Oracle® Retail Science Cloud Services

User Guide

Release 18.0.001

F17712-01

April 2019

ORACLE®

Oracle Retail Science Cloud Services User Guide, Release 18.0.001

F17712-01

Copyright © 2019, Oracle and/or its affiliates. All rights reserved.

Primary Author: Judith Meskill

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

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

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

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

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

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

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

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Value-Added Reseller (VAR) Language

Oracle Retail VAR Applications

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

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

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

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

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

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

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

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

Contents

Send Us Your Comments xv

Preface xvii

1 Getting Started

About Oracle Retail Science Cloud Services 1-1
Administration 1-1
Oracle Retail Science Dashboard 1-1
Process 1-2
Overview of the User Interface 1-7
Process Train 1-7
View Menu 1-7
Embedded Help 1-7
Process Indicator 1-8
Search 1-8
Icons 1-8
Buttons 1-9
Browser Settings 1-10
Setting Up Internet Explorer 1-10
Supported Characters for Text Entry 1-11
Histograms 1-11

2 Customer Decision Trees

Introduction 2-1
Overview of CDT Process 2-1
Overview Tab 2-2
CDT Generation Stage Status 2-2
Aggregate Statistics 2-3
Calculation Report 2-4
CDT Score Distribution Histogram 2-6
Generate CDT Tab 2-7
Data Setup 2-7
Category Selections 2-8
Time Interval Setup 2-9
Data Filtering 2-10

Filter Setup 2-10
Data Filtering Summary 2-11
Data Filtering Histograms 2-12
Calculation 2-13
Version Setup 2-14
Category Attributes Setup 2-15
Calculation Report 2-16
CDT Score Distribution Histogram 2-17
Evaluation 2-18
Pruning Setup 2-18
Pruning Results 2-19
Evaluation Report 2-20
Evaluation Histograms 2-21
Escalation 2-22
Setup Escalation 2-22
Escalation Report 2-23
Completion of Process 2-24
Manage CDTs Tab 2-24
Browse by Category or Version 2-26
Search 2-26
Create CDT 2-26
Deleting a CDT 2-27
Using the CDT Editor 2-27
Add a Child 2-28
Editing Nodes 2-28
Deleting Nodes 2-28
Copying/Pasting in Customer Decision Tree Editor 2-29
Viewing a Customer Decision Tree 2-29
Comparing Two CDTs 2-29

3 Demand Transference

Introduction 3-1
Overview of DT Process 3-1
Overview Tab 3-2
Generation Stage Status 3-2
Aggregate Statistics 3-3
Calculation Report 3-4
Generate Models Tab 3-4
Data Setup 3-5
Category Selections 3-5
Data Filtering 3-7
Filter Setup 3-8
Data Filter Summary 3-8
Data Filtering Histograms 3-9
Similarity Calculation 3-10
Version Setup 3-11
Category Attribute Setup 3-11

Transaction-Based Similarity Availability Per Category 3-13
Similarity Display 3-13
Elasticity Calculation 3-14
Calculation Report 3-15
Pruning Report 3-15
Assortment Elasticity Histogram 3-17
Escalation 3-17
Setup Escalation 3-17
Escalation Report 3-18
Completion of Process 3-19
Manage Models Tab 3-20
Time Interval Setup 3-20
Calculate Substitutable Demand Percentages 3-21
Browsing and Searching 3-21
Escalation Report 3-22
Histogram 3-23

4 Advanced Clustering

Introduction 4-1
Features 4-1
Overview of Advanced Clustering Process 4-2
Cluster Criteria Overview Tab 4-2
Clustering Criteria 4-3
Generate Store Clusters Tab 4-4
Cluster Criteria Stage 4-4
All Cluster Criteria 4-4
All Cluster Criteria Filter 4-6
All Cluster Criteria Summary 4-6
All Cluster Criteria Scenario List 4-6
Cluster Criteria 4-7
Effective Period 4-8
Summarization 4-9
Source Time Period 4-9
Contextual Area 4-10
Explore Data 4-11
Summary 4-11
View Stores 4-12
Product Performance 4-12
Contextual Area 4-13
Cluster Setup Stage 4-15
Summary 4-16
Scenario Definition Section 4-16
Cluster Results Stage 4-20
Summary 4-20
Scenario Results Section 4-21
Export to Excel 4-24
Scenario List 4-25

Scenario Compare 4-26
Scenario System Recommendations 4-27
Scenario Optimality 4-28
Scenario Rank 4-28
Outlier Indicator 4-28
Cluster Scores 4-28
New Stores and Stores with a Poor History 4-28
Insights Stage 4-29
Summary 4-30
Contextual Information 4-31
Manage Store Clusters Tab 4-31
Contextual Information 4-32
Scenario Compare 4-34
Attribute Analysis 4-35
Cluster Overrides 4-36
Rename Cluster 4-36
Create New Manual Cluster 4-36
Delete Cluster 4-37
Move Stores to Clusters 4-37
Cluster Review, Approve, and Adjust 4-38
Nested vs. Mixed Attribute 4-38
Nested 4-38
Mixed Attributes 4-39

5 Customer Segmentation

Introduction 5-1
Features 5-1
Overview of the Customer Segmentation Process 5-2
Customer Segmentation Overview Tab 5-2
Segmentation Criteria 5-3
Version List Tab 5-4
Version Actions 5-5
Version Criteria Pop-up 5-6
Source Time Period 5-6
Top Categories 5-6
Generate Customer Segments Tab 5-7
Segment Criteria Stage 5-8
All Segment Criteria 5-8
All Segment Criteria Filter 5-9
All Segment Criteria Summary 5-10
All Segment Criteria Scenario List 5-10
Segment Criteria 5-11
Effective Period 5-12
Summarization 5-12
Source Time Period 5-12
Contextual Area 5-13
Explore Data 5-13

Summary 5-14
Attribute Mining 5-14
Attribute Importance 5-14
Attribute Correlation 5-16
Filter 5-16
Correlation Matrix 5-16
Contextual Area 5-17
Segment Setup Stage 5-18
Scenario Definition Tab 5-19
Scenario List Tab 5-21
Scenario Compare Tab 5-22
Segment Results Stage 5-22
Summary 5-23
Scenario Results Section 5-23
Scenario List 5-27
Scenario Compare 5-28
Scenario System Recommendations 5-29
Scenario Optimality 5-29
Scenario Rank 5-29
Outlier Indicator 5-30
Segment Rules 5-30
Segment Scores 5-30
New Customers 5-31
Generate Store Profile 5-31
Segment Insights Stage 5-31
Customer Segment Analysis 5-31
Category Analysis 5-32
Opportunity Analysis Tab 5-32
Summary 5-34
Cross Segment Spend and Margin Analysis BI 5-34
RFM Analysis BI 5-35
Customer Segment vs. All Customer Sales Trends BI 5-36
Customer Attribute Analysis BI 5-36
Category Analysis BI 5-37
Trend Category Analysis BI 5-38

6 Assortment and Space Optimization

Overview 6-1
Administering ASO 6-1
Users and Roles 6-2
Overview of the User Interface 6-4
Icons 6-5
Buttons 6-6
Pull-Down Menus 6-7
Histograms 6-8
Charts 6-8
Process Train 6-8

Embedded Help 6-8
Process Indicator 6-8
Search 6-8
Workflow 6-8
Optimization 6-9
Space Cluster 6-9
Run Templates 6-9
Validation 6-9
Planogram Mappings 6-10
Smart Start Process 6-10
Fixture Types 6-10
Virtual Planogram 6-10
Planogram Set 6-10
Assortment Finalization 6-11
Assortment List 6-11
Assortment List 6-11
About Rollbacks 6-13
Assortment Analysis 6-13
Optimization List 6-17
Optimization List Interface 6-17
Create a Run 6-19
Copy a Run 6-19
Open a Run 6-19
Delete a Run 6-19
Execute Space Optimization 6-19
Optimization Setup 6-20
Summary 6-20
Setup 6-21
Review Optimization Locations 6-23
Contextual Information 6-26
Fixture and Product Data 6-28
Summary 6-28
Planogram Fixtures 6-29
Product Merchandise Data 6-37
Demand and Replenishment Data 6-42
Objective and Constraints 6-44
Summary 6-45
Objectives and Constraints Summary 6-45
Select Locations and Products 6-47
Objectives and Constraints Details 6-48
Contextual Information 6-58
Results and Analysis 6-59
Optimization Results 6-60
Included Products 6-61
Dropped Products 6-61
Alerts 6-61
Contextual Information 6-62

Planograms 6-62
Validation 6-65
Details About Errors and Warnings 6-66
Warnings 6-66
Correcting Errors 6-70
Assortment Mapping 6-71
Assortment Mapping Interface 6-71
Select Assortment 6-71
Assortment to Planogram Mapping 6-72
Mapping Errors 6-74

7 Attribute Extraction

Introduction 7-1
Overview Tab 7-1
Edit Labels Tab 7-2
Annotation Tab 7-3
Unigraphs 7-3
Assigning User Labels and Reviewing Machine Labels 7-3
Finding New Values 7-5
Digraphs 7-6
Assigning User Labels to Digraphs 7-6
Errors Tab 7-7
Normalization Tab 7-8
Results Tab 7-11
Status 7-11

8 Affinity Analysis

Introduction 8-1
User Interface 8-2
Assortment Recommender 8-4
Using Assortment Recommender Results 8-9
Market Basket Analysis 8-10
Using Market Basket Rules 8-15
Affinity Analysis 8-16
Affinity Analysis Overview 8-16
Compare Results 8-23
Create Custom Run 8-25
Manage Affinity Analysis 8-27
Compare Results 8-27

9 Social Analytics

Introduction 9-1
Overview Screen 9-1
Card 1 9-1
Card 2 9-2
Card 3 9-2

Social Analytics Screen 9-2
Trend Analysis Tab 9-2

10 Offer Optimization

Introduction 10-1
Administering Offer Optimization 10-1
Users and Roles 10-1
Overview of the User Interface 10-3
Icons 10-4
Buttons 10-5
Pull-Down Menus 10-6
Charts 10-7
Process Train 10-7
Embedded Help 10-7
Process Indicator 10-7
Search 10-7
Workflow 10-8
Optimization 10-8
Offer Optimization Overview 10-9
Create a Run 10-11
Copy a Run 10-11
Open a Run 10-11
Delete a Run 10-12
Offer Optimization Run Tab 10-12
Setup 10-13
Setup 10-13
Objectives and Budget 10-14
Review Information 10-15
Rules 10-16
Temporal Rules 10-17
Promotion Rules 10-21
Markdown Rules 10-22
Sell-Through Target 10-24
Price Ladder 10-26
Product Groups 10-28
Results 10-30
Possible Actions 10-30
Revenue Tile 10-31
Promotion Tile 10-32
Markdown Tile 10-34
Targeted Tile 10-36
Offers Tile 10-38
Export 10-38
Innovation Workbench 10-39
Integration with Oracle Digital Assistant 10-39

11 Size Profiles

Overview 11-1
Navigation 11-1
All Runs 11-2
Creating a Run 11-2
Information Tab 11-3
Size Ranges Tab 11-3
Parameters Tab 11-4
Reviewing a Run 11-5
Run Output 11-5
Size Profiles Tab 11-6
Input Tab 11-8
Graphics 11-8
Submitting a Run 11-9
Approved Profiles 11-9
Approved Profiles Table 11-10
Assigned Tasks 11-11
Assigned Tasks Table 11-11

12 Data Management

13 Innovation Workbench

Process 13-1
Workspace List 13-1
SQL Workshop 13-2
Application Builder 13-5
Manage Notebook 13-6
Features 13-6

14 Oracle Digital Assistance

Transactional Digital Assistance 14-1
Q&A Digital Assistance 14-2
Voice Component 14-3

Send Us Your Comments

Oracle Retail Science Cloud Services User Guide, Release 18.0.001

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

This guide describes the Oracle Retail Science Cloud Services user interface. It provides step-by-step instructions to complete most tasks that can be performed through the application.

Audience

This User Guide is intended for retailers and analysts.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info or visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs if you are hearing impaired.

Related Documents

For more information, see the following documents in the Oracle Retail Science Cloud Services documentation set:

- *Oracle Retail Science Cloud Services Administration Guide*
- *Oracle Retail Science Cloud Services Implementation Guide*
- *Oracle Retail Science Cloud Services Release Notes*
- *Oracle Retail Insights Cloud Service Suite/Oracle Retail Science Cloud Services Data Interface*

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

Improved Process for Oracle Retail Documentation Corrections

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

Oracle Retail documentation is available on the Oracle Technology Network at the following URL:

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

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

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

Oracle Retail Documentation on the Oracle Technology Network

Oracle Retail product documentation is available on the following web site:

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

(Data Model documents are not available through Oracle Technology Network. You can obtain these documents through My Oracle Support.)

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.
italic	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.

1

Getting Started

This chapter provides an overview of Oracle Retail Science Cloud Services.

About Oracle Retail Science Cloud Services

The Oracle Retail Science Cloud Services is an analytical product that consists of the following modules: Customer Decision Tree Science Cloud Service, Demand Transference Science Cloud Service, Advanced Clustering Cloud Service, Customer Segmentation, and Assortment and Space Optimization Cloud Service. You may have access to all of these modules or to a subset of them.

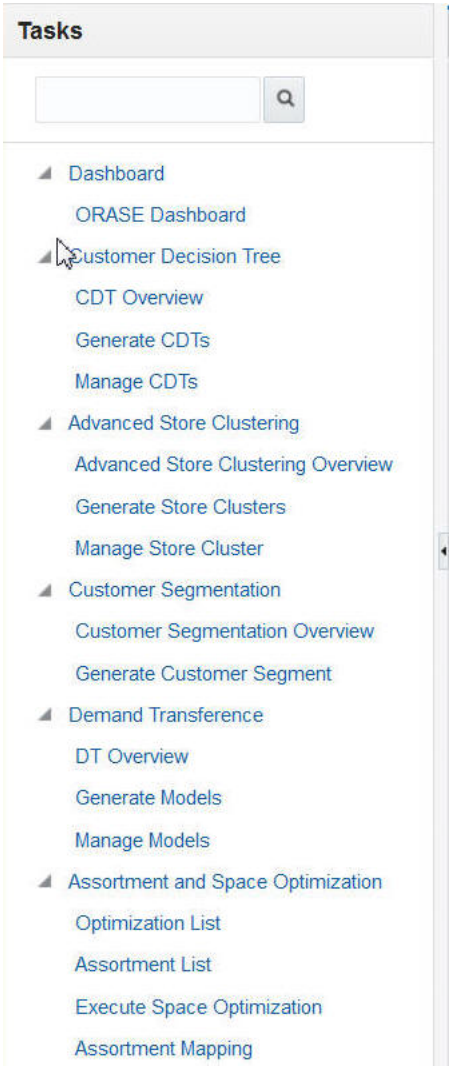
Administration

For information about the administration of Oracle Retail Science Cloud Services, see *Oracle Retail Science Cloud Services Administration Guide*.

Oracle Retail Science Dashboard

The dashboard provides access to all Oracle Retail Science modules. The main tasks associated with each module are listed in the panel on the left, as shown in Figure 1–1. Each module has its own tab.

