle Retail Predictive Planning through an automatic approval of a forecast that is fed directly in the planning system. This allows you to accept all or part of Sales Value forecast. Once that decision is made, the balance of business measures are planned within Oracle Retail Predictive Planning.

Halo

Used to explain the bias shown by customers towards certain products because of a favorable experience with other products.

Holt's Model of Exponential Smoothing

See [Trend Exponential Smoothing](#) or [TrendES](#).

Interactive Forecasting

A workbook in RDF that is used to simulate forecast by modifying parameters such as Forecast Method and History Start Date.

IntermittentES or Intermittent Exponential Smoothing

RDF fits the data to the Croston's model of exponential smoothing. This method should be used when the input series contains a large number of zero data points (that is, intermittent demand data). The original time series is split into a Magnitude and Frequency series, and then the SimpleES model is applied to determine level of both series. The ratio of the magnitude estimate over the frequency estimate is the forecast level reported for the original series.

Like-item or Like SKU

An item that is used as a model to forecast a new item introduction.

Lost Sales

Periods in sales data in which there was no inventory to meet consumer demand.

Measure

Any item of data that can be represented on a grid in a view.

Measure Description

The description of the measure that can be viewed in a workbook. This description may contain relationships and calculations.

Measure Function

Internal functions that can be used to simplify building calculations for a measure.

Measure Identifier

The combination of role, version, metric, and units that uniquely specifies a single measure.

Metric

A measure definition with the role, version, and units omitted.

Moving Average

For each period t , the moving average method takes the average of the periods from $t-2$ to $t+2$ as the smoothed baseline.

Multiplicative Seasonal

Also referred to as Multiplicative Winters Model, this model extracts seasonal indices that are assumed to have multiplicative effects on the un-seasonalized series.

Preprocessing

In RDF, Preprocessing refers to a module that processes data before forecasts are generated to adjust for situations, such as lost sales and unusually high demand.

Profile

Spreading ratios that are used in the Curve process. Typical profiles can include store participation, size distribution, and time (phase-to-week) profiles, as well as other information. Profiles are generated using historical data and phase definitions based on your system configuration.

Profile Based

RDF generates a forecast based on a seasonal profile that can be created in Curve or a legacy system. Profiles can also be copied from another profile and adjusted. Using historic data and the profile, the data is de-seasonalized and then fed to the SimpleES method. The Simple forecast is then re-seasonalized using the profiles.

Profile Spread

Used at the final-level to utilize a profile (either generated externally or with Curve) to determine the spreading ratios from the Source-level forecast down to the Final-level forecast.

Promote

Promote is an optional add-on automated predictive solution that allows you to incorporate the effects of promotional and causal events, such as radio advertisements and holiday occurrences, into your time series forecasts. The promotional forecasting process uses both past sales data and promotional information to forecast future demand.

Promotion Planning

A workbook and simulation process used within the context of promotional forecasting. Promotion planning involves specifying whether the event status for a particular promotional variable is active (on) or inactive (off) for a specific product/location/calendar combination. When past promotional events are represented as accurately as possible, the modeling routine can more precisely detect correlation between event occurrences and changes in sales values.

Promotional Effectiveness

A workbook used in the context of promotional forecasting. This workbook allows you to analyze the effects of promotions on items at both the micro and the macro level. What if analysis can also be performed on the results of promotional forecasts, as you can modify future and past promotional inputs, the system-estimated effects of promotions, and the promotional forecasts themselves.

Promotional Forecasting

Promote's forecasting technique (also referred to as Causal forecasting) uses promotional factors and events to predict future demand. Promotion events are events, such as advertisements, holidays, competitor information, and other factors that affect the normal selling cycle for a business.

Promotion Group

A set of products or locations that are believed to exhibit similar effects during common causal events. Promotion groups should be established to maximize the number of time series for each group (so each promotional event can be evaluated from as many different observations as possible), while ensuring that each time series is affected by causal events to the same degree.

SeasonalES Method

A combination of several Seasonal methods. This method is generally used for known seasonal items or forecasting for long horizons. This method applies the Multiplicative Seasonal model unless zeros are present in the data, in which case the Additive Winters model of exponential smoothing is used. If less than two years of data is available, a Seasonal Regression model is used. If there is too little data to create a seasonal forecast (in general, less than 52 weeks), the system selects from the SimpleES, TrendES, and IntermittentES methods.

Seasonal Regression

Seasonal Regression cannot be selected as a forecasting method, but it is a candidate model used when the SeasonalES method is selected. This model requires a minimum of 52 weeks of history to determine seasonality. Simple Linear Regression is used to estimate the future values of the series based on a past series. The independent variable is the series history one-year or one cycle length prior to the desired forecast period, and the dependent variable is the forecast. This model assumes that the future is a linear combination of itself one period before plus a scalar constant.

Simple/IntermittentES Method

A combination of the SimpleES and IntermittentES methods. This method applies the SimpleES model unless a large number of zero data points are present. In this case, the Croston's model is applied.

SimpleES or Simple Exponential Smoothing Method

RDF uses a simple exponential smoothing model to generate forecasts. SimpleES ignores seasonality and trend features in the demand data, and it is the simplest model of the exponential smoothing family. This method can be used when less than one year of historic demand data is available.

Simple Moving Average

See [Moving Average](#).

Sister Store

A store that is used as a model to forecast a new store.

Source Level Forecast

The level at which the aggregate, more robust forecast is run.

Time Series

Set of evenly spaced numerical data obtained by observing response variable at regular time periods. This data is used to forecast based only on past values. It assumes that factors influencing past and present continues influence in future.

Training Window

The number of weeks of historical sales data to use in generating a forecast.

Trend Exponential Smoothing or TrendES

Also referred to as Holt's Model, RDF fits the data to the Holt model of exponential smoothing. The Holt model is useful when data exhibits a definite trend. This method separates out base demand from trend and then provides forecast point estimates by combining an estimated trend and the smoothed level at the end of the series.

Wizard

A set of windows that guide you through the process of creating a new workbook or performing other actions in a solution by asking you various questions and having you select values.

Workbook

The framework used for displaying data and user functions. Workbooks are task-specific and may contain one or more views. Users define the format of their workbooks. Also see [Workbook Template](#), [View](#).

Workbook Template

The framework for creating a workbook. You build each new workbook from an existing workbook template, such as Pre-Season Financial Plan or Forecasting Administration. Several workbook templates are supplied with the Oracle Retail Predictive Solutions and are available for selection when you choose **File - New** to create a new workbook.

View

A multidimensional spreadsheet used to display workbook-specific information. View data can also be displayed in chart format.