Figure 1–1 Oracle Retail Science Tasks



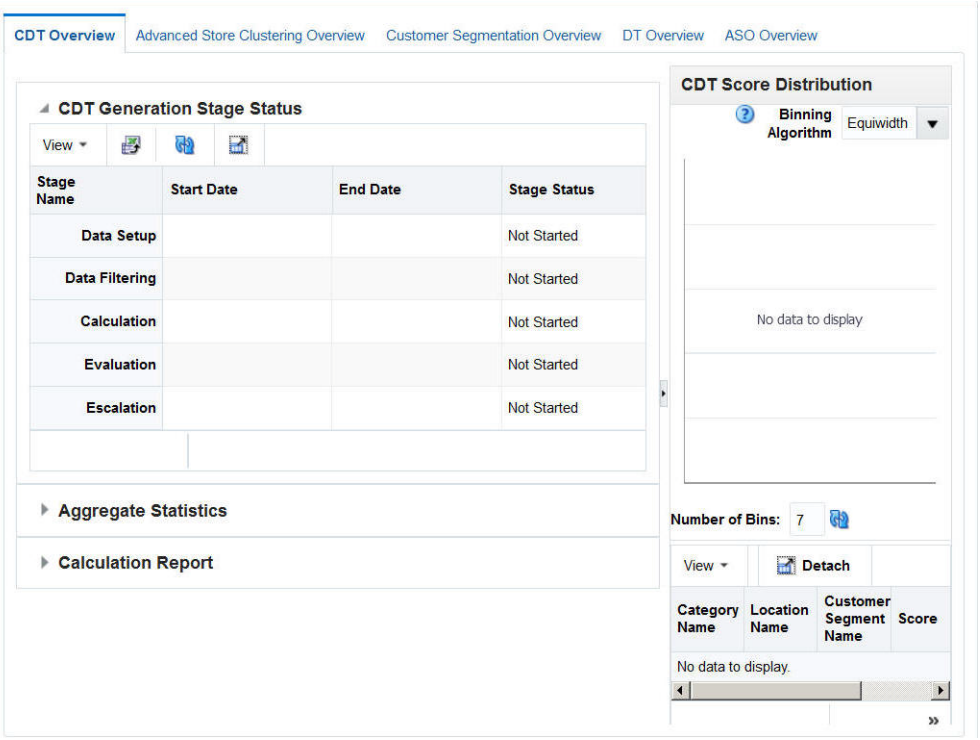
Process

This guide contains a chapter for each of the modules with detailed instructions. In general, each module contains a set of tabs, including an overview tab and tabs to use the module functionality.

Customer Decision Tree

The CDT overview tab is shown in Figure 1–2.

Figure 1–2 CDT Overview Tab



When you use the CDT Cloud Service module, you following this general iterative process to create and manage CDTs:

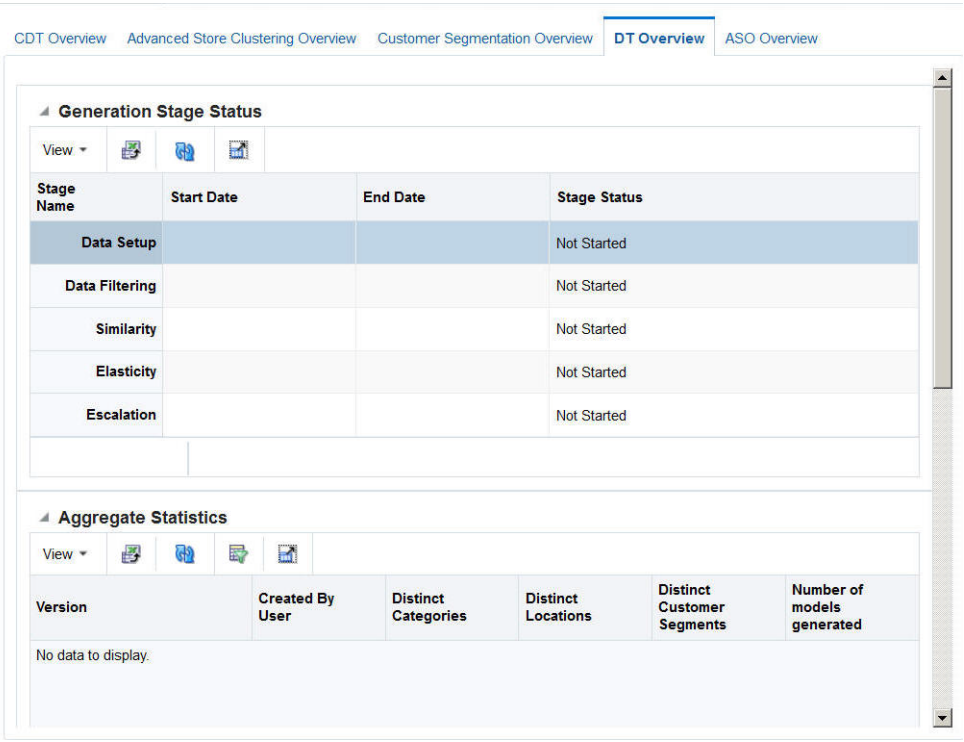
- Setup. You define the domain of calculation by selecting the categories and time frame for the calculation of the CDTs.
- Data Filtering. You configure filters that remove input data that might cause errors in the calculation or that might lead to inaccurate or unreliable answers.
- Calculation. You calculate a specific version of a CDT. Different versions of the CDT based on different configurations can be calculated and compared.
- Evaluation. You examine the results of the calculation and determine the reliability and accuracy of the answers. You can prune inaccurate or unreliable results.
- Escalation. When you are satisfied with the results, you can set the escalation path to fill in the holes for partitions whose CDTs were removed during pruning by setting up a search path through the segment hierarchy and the location hierarchy. Then, you can set a version of a CDT as complete.

Within each stage, you set the required parameters. In some stages, once you do that you click **Run** to initiate the process. To determine when a run is complete, go to the Overview tab and click the **Refresh** icon in order to see the current status of each stage. If you make changes to a stage, you must re-run that stage and all the relevant stages that follow it as the results of those stages are made invalid by the changes you made.

Demand Transference

The DT overview tab is shown in Figure 1–3.

Figure 1–3 DT Overview



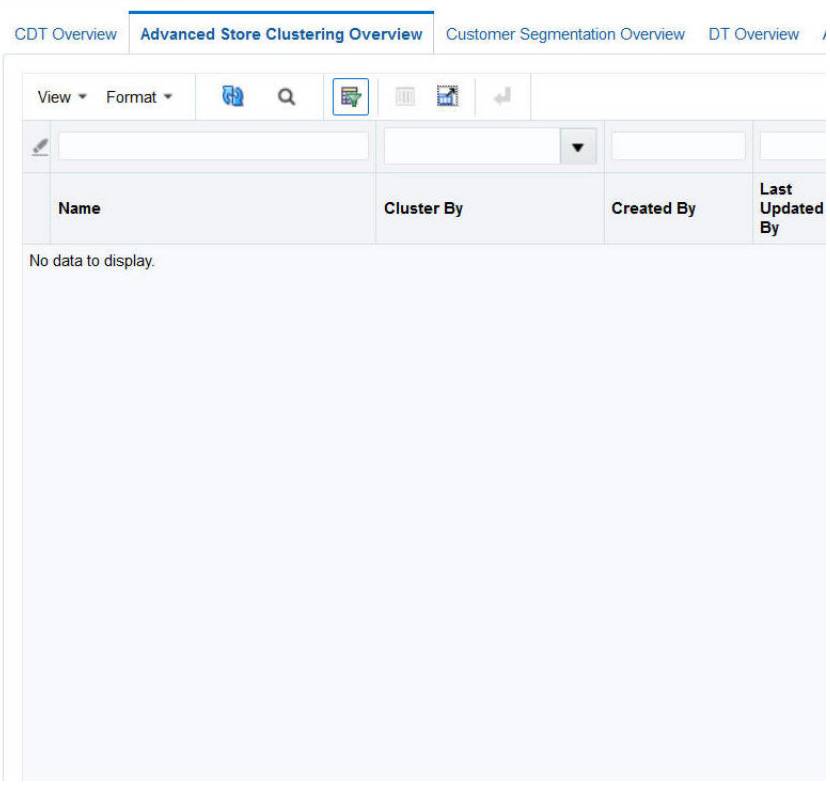
When you use the DT Cloud Service module, you follow this general iterative process to create and manage DT models:

- Data Setup. You define the categories to be used in the DT calculation.
- Data Filtering. You configure filters that remove input data that might cause errors in the calculation or that might lead to inaccurate or unreliable results.
- Similarity Calculation. You calculate similarities and assess the results of the calculation.
- Elasticity Calculation. You calculate the assortment elasticities and assess the results in terms of substitutable demand, which is the percentage of demand of a SKU that is retained when the SKU is deleted from the stores where it is selling.
- Escalation. When you are satisfied with your results, you can set the escalation path to fill in the holes for partitions whose DT models were removed during pruning by setting up a search path through the segment hierarchy and the location hierarchy. Then you can set a version of the DT model as complete.
- Manage Models. Use this tab to set time intervals for evaluating your results and to override the value for the maximum substitutable demand percentage.

Advanced Clustering

The AC overview tab is shown in Figure 1–4.

Figure 1–4 Advanced Clustering Overview Tab



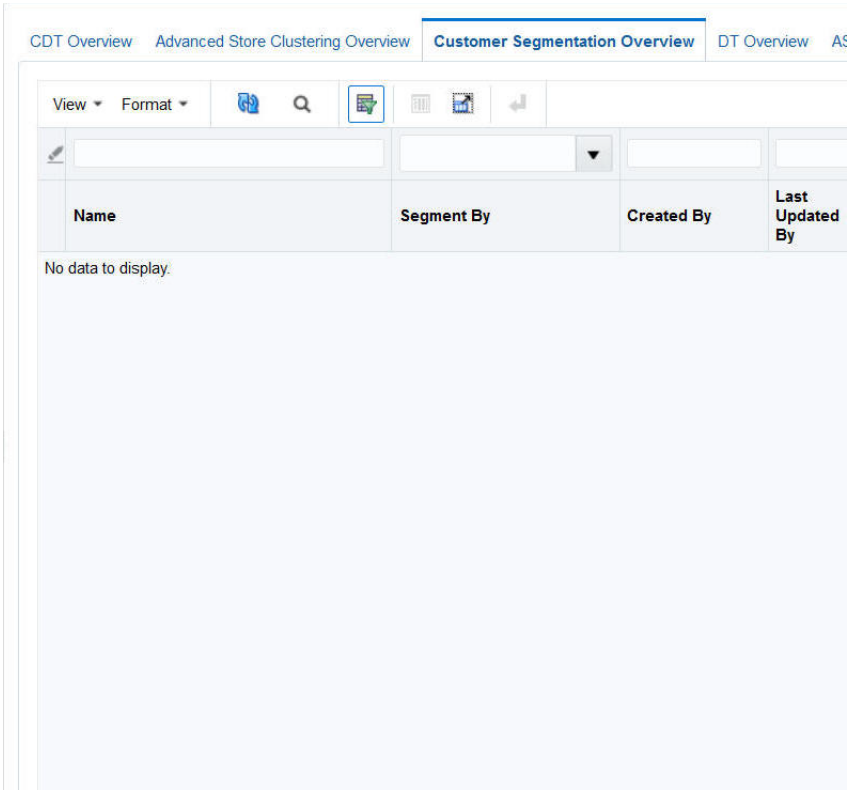
To use Advanced Clustering, follow this general process to create and manage clusters, working in the Generate Store Clusters tab and the Manage Store Clusters tab:

- Cluster Criteria. View all available clusters for specified merchandise, location, and calendar. Review cluster criteria or scenario details for each cluster. Use an existing cluster as the basis for creating a new cluster.
- Explore Data. Examine the input data for the cluster. Review multiple data points and significant attributes using the contextual information.
- Cluster Setup. Define multiple what-if scenarios. Such scenarios can be compared with one another throughout the clustering process.
- Cluster Results. View the scenario results and compare scenarios.
- Cluster Insights. Gain an understanding about cluster results and cluster performance prior to approval by examining the contextual information.
- Manage Store Clusters. Manage existing cluster criteria. Perform manual overrides and approve clusters.

Customer Segmentation

The Customer Segmentation overview tab is shown in Figure 1–5.

Figure 1–5 Customer Segmentation Overview Tab



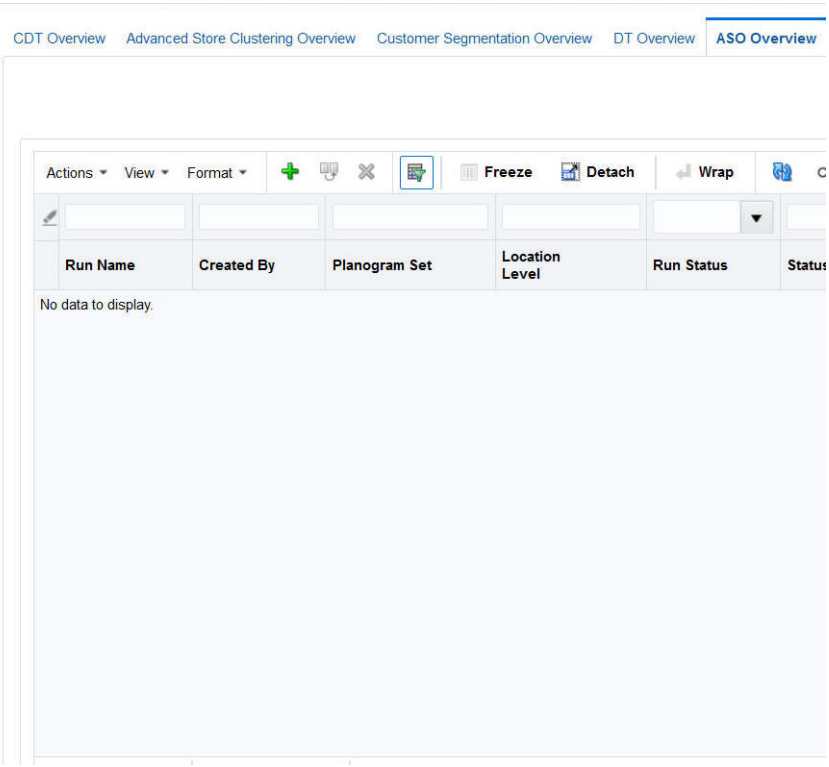
To use Customer Segmentation, follow this general process to create segments, working in the Generate Customer Segment tab:

- Segment Criteria – View all available segments for the specified merchandise, location, and calendar. Review the segment criteria or scenario details for each segment. Use an existing segment as the basis for creating a new segment.
- Explore Data – Examine attributes and their summaries for the segment. Review multiple significant attributes and their correlations using the contextual information.
- Segment Setup – Define multiple what-if scenarios. Such scenarios can be compared with one another throughout the segmentation process.
- Segment Results – View the scenario results and compare scenarios.
- Segment Insights – Gain an understanding of segment results and segment performance prior to approval by examining the contextual information.

Assortment and Space Optimization

The ASO overview tab is shown in Figure 1–6.

Figure 1–6 ASO Overview Tab



ASO is used to optimize a category manager's assortment plan by creating optimal planograms. The category manager creates a preliminary assortment and wants to determine how well that list of products can fit in stores, given the available space, product sizes, and merchandising goals, constraints, and rules. The application creates virtual planograms that organize products onto fixtures in a way that best achieves the optimization objectives. Once the manager finalizes the optimization results, the product level data can be exported for use in planning applications.

Overview of the User Interface

The user interface includes functionality described in this section.

Process Train

You can use the process train to navigate through the stages of each module. You can also use the **Back** button and the **Next** button to move through the train. The color changes for each stage once you visit that stage. Certain stages require you to run that stage before you can go to the next stage.

View Menu

The View Menu provides access to a variety of functionality that you can use to customize the display of the tables in the user interface.

Embedded Help

Embedded help, which you access by clicking the **Question Mark** icon, provides additional information about the type of details required by certain fields.

Process Indicator

At the top of the user interface, in the right-hand corner, is a process indicator that you can use to monitor the status of a user action such as clicking **Next** to go to the next stage.

Search

In certain cases, you can customize your search, using advanced search capabilities to specify the search criteria.

Icons

The following icons are used in the user interface.

Table 1–1 Icons Used in the User Interface

Icon	Icon Name	Icon Description
	Add	Add a category for CDT or DT. Add a cluster in Advanced Clustering. Create new scenarios in Advanced Clustering.
	Approve Version	Approve a version of a CDT.
	Calculate	Initiate calculation of substitutable demand percentages in DT.
	Compare Two CDTs	Look at two CDTs side-by-side in the CDT Editor.
	Complete	Indicates the CDT is ready to be activated.
	Delete an entry in a table. Delete a scenario in Clustering.	Delete.
	Detach	Detach the table from the user interface for better viewing.
	Duplicate	Make a copy of a scenario.
	Edit	Edit the category attributes (CDT and DT)
	Embedded Help	Indicates that embedded help is available for the adjacent field.
	Execute	Execute the scenario.
	Export to Excel	Export the selected data to Excel.
	Go To Top	Adjusts table display.
	Go Up	Adjusts table display.

Table 1–1 (Cont.) Icons Used in the User Interface

Icon	Icon Name	Icon Description
	No	Indicates No in Advanced Clustering.
	Query By Example	Provides access to a text entry field at the top of each column that you can use to search for data by an initial set of characters.
	Refresh	Update the table display or the stage status.
	Revert	Reverts the DT calculation.
	Save	Save the scenario.
	See Similarities	See similarities in DT Similarities Display.
	Set Version As Complete	Set a CDT version as complete.
	Select Date	Access a calendar in order to select a specific date.
	Show As Top	Adjusts table display.
	View One CDT	Access a CDT in the CDT Editor.
	Withdraw From Approved Version	Un-approve a version of a CDT.
	Yes	Acts as an indication that something exists in Advanced Clustering.

Buttons

Buttons are used for navigation and to perform certain actions.

Table 1–2 Buttons

Name	Description
Action	Provides access to Save, Approve, and Reject.
Advanced	Provides access to advanced search functionality.
Approve	Used to approve a CDT or a cluster.
Back	Used to navigate the process train.
Cancel	Cancels the action.
Complete	Used to make a CDT or DT model active.
Next	Moves to the next stage.
Reject	Used to navigate the process train.
Reset	Resets the values to the original ones.
Run	Initiates a run.

Table 1–2 (Cont.) Buttons

Name	Description
Search	Provides access to search functionality.
Stop	Stops a process.

Browser Settings

The supported browsers include Internet Explorer 11 and Google Chrome 28 or newer.

Concurrent Browser Sessions

Users should not log into more than one browser session at the same time using the same user name.

Localization

The default language for the application is English. If you are using a different language on your computer, you should adjust the language settings on your browser as appropriate.

Setting Up Internet Explorer

To configure Internet Explorer:

1. Open Internet Explorer.
2. From the **Tools** menu, select **Internet Options**.
3. From the **Internet Options** dialog box, click the **Security** tab.
4. From the **Security** tab, click **Local intranet**, or, if you have been instructed to do so by your Systems Administrator, **Trusted sites**, and then click the **Sites** button.

   If you selected Local intranet, go to step 5. If you selected Trusted sites, go to Step 6.
5. On the **Local Intranet** dialog box, click the **Advanced** button.
6. On the resulting **Local intranet** or **Trusted sites** dialog box, add the application URL if it is not already listed.

   To do so, type the application URL in the **Add this Web site to the zone** text box. Click **Add**. When the URL appears in the Web sites list, click **OK**.
7. If the **Local Intranet** dialog box from step 5 is still open, click **OK** to close it.
8. Based on the selection your made in step 4, from the **Security Tab** of the **Internet Options** dialog box, select either **Local intranet** or **Trusted sites**. Click the **Custom Level** button.
9. In case you have Pop-up Blocker enabled, add the host name from the application URL as an exception using the following steps:
   a. On the **Internet Options** dialog box, click the **Privacy** tab.
   b. On the **Privacy** tab, in the **Pop-up Blocker** section, click **Settings**.
   c. On the **Pop-up Blocker Settings** dialog box, enter the host name in the **Address of website to allow** field, and click **Add**.
   d. Click **Close**.
10. On the Internet Options dialog box, click **OK** to return to the browser.

Supported Characters for Text Entry

The following characters are valid for text input: all letters, all numbers, and the following characters: '_', '#', '%', '*', '$', ' ', ',' & '-'

Histograms

Certain stages have associated histograms that can help you analyze the data presented in that stage. You can adjust the way the histogram presents the data in two ways. You can select the number of bins that are used to display the data. In addition, you can select how the bins are defined: Equiwidth or Custom. Each of these options uses a specific algorithm to determine how the bins are defined.

The Equiwidth approach takes the minimum and maximum values in a set of numbers and divides that range into equally sized bins. For example, using the numbers from 1 to 100 with 10 bins, the histogram shows bins for 1-10, 11-20, and so on. If specific bins have no value represented (for example, if all the values are in the range of 1-10 and 91-100), then the histogram will not show that bin in the UI. Additionally, the histogram data series ranges are shown using the actually minimum and maximum values for each of the bins. So rather than showing a range of 1-10, if the only value available was a 5, then the range for the first bin would appear as 5-5.

In the Custom approach, each of the bins has an equal number of values represented, while the minimum and maximum number associated with the bin is adjusted. However, the bins are defined using distinct values instead of all the available values. The bins may or may not be of equal height, depending on how diverse the numbers are.

The two approaches differ in that the Custom approach only shows fewer bins than requested if there are fewer distinct values than what was requested for the number of bins.

The two approaches are similar in that both handle the Min/Max value display in a similar manner, using actual data values that are associated with the bin.

To determine which approach to use, you should consider what type of data you are trying to see and the amount of detail you want. For example, if you are trying to set a data filter value, and you want to do so using a common value, you may be able to see where the majority of the data falls using one of the algorithms, while the other algorithm may help you pinpoint a specific value within the range. The Equiwidth approach is negatively affected by values that are at the extreme ends of a value being binned. This can cause the majority of the data values to appear in a single bin. The Custom approach puts a greater emphasis on a value that is repeatedly found in a dataset. Depending on the values being charted, you may find that one of the approaches presents better data than the other approach.

2

Customer Decision Trees

This chapter describes the use of the Customer Decision Tree Science Cloud Service module.

Introduction

A customer decision tree is a decision support tool that uses a tree-like graph to model a customer hierarchical decision-making process for a specific product. The branches of the tree provide a visual representation of the choices a customer makes and the order of importance of various product characteristics. Transaction data is used in the analysis.

A customer decision tree identifies the decisions a customer makes when choosing a particular product. The decision tree is produced by algorithms that analyze historical customer sales data. It illustrates how customers shop and how they evaluate the importance of different product attributes when making buying decisions. Such information can be useful to a retailer in terms of product selection and display.

The CDT Cloud Service module consists of three tabs: Overview, Generate CDTs, and Manage CDTs. You use the Overview tab to keep track of the status of each stage within the Generate CDTs tab. You use the Manage CDTs tab to assess the set of CDTs you and other users have created using the Generate CDTs tab, setting versions as complete, approving versions, and deleting versions, as necessary.

Overview of CDT Process

When you use the CDT Cloud Service module, you follow this general iterative process to create and manage CDTs:

- Setup. You define the domain of calculation by selecting the categories and time frame for the calculation of the CDTs.
- Data Filtering. You configure filters that remove input data that might cause errors in the calculation or that might lead to inaccurate or unreliable answers.
- Calculation. You calculate a specific version of a CDT. Different versions of the CDT based on different configurations can be calculated and compared.
- Evaluation. You examine the results of the calculation and determine the reliability and accuracy of the answers. You can prune inaccurate or unreliable results.
- Escalation. When you are satisfied with the results, you can set the escalation path to fill in the holes for partitions whose CDTs were removed during pruning by setting up a search path through the segment hierarchy and the location hierarchy. Then, you can set a version of a CDT as complete.

Within each stage, you set the required parameters. In some stages, once you do that you click **Run** to initiate the process. To determine when a run is complete, go to the Overview tab and click the **Refresh** icon in order to see the current status of each stage. If you make changes to a stage, you must re-run that stage and all the relevant stages that follow it as the results of those stages are made invalid by the changes you made.

Overview Tab

The Overview tab displays information that you can view and use to monitor the progress of the CDT stages as well as to view some aggregate statistics and the CDT results from the last successful run.

This tab contains the following sections:

- CDT Generation Stage Status
- Aggregate Statistics
- CDT Calculation Report
- CDT Score Distribution

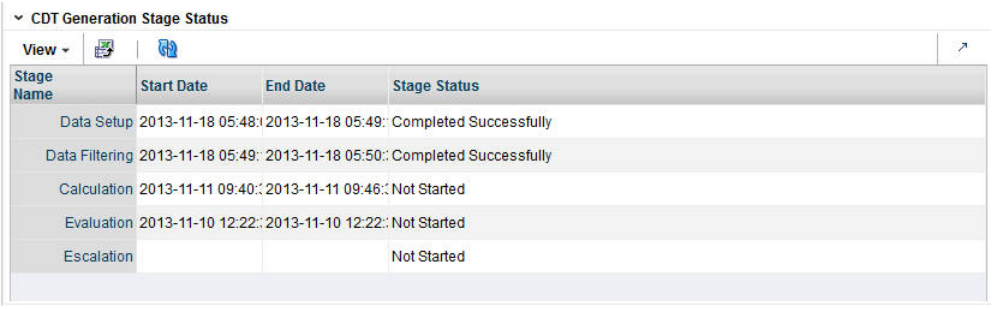
CDT Generation Stage Status

The CDT Generation Stage Status table displays the status of each of the CDT stages.

Click the **Refresh** icon to update the fields and see the latest status for each stage.

In cases where the status is either Failed or Stopped/Paused, or the run never completes, you should consult the database logs in RSE_LOG_MSG as well as the WebLogic console logs in your WebLogic domain in order to troubleshoot the problem.

Figure 2–1 CDT Generation Stage Status



You can use the status information to monitor the progress of each stage. It contains the following fields, which can be arranged and viewed, but not modified.

Table 2–1 CDT Generation Stage Status Fields

Field Name	Description
Stage Name	A row exists in the table for each CDT stage that provides detailed status information. The five stages are Setup, Data Filtering, Calculation, Evaluation, and Escalation.
Start Date	The date and time when a run for the stage most recently started.
End Date	The date and time when a run for the stage most recently ended.

Table 2–1 (Cont.) CDT Generation Stage Status Fields

Field Name	Description
Stage Status	The current status of the stage: Not Started, Not Started (Scheduled for Later), Processing, Completed Successfully, Completed with Errors, Stopped/Paused, Cancelled, or Failed.

The following provides an explanation of the different values for the status of a stage.

Table 2–2 Status Explanations

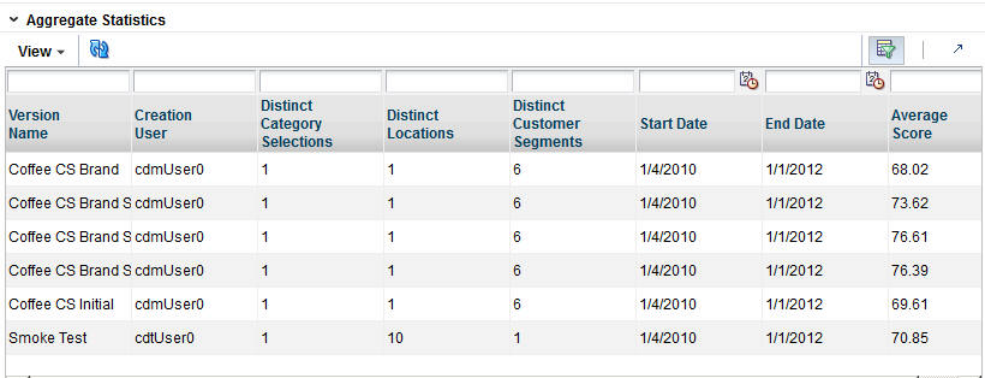
Stage Status	Description
Not Started	This indicates that the stage has not yet been initiated. This can occur when you first begin to create a CDT or when you update an existing version.
Not Started (Scheduled for Later)	The stage is scheduled to start after the previous stage is complete.
Processing	This indicates that the stage is currently being processed.
Stopped/Paused	This indicates either that the stage has been stopped by the application or that the user has chosen to stop the process (for example, to make a change to an option). In the later case, once any background processing has stopped, the user can re-run the stopped stage.
Completed Successfully	This indicates that the stage has been successfully processed.
Completed with Errors	This indicates that the stage was able to complete the processing of all requested CDTs, but one or more CDTs encountered errors during the run and were not able to complete successfully. This most commonly occurs when the data used for a CDT is not available or is too sparse to produce a result.
Failed	This indicates that a problem occurred during the processing. If the failure occurred during the Calculation stage, and only a few CDT results did not produce a CDT score, then the user can continue without those results since pruning will remove them. If the failure occurred during any of the other stages, the user must re-run the stage.
Cancelled	This indicates that the database has cancelled the execution of the stage either because of missing data or an exception. The user should review the RSE_LOG_MSG log to determine the problem.

Aggregate Statistics

The Aggregate Statistics table displays statistical details about the existing CDT versions.

Click the **Refresh** icon to update the fields and see the latest information in this table.

Figure 2–2 CDT Aggregate Statistics



Each existing version has a row in the table. The table contains the following fields, which can be arranged and viewed, but not modified.

Table 2–3 Aggregate Statistics Fields

Field Name	Description
Version Name	User-created name that identifies this CDT version. The name must be unique to the user.
Creation User	The login name of the person who created this version.
Distinct Category Selections	The number of categories associated with this version.
Distinct Locations	The number of locations associated with this version.
Distinct Customer Segments	The number of customer segments associated with this version.
Start Date	The beginning date associated with the CDTs in this version.
End Date	The end date associated with the CDTs in this version.
Average Score	The average of the CDT scores for a version. A CDT score ranges from 0 to 100. A higher value indicates a better score.

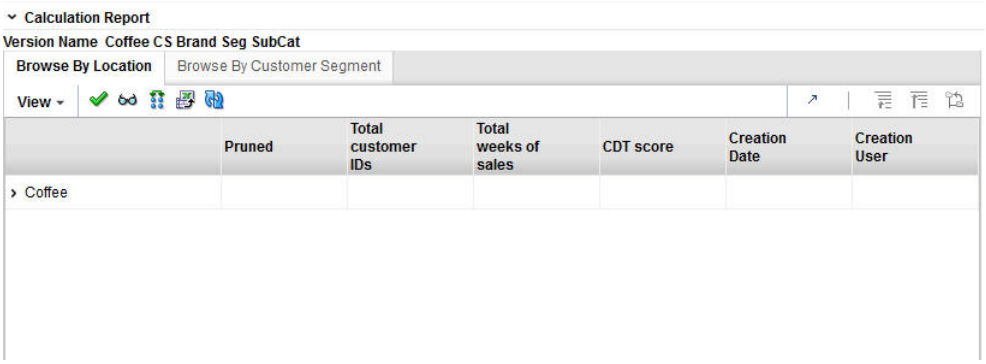
Calculation Report

The Calculation Report displays the CDT results for the indicated version from the last successful run, if one has occurred.

Here you can review information by Location or by Customer Segment. You can also set a version as complete, view and edit a CDT, or compare two CDTs.

Click the **Refresh** icon to update the fields and see the latest information in this table.

Figure 2–3 CDT Calculate Report



The Calculation Report contains the following fields:

Table 2–4 Calculation Report Fields

Field Name	Description
Partition	A partition is the combination of category, segment, and location. The column identifies the names of all nodes in the tree structure in the table. The node type may be either category, location, or customer segment.
Pruned	The number of CDTs removed from the list of usable CDTs. CDTs are removed that do not meet the filtering thresholds. During the Escalation stage, the escalation process makes adjustments for the CDTs that are removed.
Total Customer IDs	The number of customers used in the calculation of the CDT.
Total Weeks of Sales	The total number of weeks of sales used in the calculation of the CDT. This provides an indication of the amount of sales data used in the calculation. All the CDTs in a given version should have a similar value; if this is not the case, the results should be evaluated.
CDT Score	A confidence score assigned to help in assessing the quality of the CDT. A CDT score ranges from 0 to 100. A higher value indicates a better score.
Creation Date	The date when the version whose data is displayed was created.
Creation User	The login name of the user who created the version.

The following icons are available for the Calculation Report:

Set Version as Complete

Each partition must have an active CDT. The active CDT is used by other applications that require a CDT.

All CDTs within a version are activated when you click **Set Version as Complete**.

If a CDT exists in more than one version, then the most recently activated version takes precedence.

If a version contains multiple categories and is Complete and then a different version with a partial overlap of categories is later marked as Complete, only the overlapping categories are replaced in the new version.

If a version is overwritten, then the CDTs in that version are no longer active.

When a CDT is at risk of going from active to inactive, you will see a warning message.

To determine whether or not a CDT is active, go to the Manage CDTs tab. The Browse by Categories or by Versions table displays a flag that indicates whether or not a CDT is active.

View One CDT

This icon provides access to the CDT Editor. For details about this functionality, see Using the CDT Editor.

Compare Two CDTs

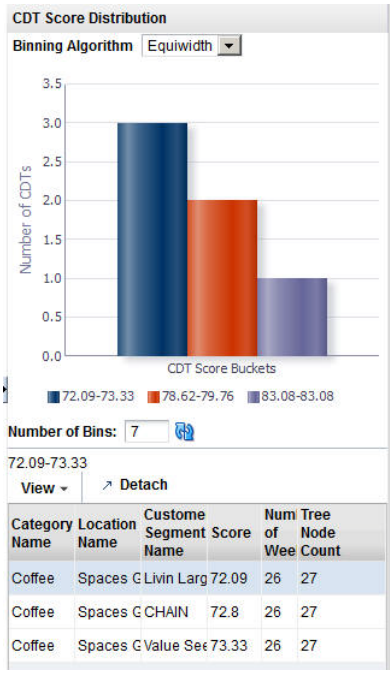
This icon provides access to CDT Compare. For details about this functionality, see Comparing Two CDTs

CDT Score Distribution Histogram

The CDT Score Distribution histogram displays, for the current calculation and the current user, the range of CDT scores and their frequencies. CDT scores range from 0 to 100; a higher number indicates a better score. You can use the information in this histogram to assess the quality of your CDTs.

For detailed information about using the histograms, see Chapter 1, "Getting Started.".

Figure 2–4 CDT Score Distribution



Below the histogram you see a table that provides details about the histogram. This information can help you assess the quality of a CDT. Click on a specific bin in the histogram to populate the table with information about that bin. You see a list of all the specific partitions for the bin, along with the score for each, the number of weeks of data used to calculate the score, and the number of tree nodes for each partition for the selected bin.

Table 2–5 CDT Score Distribution Fields

Field Name	Description
Category Name	The category name of the partition for the specific CDT score.
Location Name	The location name of the partition for the specific CDT score.
Customer Segment Name	The customer segment name of the partition for the specific CDT score.
Score	The CDT score for the defined partition. A CDT score ranges from 0 to 100. A higher value indicates a better score.
Number of Weeks	The number of weeks of data used to calculate the CDT score.
Tree Node Count	The number of tree nodes present in the CDT tree structure used to calculate the CDT score.

Generate CDT Tab

The Generate CDTs tab is used to configure, run, evaluate, modify, and deploy a CDT. The process is divided into five stages that must be run in order. You can return to a stage you have already completed and make changes, but if you do, you must rerun that stage and all the stages that follow that stage, as the calculations are invalidated by the modifications you just made to the settings in the stage you changed.

To monitor the progress of any stage, go to the Overview tab and click the **Refresh** icon.

The five CDT stages are:

Table 2–6 Generate CDTs Tab: Stages

Stage Name	Description
Data Setup	Select CMPO groups and define the time intervals for the data to be used in the calculation.
Data Filtering	Filter out input data that may result in inaccurate or unreliable answers.
Calculation	Calculate the CDT.
Evaluation	Assess the reliability and accuracy of the results of the calculation. Prune unreliable or inaccurate answers.
Escalation	Set the escalation path for the CDT. Use the Escalation Report to evaluate the results of the escalation.

Data Setup

The Data Setup stage provides two sections to configure: Category Selections and Time Interval Setup. In this stage, you select the category or categories you want to calculate CDTs for and specify the time interval for the calculation.

Process

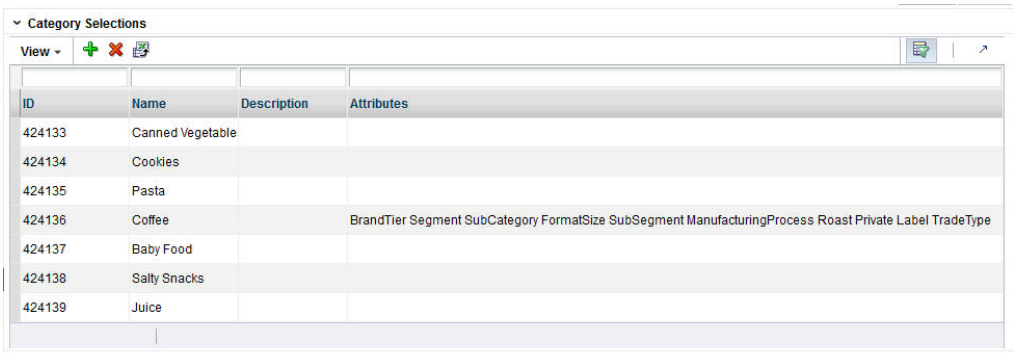
Here is the high-level process for setting up the data for CDT.

1. Select the category or categories you want to calculate CDTs for.
2. Set up the time intervals for the CDT calculation.
3. Click **Next** to go to the Data Filtering stage.

Category Selections

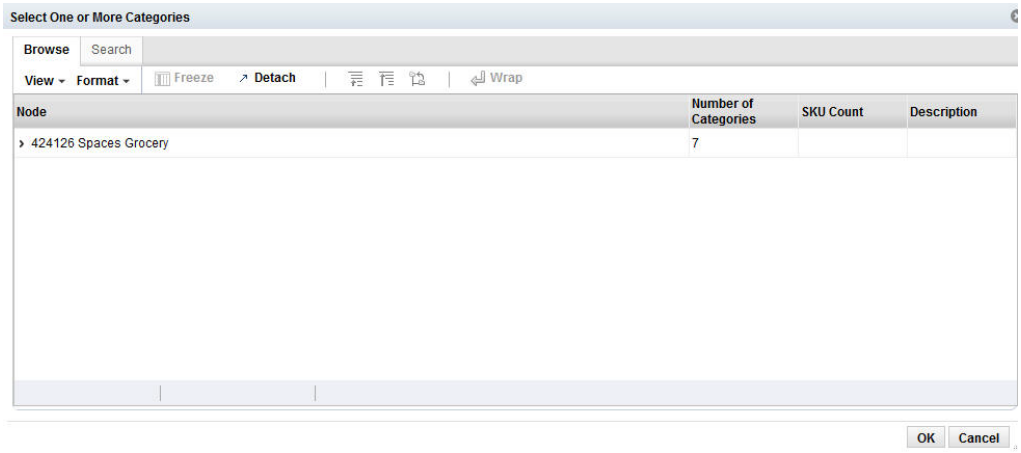
Use this table to add categories you want to include in the CDT calculation or delete categories that you want to remove from the CDT calculation.

Figure 2–5 CDT Category Selections



To display a list of available nodes and the categories included in those nodes, click the **Add** icon. You see the Select One or More Categories dialog box, which contains two tabs, Browse and Search. You can use either tab to find the categories you are looking for.

Figure 2–6 CDT Categories Browse Tab



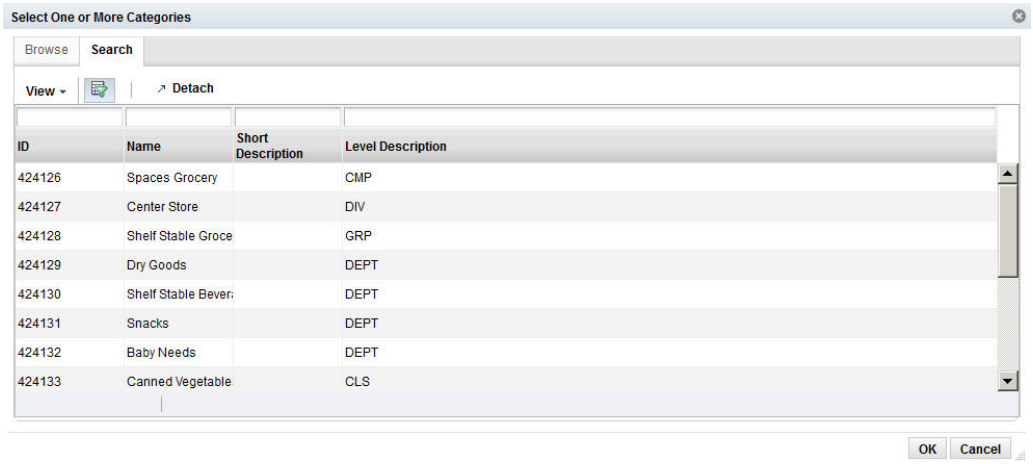
The Browse tab displays a table with the following fields:

Table 2–7 Category Selections: Browse

Field Name	Description
Partition	The node tree structure can be expanded in order to view its categories.
Number of Categories	The number of categories within the node that has been selected. The number can help you understand the amount of processing required for the calculation.
SKU Count	The number of SKUs in a category. A category with too few SKUs may not produce good CDTs.
Description	A description that provides additional information about the category.

Select the category or categories within the node that you want to be part of CDT and click **OK**.

Figure 2–7 CDT Categories Search Tab



The Search tab displays a table with the following fields:

Table 2–8 Category Selections: Search

Field Name	Field Description
ID	An external code used to identify the category in other systems such as CMPO.
Name	The category name.
Short Description	A description that provides additional information about the category.
Level Description	A description of the level of the merchandise hierarchy that the node belongs to.

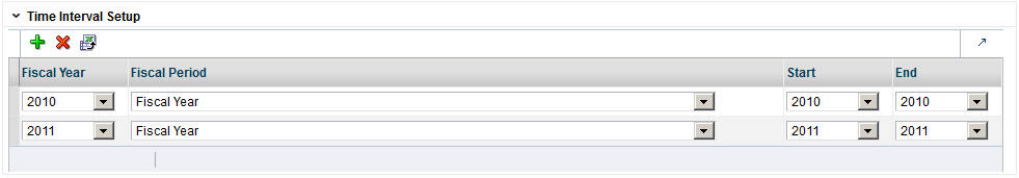
Select the category or categories within the node that you want to be part of the CDT and click **OK**.

Your selections are displayed in the Category Selections table.

Time Interval Setup

The time interval defines the time span for the historical data that is used to calculate the CDT. Only the specified weeks of sales data are used in the calculation. A group of intervals can be defined. Gaps between intervals are permitted; however, intervals cannot overlap. A six-month period is recommended.

Figure 2–8 CDT Time Interval Setup



To define a time interval for the CDT to be generated, click the **Add** icon. A row displays in the dialog box. Select from the following drop-down menus in order to define the time interval:

Table 2–9 Time Interval: Menus

Menu Name	Description
Fiscal Year	The fiscal year for the time interval.
Fiscal Period	The fiscal period within the fiscal year (Fiscal Quarter, Fiscal Period, or Fiscal Week).
Start	This defines when the time interval specified in Fiscal Period begins.
End	This defines when the time interval specified in Fiscal Period ends.

After selecting the category and specifying the time interval, click **Next** to go to the Data Filtering stage.

Data Filtering

The Data Filtering stage applies to all the categories and time intervals that you select in the Data Setup stage.

Process

Here is the high-level process for setting up and running data filtering.

1. Enter the appropriate values into the Filter Setup text entry boxes.
2. Click **Run** in order to filter the data.
3. Review the filtering results in the Data Filtering Summary table and the Data Filtering histograms.
4. After reviewing the results, if necessary, make changes to the values for the filters in Filter Setup and re-run the stage.
5. When you are satisfied with the results, click **Next** to go to the Calculation stage.

Filter Setup

You configure the following filters in order to filter out data you consider unacceptable for the calculation of the CDT. Note that the two attribute filters listed in Table 2–10 are stored and used during the Calculation stage. You also set additional data filtering parameters in the Calculation stage.

Figure 2–9 CDT Filter Setup

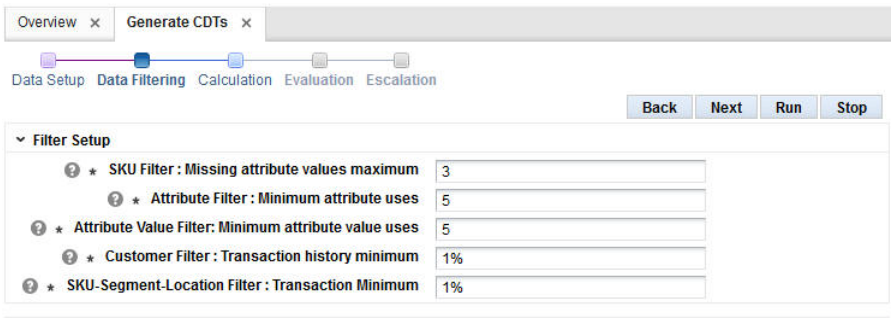


Table 2–10 Data Filters

Data Filter Name	Data Filter Description
SKU Filter: Missing attribute values maximum	Each SKU is defined by its attribute values. If a certain absolute value for the attribute values is not defined, then the product definition is not accurate. A SKU with too many missing attribute values should be filtered out. The default value is 25 for the total attribute values (that is, a SKU with greater than 25 missing attribute values is not included in the calculation of the CDT).
Attribute Filter: Minimum attribute uses	An attribute that is used by only a few SKUs should be filtered out. The default value is an absolute value of 5 for the total SKUs in the category (that is, the data for an attribute that is used by fewer than five of the SKUs is not included in the calculation of the CDT). This filter does not remove t-log level data, but removes the attribute and attribute values from the CDT creation process.
Attribute Value Filter: Minimum attribute value uses	An attribute value that is used by only a few SKUs should be filtered out. The default value is an absolute value of 5 for the SKUs in a category (that is, the data for an attribute value that is used by fewer than five of the SKUs is not included in the calculation of the CDT). This filter does not remove t-log level data, but removes the attribute and attribute values from the CDT creation process.
Customer Filter: Transaction history minimum	Customers with short transaction histories are considered outliers. You assign a percentage value that is applied to the median number of transactions for all customers. Such customers are filtered out. The default value is 10% (that is, a customer who has fewer than 10% of the median number of transactions for all customers is not included in the calculation of the CDT).
SKU-Segment-Location Filter: Transaction minimum	SKUs that have few transactions for a given location-segment partition are considered outliers. You assign a percentage value that is applied to the median number of transactions for the SKUs in a specific partition. Such transactions are filtered out. The default value is 10% (that is, a SKU that is involved in fewer that 10% of the median number of transactions for a specific partition is not included in the calculation of the CDT).

Once you have configured the filters, click **Run** to start the filtering process.

Data Filtering Summary

The following information is provided after the filtering is complete and quantifies the amount of data filtered out for the three indicated filters for sales units, sales amounts, transaction counts, SKU counts, and customer counts. Use this information to assess the effects of the pruning.

Figure 2–10 CDT Data Filtering Summary

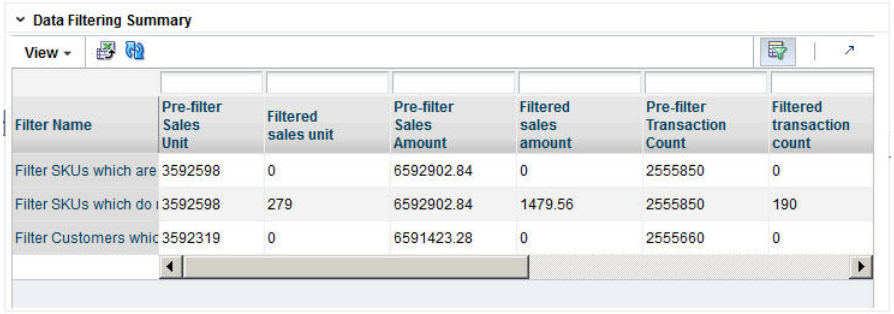


Table 2–11 Filter Data Summary Fields

Field Name	Field Description
Filter Name	The following filter names are listed: Filter SKUs that are missing too many attribute values. This maps to the SKU Filter: Missing attribute values maximum in the Filter Setup. Filter SKUs that do not have enough sale transaction history. This maps to the Customer Filter: Transaction history minimum in the Filter Setup. Filter Customers that do not have typical transaction history. this maps to the SKU-Segment-Location Filter: Transaction minimum in the Filter Setup.
Pre-filter Sales Unit	Amount prior to the application of the filter.
Filtered Sales Unit	Amount remaining after the application of the filter.
Pre-filter Sales Amount	Amount prior to the application of the filter.
Filtered Sales Amount	Amount remaining after the application of the filter.
Pre-filter Transaction Count	Amount prior to the application of the filter.
Filtered Transaction Count	Amount remaining after the application of the filter.
Pre-filter SKU Count	Amount prior to the application of the filter.
Filtered SKU Count	Amount remaining after the application of the filter.
Pre-filter Customer Count	Amount prior to the application of the filter.
Filtered Customer Count	Amount remaining after the application of the filter.

Data Filtering Histograms

The following histograms illustrate the effects of filtering. You can use the information displayed in the histograms to adjust the configuration of the filters in order to eliminate outlier data. If you modify the filters, you must re-run the stage.

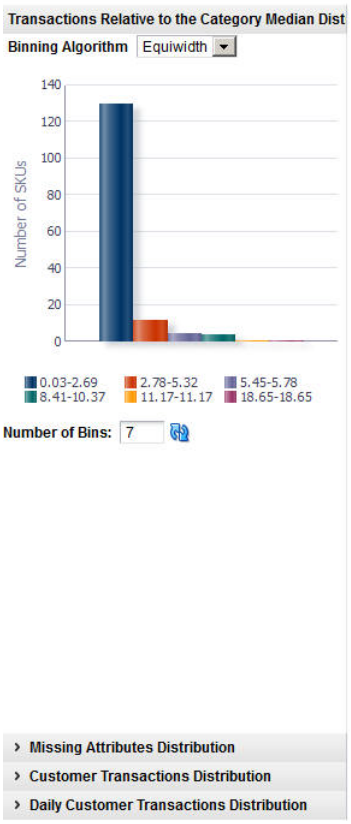
Table 2–12 Filter Histograms

Histogram Name	Description
Transactions Relative to the Category Median Distribution	Displays the number of transactions relative to the category median for a SKU's category.
Missing Attribute Distribution	Displays the number of attributes that have missing values per number of SKUs.

Table 2–12 (Cont.) Filter Histograms

Histogram Name	Description
Customer Transaction Distribution	Displays the number of transactions in historical data per number of customers.
Daily Customer Transaction Distribution	Displays the relative number of transactions per day per number of customers.

Figure 2–11 CDT Transactions Relative Histogram



Once you are satisfied with the pruning results, click **Next** to go to the Calculation stage.

Calculation

The Calculation stage creates the CDTs for all of the partitions that you selected in the Data Setup stage. A separate CDT calculation occurs for each partition. For example, if you selected two CMPO nodes, and the system also has two customer segments and three locations, then that is a total of 2 x 2 x 3 = 12 CDTs at the lowest level. In addition, there are CDTs at the higher levels (for example, above location).

Process

Here is the high-level process for calculating CDTs.

1. Enter a unique name for the version of the CDTs to be calculated.
2. Enter values for lowest tree level and minimum percentage of SKUs.
3. Use the check boxes to indicate whether or not only the top level processing calculation should occur for location or customer segment.

4. Customize the ranking of the category attributes if necessary.
5. Click **Run** to start the calculation.
6. Review the calculation results in the Calculation Report and the CDT Score Distribution histogram.
7. After reviewing the results, if necessary, make changes to the values for the calculation in Version Setup and re-run the stage.
8. When you are satisfied with the results, click **Next** to go to the Evaluation stage.

Version Setup

At the top of this stage you see three text boxes and two check boxes that you use to configure the levels at which the calculation occurs.

Figure 2–12 CDT Version Setup

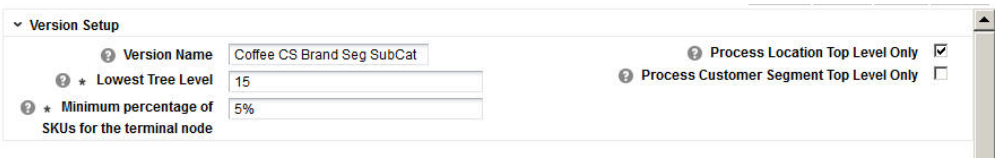


Table 2–13 Calculation Stage: Fields

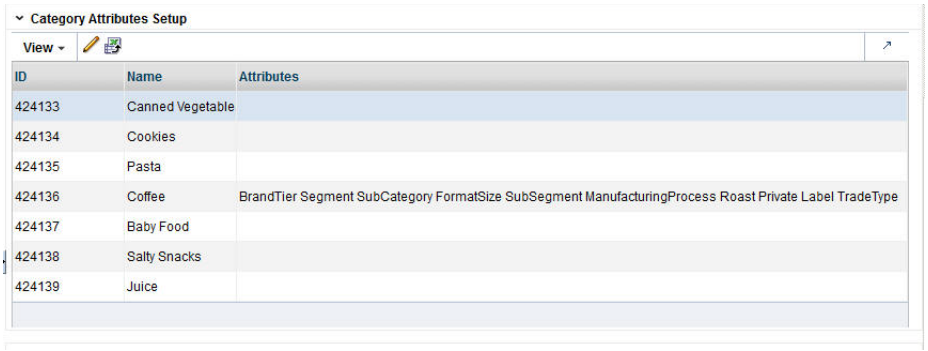
Field Name	Field Description
Version Name	Assign a name to each version of a CDT calculation. This allows you to create and save more than one version of a CDT. The version name you assign here is used in the Calculation Report, Aggregate Statistics table, and in the Manage CDTs tab. Version names can be re-used; however, if the version name in question has active CDTs, then you will see a warning that the active CDTS will be removed from the version if you do re-use the version name.
Lowest Tree Level	Use this option to define the maximum number of levels of CDTs to be calculated. The default value is 15. The number of levels in a CDT can at most be equal to the number of attributes used by the SKUs in the category plus 1. (The number 1 is used for the top level of the tree.) When determining the maximum number of levels, you should consider how many attributes should be represented in the tree. (For example, if you only want to see the top six attributes in the CDT, then set the lowest level to 7.)
Minimum Percentage of SKUs for the Terminal Node	Use this option to make sure the terminal nodes have a sufficient number of SKUs. For example, if the current CDT partition contains 100 SKUs, then any one node in the tree containing fewer than 5 SKUs will be considered a terminal node and the tree will not expand further along that branch of the tree.
Process Location Top Level Only	Check this option if you want CDTs to be calculated for the Location Chain *only*. You can select this option in order to decrease the amount of time it takes the system to perform the calculation.
Process Customer Segment Top Level Only	Check this option if you want CDTs to be calculated for the Customer Segment Chain *only*. You can select this option in order to decrease the amount of time it takes the system to perform the calculation.

Category Attributes Setup

The application determines a specific ranking order for the category attributes. You can optionally change this order and create your own ranking of the attributes for a category by ordering the attributes from most important to least important. The Category Attributes Setup table lists the categories you are calculating CDTs for. Use this list to select a category to edit.

Click the **Edit** icon to access this functionality. You can adjust the ranking for all the attributes or a subset of the attributes.

Figure 2–13 CDT Category Attributes Setup



The Category Attributes Setup dialog box contains the following fields. For each attribute you want to rank, select the Top Attribute check box and enter a value for the Attribute Rank. Click **OK**.

Figure 2–14 CDT Category Attributes Setup

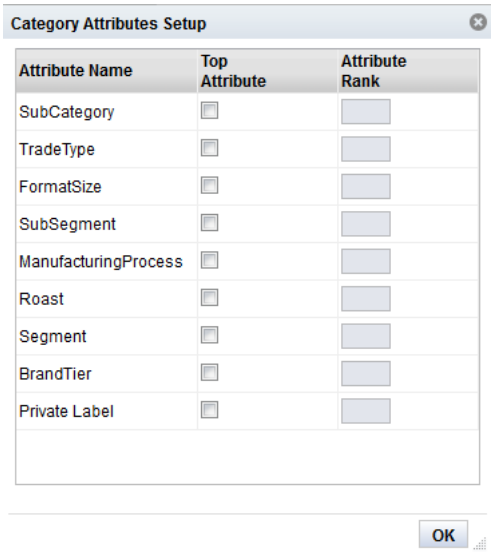


Table 2–14 Category Attributes Setup Dialog Box

Field	Description
Attribute Name	Identifies the attribute.

Table 2–14 (Cont.) Category Attributes Setup Dialog Box

Field	Description
Top Attribute	Check to enable ranking. This indicates that this attribute must be one of the first attributes in the resulting tree. This box is checked automatically for functional fit attributes and cannot be deselected. A functional fit attribute is one for which a customer cannot choose a different product with a different attribute value. For example, windshield wiper blades must be selected according to the car manufacturer's requirements. The numerical rank is optional. If you do not enter rank values, then the system will determine the order of all the top attributes first and then put any remaining attributes below the top attributes in the tree.
Attribute Rank	Enter a value for each attribute, starting with the value 1, to indicate the relative rank of the attribute. A value can be entered only if the attribute has been identified as a top attribute. You do not have to specify the rank if Top Attribute is checked. If multiple attributes are checked, then the system rank orders them.

Once you have configured the parameters for the calculation, click **Run** to start the process.

Calculation Report

The Calculation Report displays the CDT results for the indicated version from the last successful run (if one has occurred). Here you can review information by Location or by Customer Segment. You can also set a version as active, view and edit a CDT, or compare two CDTs.

Figure 2–15 CDT Calculation Report



Table 2–15 Calculation Report Fields

Field Name	Description
Partition	A partition is the combination of category, segment, and location. The column identifies the names of all nodes in the tree structure in the table. The node type may be either category, location, or customer segment.
Pruned	The number of CDTs removed from the list of usable CDTs. CDTs are removed that do not meet the filtering thresholds. During the Escalation stage, the escalation process makes adjustments for the CDTs that are removed.
Total Customer IDs	The number of customers used in the calculation of the CDT.

Table 2–15 (Cont.) Calculation Report Fields

Field Name	Description
Total Weeks of Sales	The total number of weeks of sales used in the calculation of the CDT. This provides an indication of the amount of sales data used in the calculation. All the CDTs in a given version should have a similar value; if this is not the case, the results should be evaluated.
CDT Score	A confidence score assigned to help in assessing the quality of the CDT. A CDT score ranges from 0 to 100. A higher value indicates a better score.
Creation Date	The date when the version whose data is displayed was created.
Creation User	The login name of the user who created the version.

The following icons are available for the Calculation Report:

Set Version as Complete

Each partition must have an active CDT. The active CDT is used by other applications that require a CDT.

All CDTs within a version are activated when you click **Set Version as Complete**.

If a CDT exists in more than one version, then the most recently activated version takes precedence.

If a version contains multiple categories and is Complete and then a different version with a partial overlap of categories is later marked as Complete, only the overlapping categories are replaced in the new version.

If a version is overwritten, then the CDTs in that version are no longer active.

When a CDT is at risk of going from active to inactive, you will see a warning message.

To determine whether or not a CDT is active, go to the Manage CDTs tab. The Browse by Categories or by Versions table displays a flag that indicates whether or not a CDT is active.

View One CDT

This icon provides access to the CDT Editor. For details about this functionality, see Using the CDT Editor.

Compare Two CDTs

This icon provides access to CDT Compare. For details about this functionality, see Comparing Two CDTs

CDT Score Distribution Histogram

The CDT Score Distribution histogram displays, for the current calculation, the range of CDT scores and their frequencies. CDT scores range from 0 to 100; a higher number indicates a better score. You can use the information in this histogram to assess the quality of your CDTs.

Below the histogram you see a table that provides details about the histogram. This information can help you assess the quality of a CDT. Click on a specific bin in the histogram to populate the table with information about that bin. You see a list of all the specific category name/location name/customer segment name partitions for the bin,

along with the score for each, the number of weeks of data used to calculate the score, and the number of tree nodes for each partition.

Table 2–16 CDT Score Distribution Fields

Field Name	Description
Category Name	The category name of the partition for the specific CDT score.
Location Name	The location name of the partition for the specific CDT score.
Customer Segment Name	The customer segment name of the partition for the specific CDT score.
Score	A confidence score assigned to help in assessing the quality of the CDT. A CDT score ranges from 0 to 100. A higher value indicates a better score.
Number of Weeks	The number of weeks of data used to calculate the CDT score.
Tree Node Count	The number of tree nodes used to calculate the CDT score.

When you have finished with the calculation process, click **Next** to go to the Evaluation stage.

Evaluation

Once the CDT has been generated, you can evaluate it and make adjustments to parameters used in the CDT calculation during the Evaluation stage. Only the current version of the CDT, the one just generated in the Calculation stage, can be assessed in the Evaluation stage.

Process

Here is the high-level process for using the Evaluation stage to prune the results of the CDT calculation.

1. If you have determined that you need to prune the results of the calculation, enter appropriate values into the text boxes in the Pruning Setup area.
2. Click **Run** to start the pruning process.
3. Review the pruning results in the Pruning Results section and in the Pruning histograms.
4. After reviewing the results, if necessary, make changes to the values for the pruning and re-run the stage.
5. When you are satisfied with the results, click **Next** to go to the Escalation stage.

Pruning Setup

You can make changes by pruning the CDT based on the customer count, the SKU count, the tree-level count, and the minimum CDT score. To do this, enter a minimum value for each pruning filter.

Figure 2–16 CDT Pruning Setup



Once you have completed setting up the filters, click **Run** to begin the processing. After you review the pruning results, you can change the values for the filters if you find it necessary. Once you make changes, you must run the stage again in order to see the results of your changes.

Table 2–17 Pruning Setup: Filters

Filter Name	Filter Description
Minimum Customer Count	The minimum number of customers to be used in the CDT calculation.The default is 1000.
Minimum SKU Count	The minimum number of SKUs to be used in the CDT calculation. The default is 10.
Minimum Tree Level Count	The minimum number of levels to be used in the CDT calculation. The default is 2.
Minimum CDT Score	The minimum measure of the quality of a CDT to be used in the CDT calculation. The default is 25.

Pruning Results

The Pruning Results are located below Pruning Setup and display information that can help you assess the effects of the values you provided for the pruning filters.

Figure 2–17 CDT Pruning Results

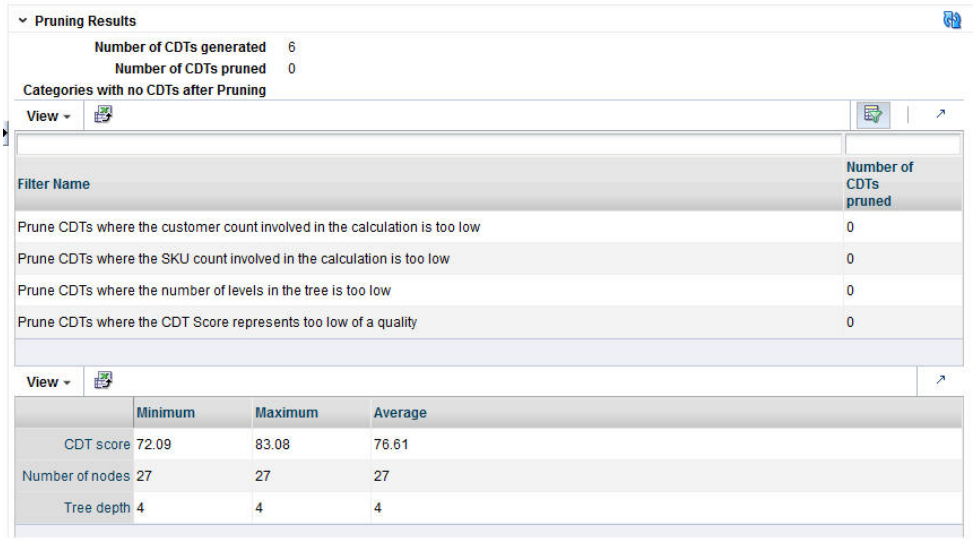


Table 2–18 Pruning Results

Field Name	Field Description
Number of CDTs Generated	The number of CDTs that were generated by the Calculation stage.
Number of CDTs Pruned	The number of CDTs that were pruned after the filters were applied.
Categories with no CDTs After Pruning	The names of the categories from which all CDTs were pruned.

Below this list is a table that identifies the number of CDTs that have been pruned by the following filters.

- Prune CDTs where the customer count involved in the calculation is too low.
- Prune CDTs where the SKU count involved in the calculation is too low.
- Prune CDTs where the number of levels in the tree is too low.
- Prune CDTs where the CDT score represents too low a quality.

A second table provides an overview of the pruning results, including the minimum, maximum, and average values for the CDT score, the number of nodes, and tree depth.

The number of nodes indicates the size of the tree. If a CDT contains few nodes, this can indicate a problem with the data or that too many nodes were excluded during the Calculation stage (because of a parameter setting).

Tree depth also indicates the size of the tree. This value can be used in conjunction with the Lowest Tree Level setting in the Calculation stage to analyze the results in terms of the number of levels in the tree.

Evaluation Report

The Evaluation Report displays the CDT results for the indicated version from the last successful run, if one has occurred. Here you can review information by Location or by Customer Segment. You can also set a version as active, view and edit a CDT, or compare two CDTs.

Table 2–19 Evaluation Report Fields

Field Name	Description
Partition	A partition is the combination of category, segment, and location. The column identifies the names of all nodes in the tree structure in the table. The node type may be either category, location, or customer segment.
Pruned	The number of CDTs removed from the list of usable CDTs. CDTs are removed that do not meet the filtering thresholds. During the Escalation stage, the escalation process makes adjustments for the CDTs that are removed.
Total Customer IDs	The number of customers used in the calculation of the CDT.
Total Weeks of Sales	The total number of weeks of sales used in the calculation of the CDT. This provides an indication of the amount of sales data used in the calculation. All the CDTs in a given version should have a similar value; if this is not the case, the results should be evaluated.
CDT Score	A confidence score assigned to help in assessing the quality of the CDT. A CDT score ranges from 0 to 100. A higher value indicates a better score.
Creation Date	The date when the version whose data is displayed was created.
Creation User	The login name of the user who created the version.

The following icons are available for the Evaluation Report:

Set Version as Complete

Each partition must have an active CDT. The active CDT is used by other applications that require a CDT.

All CDTs within a version are activated when you click **Set Version as Complete**.

If a CDT exists in more than one version, then the most recently activated version takes precedence.

If a version contains multiple categories and is Complete and then a different version with a partial overlap of categories is later marked as Complete, only the overlapping categories are replaced in the new version.

If a version is overwritten, then the CDTs in that version are no longer active.

When a CDT is at risk of going from active to inactive, you will see a warning message.

To determine whether or not a CDT is active, go to the Manage CDTs tab. The Browse by Categories or by Versions table displays a flag that indicates whether or not a CDT is active.

View One CDT

This icon provides access to the CDT Editor. For details about this functionality, see Using the CDT Editor.

Compare Two CDTs

This icon provides access to CDT Compare. For details about this functionality, see Comparing Two CDTs.

Three histograms are provided in the Evaluation stage.

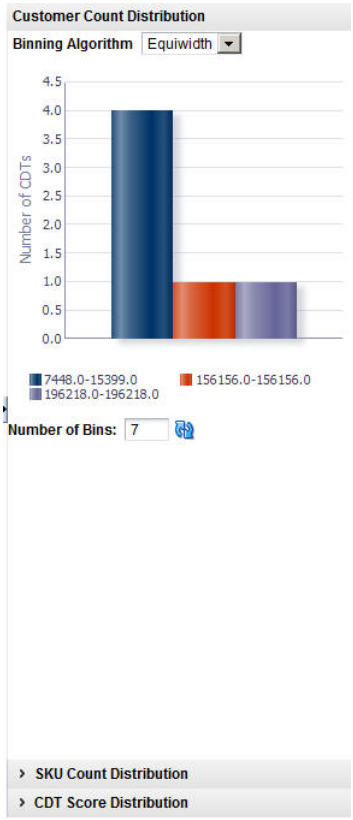
Evaluation Histograms

The following histograms are displayed in the Evaluation stage. You can use the information from the histograms to understand the data that was pruned by the filters.

Table 2–20 Evaluation Histograms

Name	Description
Customer Count Distribution	Displays the number of CDTs for a specific customer count.
SKU Count Distribution	Displays the number of CDTs for a specific SKU count.
CDT Score Distribution	Displays the CDT score distribution.

Figure 2–18 CDT Customer Count Distribution



When you have completed the evaluation of the results, click **Next** to go to the Escalation stage.

Escalation

The Escalation stage is used to fill in the holes for partitions whose CDTs were removed during pruning by setting up a search path through the segment hierarchy and the location hierarchy.

Process

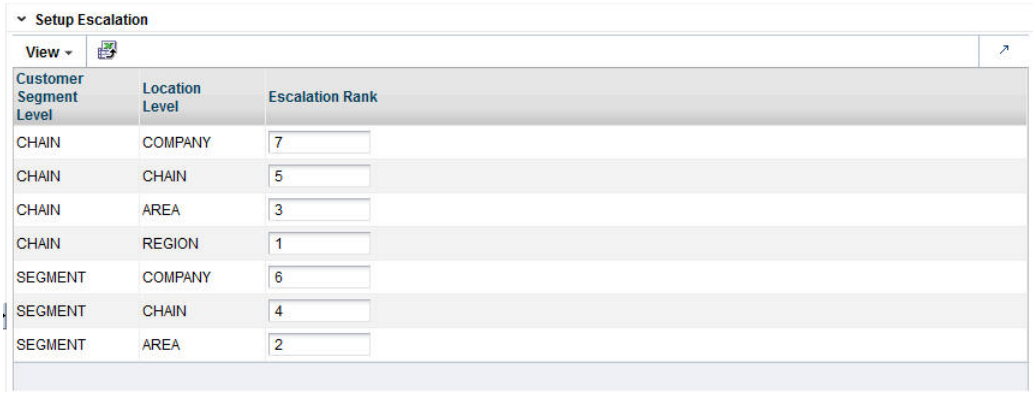
Here is the high-level process for setting up an escalation.

1. Enter a series of numbers to indicate the escalation rank, which determines the order in which the escalation occurs.
2. Click **Run** to start the escalation process.
3. Review the escalation results in the Escalation Report.
4. After reviewing the results, if necessary, make changes to the escalation ranks and rerun the stage.
5. When you are satisfied with the results, you can complete the version and make the version active so that it is available for other applications to use.

Setup Escalation

Escalation occurs along the segment hierarchy and the location hierarchy. Here is an example of an escalation path:

Figure 2–19 CDT Setup Escalation



The following fields are required to set up the escalation.

Table 2–21 Setup Escalation

Field	Description
Customer Segment Level	Identifies the customer segment level in the escalation.
Location Level	Identifies the location level in the escalation.
Escalation Rank	Used to assign the ranks for the escalation, which determines the order in which the escalation occurs.

Here is an example of an escalation path.

Table 2–22 Example of Escalation Path

Segment Level	Location Level	Escalation Rank
Segment chain	Location chain	8
Segment chain	Region	7
Segment chain	Location area	6
Segment chain	Store cluster	5
Segments	Location chain	4
Segments	Region	3
Segments	Location area	2
Segments	Store cluster	1

You fill in the order of numbers. Every row must have an ordering number, and no ordering number can be reused.

The escalation path is specific to the user and the current version that the user is working on.

The default ordering is to go up the location hierarchy first, and then up the segment hierarchy, as shown in the example above. The reason is that the segment hierarchy has only two levels, and so its top level is very general.

Escalation Report

The Escalation Report breaks down the numbers to provide counts for the number of partitions filled with higher-level CDTs and the number of partitions that have not

been changed by escalation. In addition, the fraction of CDTs for each partition is displayed.

Figure 2–20 CDT Escalation Report

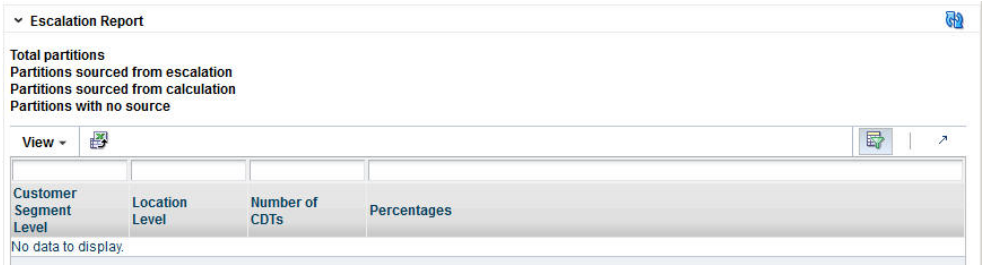


Table 2–23 Escalation Report

Field Name	Description
Total Partitions	The number of partitions in the version.
Partitions Sourced from Escalation	The number of partitions removed during escalation.
Partitions sourced from Calculation	The number of partitions removed during calculation.
Partitions with No Source	A partition that does not have CDTs assigned to it because all CDTs related to the partition have been pruned.
Customer Segment Level	Identifies the customer segment level.
Location Level	Identifies the location level.
Number of CDTs	The number of partitions that are trying to have a CDT assigned. This is generally the number of customer segments by the number of locations.
Percentages	The percentage of partitions that have been assigned a CDT from a given escalation level.

Completion of Process

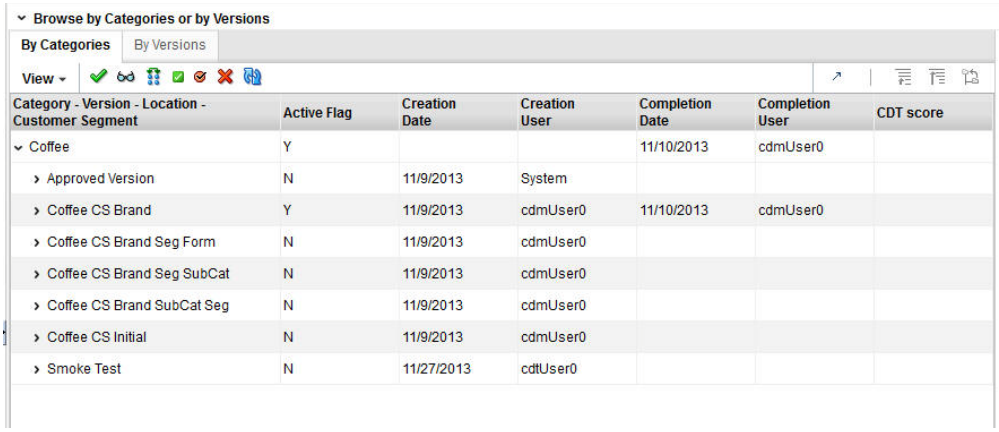
When a version is complete, the results for the version are activated so that other applications can use the information. The similarity data that has been calculated during the generation process is also activated for use.

After the completion of this step, the intermediate results from each stage is removed from the database and can no longer be used.

Manage CDTs Tab

The Manage CDTs tab is used to control which version of a CDT is active or approved. You can also access the functionality to edit a CDT or compare two CDTs.

Figure 2–21 Manage CDTs



You can find the CDT you are interested in by:

- Browsing by categories or versions
- Searching by version name or user name.

Set Version as Complete

Each partition must have an active CDT. The active CDT is used by other applications that require a CDT.

All CDTs within a version are activated when you click **Set Version as Complete**.

If a CDT exists in more than one version, then the most recently activated version takes precedence.

If a version contains multiple categories and is Complete and then a different version with a partial overlap of categories is later marked as Complete, only the overlapping categories are replaced in the new version.

If a version is overwritten, then the CDTs in that version are no longer active.

When a CDT is at risk of going from active to inactive, you will see a warning message.

To determine whether or not a CDT is active, go to the Manage CDTs tab. The Browse by Categories or by Versions table displays a flag that indicates whether or not a CDT is active.

View One CDT

This button provides access to the CDT Editor. For details about this functionality, see Using the CDT Editor.

Compare Two CDTs

This button provides access to CDT Compare. For details about this functionality, see Comparing Two CDTs.

Approve a Version

Only one approved CDT is permitted for a given CMPO node/segment/location combination. If you approve a second CDT with the same combination, it overwrites the first one.

If an Approved CDT exists, then the Approved CDT is returned to the requesting application.

If no Approved CDT exists, then the Active CDT is returned to the requesting application.

If no Approved CDT exists and no Active CDT exists, an error is returned to the requesting application.

Approval can be done for individual CDTs or for an entire version; activation is for an entire version.

An approved version is used by the CMPO application. An active version is used by Demand Transference.

Unapprove

Use this to change the status of a CDT that was previously approved. Select a single CDT record in the data table and click the **Unapprove** button to remove it from the approved version.

Delete a Version

If you delete a version, it cannot be activated. You cannot delete an approved version, but you can remove a CDT from an approved version.

Browse by Category or Version

You can browse by category or version. The following information is displayed:

Table 2–24 Browsing by Category or Version

Field Name	Field Description
Version...Customer Segment	Partition
Active Flag	Y = Version is active. N = version is not active.
Creation Date	Date on which version was created.
Creation User	User who created the version.
Activation Date	Date on which the version was made active.
Activation User	User who made the version active.
CDT Score	A confidence score assigned to help in assessing the quality of the CDT. A CDT score ranges from 0 to 100. A higher value indicates a better score.

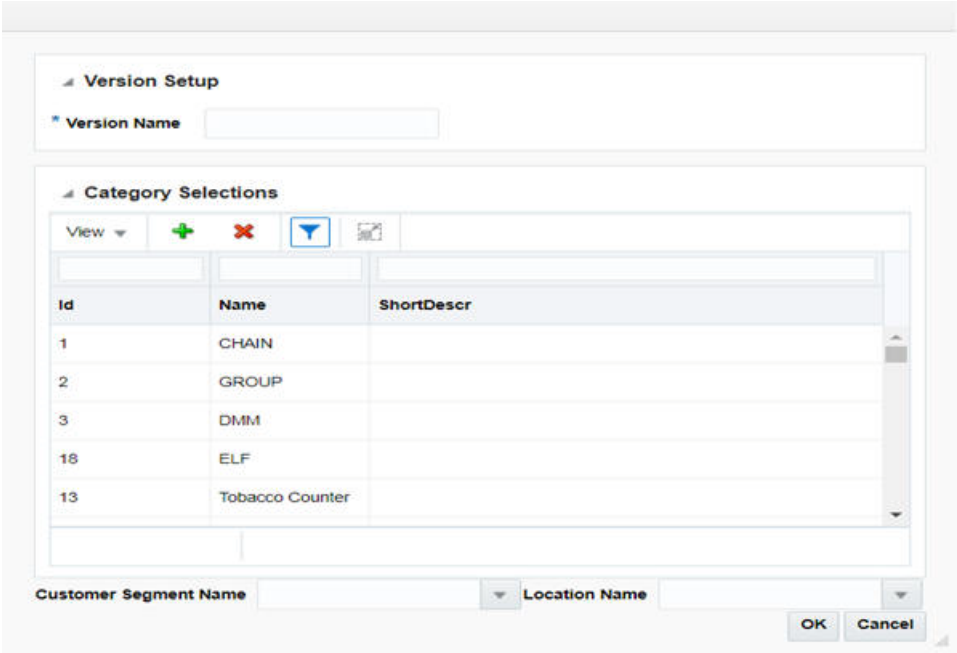
Search

You can search by version name or creation user name. In addition, you can design a search with specific criteria and save that search for future use.

Create CDT

You can create a new CDT in two ways: by modifying an existing CDT or by defining a new CDT.

Figure 2–22 Create CDT



1. Click the **Create CDT** icon, shown in Figure 2–22.
2. If you want to create a new CDT based on an existing one, select the version you want to use as a template from the list of Category Selections. If you want to create a new CDT, skip this step.
3. Enter the Version Name for the new CDT.
4. Select the Product Category, Customer Segment Name, and Location Name.
5. Click **OK**. You see the CDT Editor.
6. You can use the CDT editor, described in "Using the CDT Editor", to edit the new CDT and add a child, make changes to a branch, or change the assigned names.
7. When you have finished creating and editing the new CDT, click **Save CDT**.
8. To approve the new CDT, click **Approve CDT**.

Deleting a CDT

To delete a CDT, click the **Delete CDT** icon.

Using the CDT Editor

You can view or edit a CDT using the CDT Editor, which you access by clicking the **Edit** button. The CDT Editor displays the Customer Segment Name and Location Name for the CDT. It provides tools to navigate through the display using Zoom, Pan, Center, and Layout functionality.

The following functionality is available:

- Add a Child
- Edit
- Delete

- Delete Branch
- Copy Branch
- Paste Branch
- Save and Approve

Add a Child

To add a child node to the CDT, do the following:

1. Click **Add Child**. You see the Add CDT Node window.
2. Select an attribute from the **Attribute** drop-down menu.
3. Select a value from the **Attribute Value** drop-down menu.
4. Select the Create separate branch per value check box, if applicable. When more than one of the attribute values are selected, this field create a separate node for every attribute value selected.

   Note: A node can be created by selecting one or more attribute values. If multiple attribute values are chosen, and the check box to create a separate branch is not selected, one node that represents multiple attribute values is created.

   If one branch is created for all values of an attribute, select all attribute values and only one node will appear with the attribute value of ALL.

   If all the attribute values for an attribute are used, either by creating one node or separate nodes, a child cannot be added to the node and the **Add Child** is disabled. Creating a node with ALL does not associate the specific attribute values with the node. Instead, it encompasses any attribute value.

5. Click **OK**.

Editing Nodes

Editing is allowed on all nodes as long as there are additional attribute values available. The selections from the node being edited are preselected in the dialog. The attribute cannot be changed, so the attribute drop-down menu is not disabled. Any attribute values that are in use by other siblings are not be displayed in the list.

Deleting Nodes

Nodes can be deleted from the tree. There are two options for deleting: Delete and Delete Branch.

Selecting Delete deletes the selected node. Delete is available for any leaf node (a node without any children). It is also available for a non-leaf node that represents all the attributes for an attribute value: the children under this node move under the parent.

Selecting Delete Branch deletes the entire branch of nodes under the selected node, but not the selected node itself.

Copying/Pasting in Customer Decision Tree Editor

Copy and Paste functionality is provided to copy nodes from one branch to another. **Copy Branch** is enabled when a node is selected that has children. The Copy Branch function copies the full branch of children of the selected node, but not the selected node itself. **Paste Branch** is enabled when a branch has been copied and the selected node is a leaf node.

Viewing a Customer Decision Tree

The Customer Decision Trees have the potential to become large and occupy more space than the screen real estate allows for. The hierarchy viewer component used to display the Customer Decision Tree provides several features to assist in viewing the Customer Decision Tree effectively.

Expanding and Collapsing Branches

One way to limit the amount of space taken up by the Customer Decision Tree is to collapse branches of the tree. For any node that has children, a small triangle appears at the bottom of the box for that node. Hovering over the triangle enlarges it and displays an option to collapse that node if it is expanded and expand the node if it is collapsed.

Moving the Tree

If the entire Customer Decision Tree is not visible on one screen, the Customer Decision Tree can be moved to make other parts of the tree visible. The view can be moved by either clicking and dragging or by using the panning controls in the control bar for the hierarchy viewer.

Zooming

The hierarchy viewer provides some controls for zooming in and out to allow more or less of the tree to be in view at a time. Zooming out shrinks the size of the nodes, which may make them difficult to read.

Comparing Two CDTs

You can also compare two CDTs using the same CDT Editor functionality. You select the two CDTs you want to compare from the list. Both CDTs are displayed side by side.

3

Demand Transference

This chapter described the use of the Demand Transference Science Cloud Service module.

Introduction

Demand Transference (DT) helps you to compare products based on their similarities in order to determine what, if any, products customers might buy if the product they want to buy is for some reason unavailable. In this way, planning and ordering can be optimized. DT calculates similarities by comparing the attributes of the two products. If you are using CDT in conjunction with DT, you also have available the similarities calculated by CDT, which are based on customer-supplied transaction data.

The DT Cloud Service module consists of three tabs: Overview, Generate Models, and Manage Models. You use the Overview tab to keep track of the status of each stage during the main work you do with the application within the Generate Models tab. You use the Manage Models tab to evaluate the demand elasticity results and override the Maximum Substitutable Demand Percentage value, if that is needed.

Overview of DT Process

When you use the DT Cloud Service module, you follow this general iterative process to create and manage DT models:

- Data Setup. You define the categories to be used in the DT calculation.
- Data Filtering. You configure filters that remove input data that might cause errors in the calculation or that can lead to inaccurate or unreliable results.
- Similarity Calculation. You calculate similarities and assess the results of the calculation.
- Elasticity Calculation. You calculate the assortment elasticities and assess the results in terms of substitutable demand, which is the percentage of demand of a SKU that is retained when the SKU is deleted from the stores where it is selling.
- Escalation. When you are satisfied with your results, you can set the escalation path you can set the escalation path to fill in the holes for partitions whose DT models were removed during pruning by setting up a search path through the segment hierarchy and the location hierarchy. Then you can set a version of the DT model as complete.
- Manage Models. Use this tab to set time intervals for evaluating your results and to override the value for the maximum substitutable demand percentage.

Overview Tab

The Overview tab displays information that you can view and use to monitor the progress of the DT stages as well as to view some aggregate statistics and the DT results from the last successful run.

This tab contains the following sections:

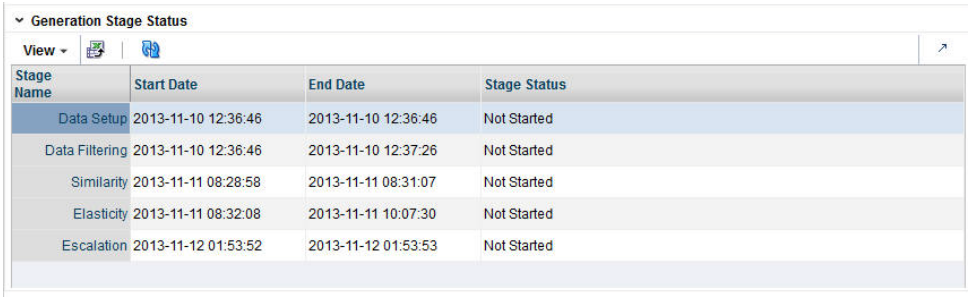
- Generation Stage Status
- Aggregate Statistics
- Calculation Report

Generation Stage Status

The Generation Stage Status table displays the current status of each of the DT stages.

Click the **Refresh** icon to update the fields and see the latest status for each stage.

Figure 3–1 DT Generation Stage Status



You can use the status information to monitor the progress of each stage. It contains the following fields, which can be arranged and viewed, but not modified.

Table 3–1 Generation Stage Status Fields

Field Name	Description
Stage Name	A row exists in the table for each DT stage that provides detailed status information. The five stages are Data Setup, Data Filtering, Similarity, Elasticity, and Escalation.
Start Date	The date and time when a run for the stage most recently started.
End Date	The date and time when a run for the stage most recently ended.
Stage Status	The current status of the stage: Not Started, Not Started (Scheduled for Later), Processing, Completed Successfully, Stopped/Paused, Cancelled, or Failed.

The following table provides an explanation of the different values for the status of a stage.

Table 3–2 Stage Status Values

Stage Status	Description
Not Started	This indicates that the stage has not yet been initiated. This can occur when you first begin to create a DT model or when you update an existing version.

Table 3–2 (Cont.) Stage Status Values

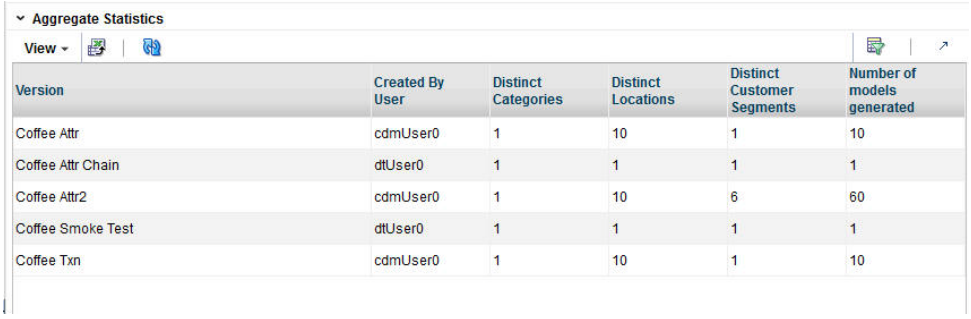
Stage Status	Description
Not Started (Scheduled for Later)	This indicates that the stage is scheduled to start after the previous stage is complete.
Processing	This indicates that the stage is currently being processed.
Stopped/Paused	This indicates either that the stage has been stopped by the application or that the user has chosen to stop the process (for example, to make a change to an option). In the later case, once any background processing has stopped, the user can re-run the stopped stage.
Completed Successfully	This indicates that the stage has been successfully processed.
Completed with Errors	This indicates that the stage was able to complete the processing of all requested CDTs, but one or more CDTs encountered errors during the run and were not able to complete successfully. This most commonly occurs when the data used for a CDT is not available or is too sparse to produce a result.
Cancelled	This indicates that the database has cancelled the execution of the stage either because of missing data or an exception. The user should review the RSE_LOG_MSG log to determine the problem.
Failed	This indicates that a problem occurred during the processing.

Aggregate Statistics

The Aggregate Statistics table displays statistical details about the existing DT model versions.

Click the **Refresh** icon to update the fields and see the latest information in this table.

Figure 3–2 DT Aggregate Statistics



Each existing version has a row in the table. The table contains the following fields, which can be arranged and viewed but not modified.

Table 3–3 Aggregate Statistics Fields

Field Name	Description
Version	User-created name that uniquely identifies this DT model version
Created By User	The user name of the person who created this version.
Distinct Categories	The number of categories associated with this version.
Distinct Locations	The number of locations associated with this version.

Table 3–3 (Cont.) Aggregate Statistics Fields

Field Name	Description
Distinct Customer Segments	The number of customer segments associated with this version.
Number of Models Generated	The number of DT models that have been calculated for this version.

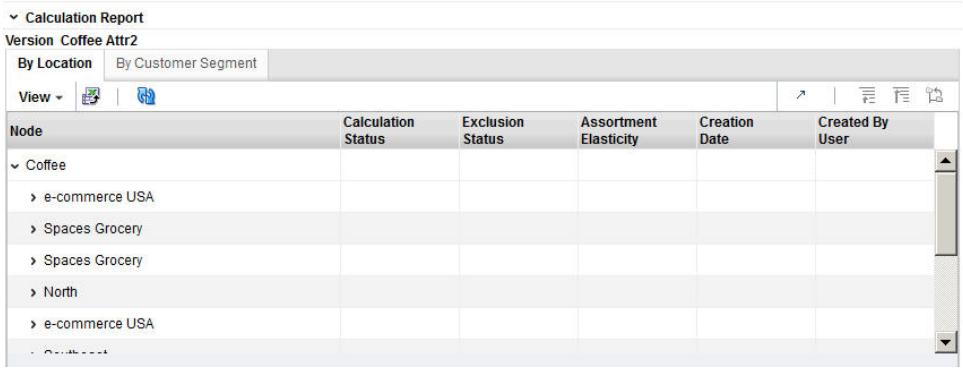
Calculation Report

The Calculation Report displays the DT model results from the last successful run, if one has occurred.

Here you can review information by Location or by Customer Segment.

Click the **Refresh** icon to update the fields and see the latest information in this table.

Figure 3–3 DT Calculation Report



The Calculation Report contains the following fields:

Table 3–4 DT Calculation Report Fields

Field Name	Description
Node	Identifies the node name.
Calculation Status	**Yes** indicates that the calculation is complete. **No** indicates that the calculation is not complete.
Exclusion Status	Yes indicates that data has been pruned. No indicates that data has not been pruned.
Assortment Elasticity	A number calculated by the application that is a parameter in the DT model. A larger magnitude indicates larger overall transference.
Creation Date	The date when the version whose data is displayed was created.
Created By User	The login name of the person who created the version.

Generate Models Tab

The Generate Models tab is used to configure, run, evaluate, modify, and deploy a DT model. The process is divided into five stages that must be run in order. You can return to a stage you have already completed and make changes, but if you do, you must re-run that stage and all the stages that follow that stage, as the calculations are invalidated by the modifications you just made to the settings in that stage.

The five stages are:

Table 3–5 Generate Models Tab: Stages

Stage Name	Description
Data Setup	Select the nodes for the DT model calculations.
Data Filtering	Filter out input data that may result in inaccurate or unreliable answers.
Similarity Calculation	Calculate the similarities in customer demand.
Elasticity Calculation	Calculate the assortment elasticities for customer demand.
Escalation	Set the escalation path for the DT model. Use the Escalation Report to evaluate the results of the escalation.

Data Setup

The Data Setup stage is used to add and delete the categories to be used in the DT model generation process.

Process

Here is a high-level process for setting up the data for DT.

1. Select the category or categories you want to calculate DT models for.
2. Click **Next** to go to the Data Filtering stage.

Category Selections

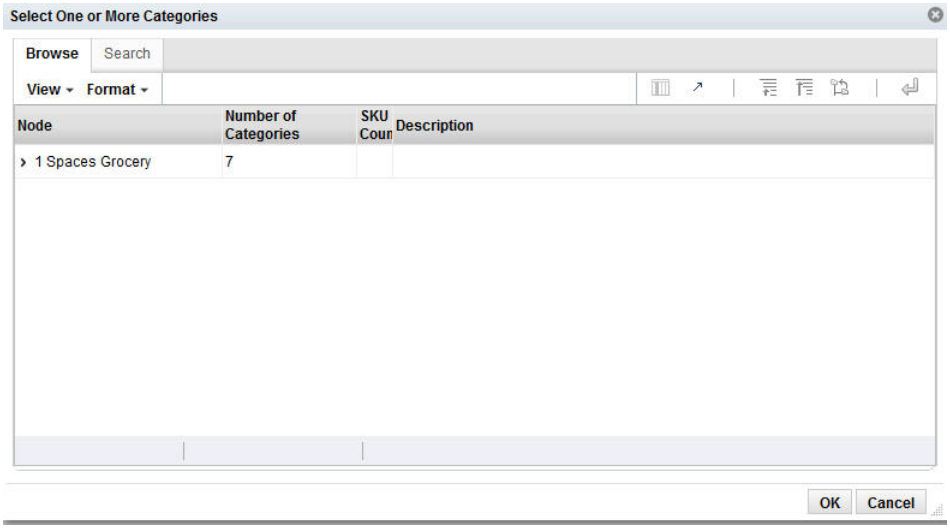
Use this table to add categories you want to include in the DT model calculation or delete categories that you want to remove from the DT model calculation.

Figure 3–4 DT Category Selections



To display a list of available nodes and the categories included in those nodes, click the **Add** icon. You see the Select One or More Categories dialog box, which contains two tabs, Browse and Search. You can use either tab to find the categories you are looking for.

Figure 3–5 DT Browse Categories



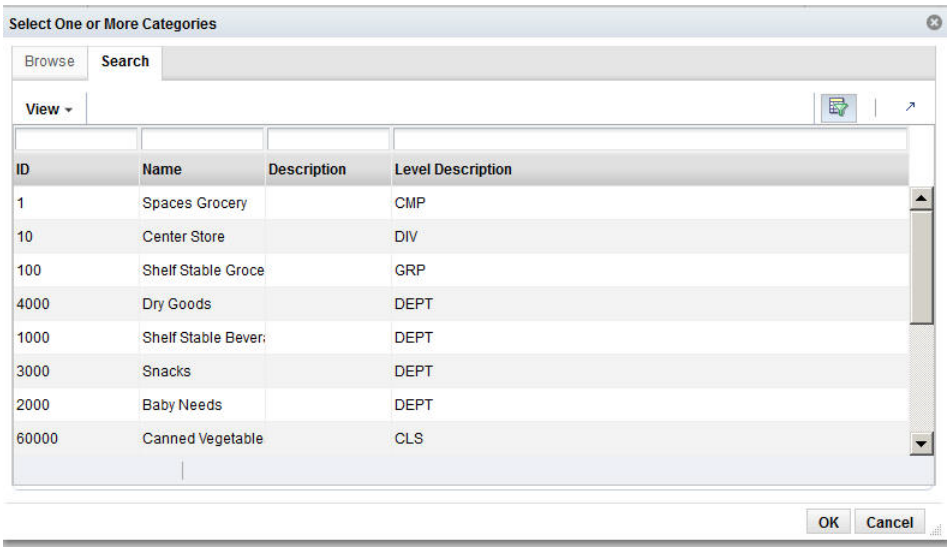
The Browse tab displays a table with the following fields:

Table 3–6 Category Selections: Browse

Field Name	Description
Node	The node tree structure can be expanded in order to view its categories.
Number of Categories	The number of categories within the node that has been selected. The number can help you understand the amount of processing required for the calculation.
SKU Count	The number of SKUs in a category. A category with too few SKUs may not produce good DT models.
Description	A description that provides additional information about the category.

Select the category or categories within the node that you want to be part of the DT model and click **OK**.

Figure 3–6 DT Search Categories



The Search tab displays a table with the following fields:

Table 3–7 Category Selections: Search

Field Name	Field Description
ID	An external code used to identify the category in other systems such as CMPO.
Name	The category name.
Short Description	A description that provides additional information about the category.
Level Description	A description of the level of the merchandise hierarchy that the node belongs to.

Select the category or categories within the node that you want to be part of the DT model and click **OK**.

Your selections are displayed in the Category Selections table.

After selecting the categories, click **Next** to go to the Data Filtering stage.

Data Filtering

The Data Filtering stage applies to all the categories that you select in the Data Setup stage. You should set the filters based on the histograms for each filter. The histograms help identify what data is actually outlier data, as compared to the rest of the data. In most cases, the default settings should be sufficient. However, if a histogram shows a flatter distribution, then you should consider modifying the default settings.

Process

Here is the high-level process for setting up and running data filtering.

1. Enter the appropriate values into the Filter Setup text entry boxes.
2. Click **Run** in order to filter the data.
3. Review the filtering results in the Data Filtering Summary table and the Data Filtering histograms.

4. After reviewing the results, if necessary, make changes to the values for the filters in Filter Setup and re-run the stage.
5. When you are satisfied with the results, click **Next** to go to the Similarity Calculation stage.

Filter Setup

You configure the following filters in order to filter out data you consider unacceptable from the calculation of the DT model.

Figure 3–7 DT Filter Setup

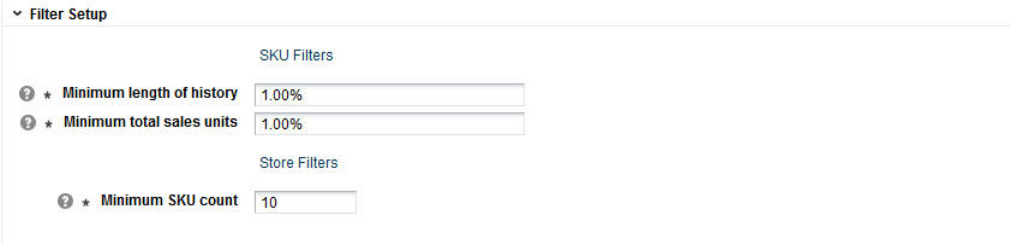
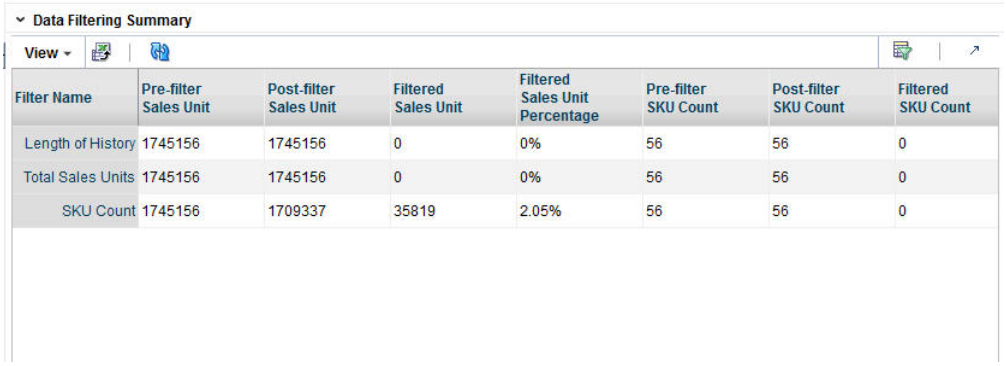


Table 3–8 Data Filters

Filter Name	Description
Minimum Length of History	This filter prunes SKU-segment-store combinations that have a short transaction history. The threshold is defined as a percentage of the median value for the category. The default value is 1%.
Minimum Total Sales Units	This filter prunes SKU-segment-store combinations that have a small number of total sales units during a given sales history for a specified customer segment and store. The threshold is defined as a percentage of the median value for the category. The default value is 1%.
Minimum SKU Count	This filter is applied after the above two filters and looks at the remaining data to determine if a store does not have enough SKUs. The threshold is defined as a set number of SKUs per store. The default value is 10 SKUs.

Data Filter Summary

Figure 3–8 DT Data Filter Summary



Filter Name	Pre-filter Sales Unit	Post-filter Sales Unit	Filtered Sales Unit	Filtered Sales Unit Percentage	Pre-filter SKU Count	Post-filter SKU Count	Filtered SKU Count
Length of History	1745156	1745156	0	0%	56	56	0
Total Sales Units	1745156	1745156	0	0%	56	56	0
SKU Count	1745156	1709337	35819	2.05%	56	56	0

The following information is provided after the filtering is complete and quantifies the amount of data filtered out for baseline history, total sales units, sales amounts, and SKU counts. Use this information to assess the effects of filtering.

Click the **Refresh** icon to update the fields and see the latest information for this table.

Table 3–9 Data Filter Summary Fields

Field Name	Field Description
Filter Name	The relevant filter of the three listed above.
Pre-filter Sales Unit	Amount prior to the application of the filter.
Post-filter Sales Unit	Amount remaining after the application of the filter.
Filtered Sales Unit	Amount filtered.
Filtered Sales Unit Percentage	Amount filtered, expressed as a percentage.
Pre-filter SKU Count	Amount prior to application of the filter.
Post-filter SKU Count	Amount remaining after application of the filter.
Filtered SKU Count	Amount filtered.
Filtered SKU Count Percentage	Amount filtered, expressed as a percentage.

Data Filtering Histograms

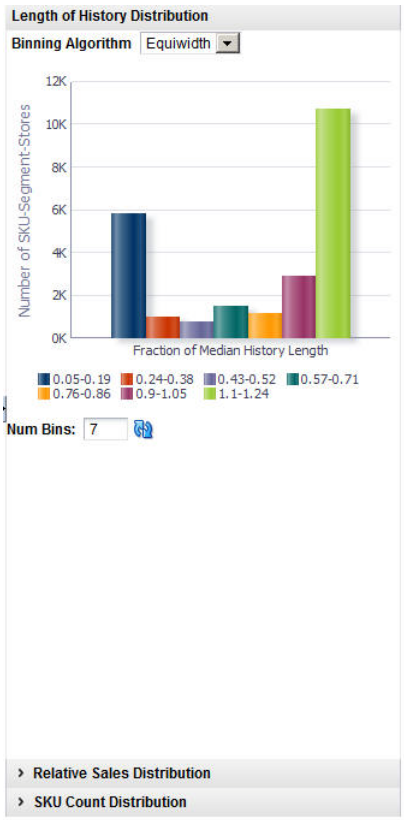
The following histograms illustrate the effects of filtering. You can use the information displayed in the histograms to adjust the configuration of the filters in order to eliminate outlier data. If you modify the filters, you must re-run the stage.

For information about adjusting the display of the histograms, see the Chapter 1, "Getting Started."

Table 3–10 Data Filtering Histograms

Histogram Name	Description
Length of History Distribution	Displays the percentage of median history length relative to the number of SKU-segment-stores.
Relative Sales Distribution	Displays the percentage of median category sales relative to the number of SKU-segment-stores.
SKU Count Distribution	Displays the percentage of the median category SKU count relative to the number of segment-stores.

Figure 3–9 DT Length of History Distribution Histogram



Similarity Calculation

Similarity in demand can be determined using either transaction-based data or attribute data. You can calculate the similarity using each type of data, if both types are available. You can only view the most recent run in the UI, so in order to compare runs, you must query the database to obtain the results from earlier runs.

Each category has its own set of similarities, relevant to the SKUs that are in that category. A similarity is calculated for each pair of historical SKUs in a category.

If transaction-based similarities are available, it is recommended that you use them instead of attribute-based similarities. Note that transaction-based similarities are only available through the Customer Decision Tree application.

Process

Here is the high-level process for calculating similarities:

1. Enter a unique name for the version of the DT model to be calculated.
2. Select the source of the data to be used: transaction-based or attributed-based. Transaction-based data is only available from the Customer Decision Tree generation.
3. Use the check boxes to indicate whether or not only top level processing should occur for location or customer segment.
4. Customize the ranking of the category attributes if necessary.
5. Click **Run** to start the calculation.

6. Review the calculation results in the Similarity Display.
7. After reviewing the results, if necessary make changes to the values for the calculation in the Version Setup and Category Attribute Setup and then re-run the stage.
8. When you are satisfied with the results, click **Next** to go to the Elasticity Calculation stage.

Version Setup

At the top of this stage you see two text boxes and two check boxes that you use to configure the parameters for the calculation.

Figure 3–10 DT Version Setup



Table 3–11 Version Setup: Fields

Field Name	Field Description
Version Name	Assign a name to each version of a DT model calculation. This allows you to create and save more than one version of a DT model. The version name you assign here is used in the Calculation Report, Aggregate Statistics table, and in the Manage DTs tab. Version names can be re-used; however, if the version name in question has active DT models, then you will see a warning that the active DT models will be removed from the version if you do re-use the version name.
Select the Source of Similarities	Use this option to define the type of data used in the calculation: transaction data or attribute data. Transaction-based data uses similarities calculated by Customer Decision Tree using transaction-based data. Attribute-based data calculates similarities within DT based on the attribute values associated with every SKU in the category.
Process Location Top Level Only	Check this option if you want DT models to be calculated for the Location Chain *only*. You can select this option in order to decrease the amount of time it takes the system to perform the calculation.
Process Customer Segment Top Level Only	Check this option if you want DT models to be calculated for the Customer Segment Chain *only*. You can select this option in order to decrease the amount of time it takes the system to perform the calculation.

Once you select the source for the similarities, you will see either Category Attribute Setup, if you have selected to use attribute-based similarities, or Transaction-based Similarity Availability Per Category, if you have selected to use transaction-based similarities.

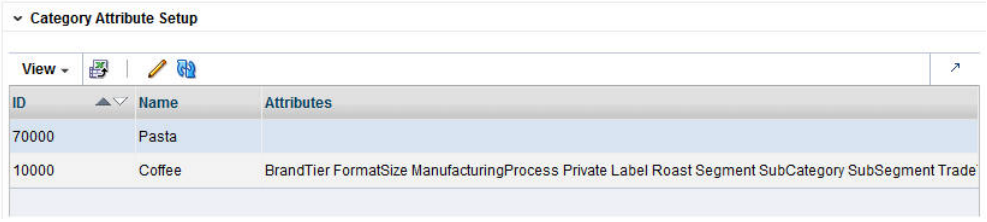
Category Attribute Setup

The application, using historical data, determines a specific weight for the category attributes. You can optionally change this weight and assign your own weight to the attributes for a category.

The weights indicate the importance of the attribute to the customers when they are making purchasing decisions. The attribute with the highest weight is the one the customer considers first when making a purchase. The system-generated weights are determined by the application from historical sales data. However, if a user disagrees with those weights, the user can override them. For example, in the case of coffee, the system may assign a weight of 0.7 to brand and 0.2 to size. This indicates that brand is historically more important to the customer than size when purchasing coffee. If the user disagrees with this analysis and thinks that brand and size are actually much closer together, the user can assign a weight of 0.5 to brand and 0.4 to size.

The Category Attribute Setup table displays the following:

Figure 3–11 DT Category Attribute Setup



Highlight the category you want to adjust the weights for and click the **Edit** icon. You see the Category Attributes Setup dialog box.

Figure 3–12 DT Category Attribute Setup

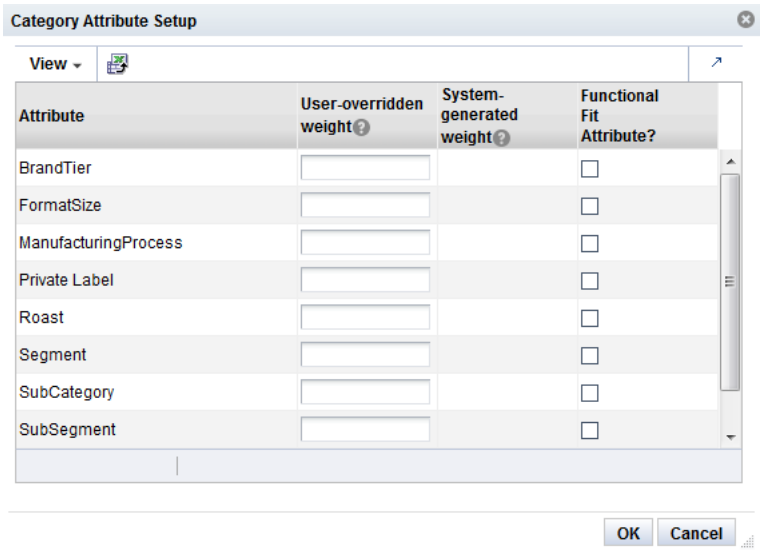


Table 3–12 Category Attributes Setup

Field	Description
ID	An external code used to identify the category in other systems such as CMPO.
Name	The category name.
Attribute	The specific attribute you are configuring.

The Category Attributes Setup pop-up lists the categories you are calculating DT models for. The system-assigned weights are also displayed.You can adjust the weight for all the attributes or a subset of the attributes.

The Category Attributes Setup dialog box contains the following fields. For each attribute you want to assign a custom weight to, enter a number between 0.000 and 1.000. For attributes that have no substitutes (such as windshield wipers of a specific length), the Functional Fit check box is checked by the system, so that similarities are not calculated for these attributes. When you are finished configuring the category attributes, click **OK**.

Table 3–13 Category Attribute Setup Fields

Field	Description
Attribute	The category attribute to assign a weight to.
User-Overridden Weight	The user-defined weight for the attribute.
System-Generated Weight	The system-generated weight for the attribute.
Functional Fit Attribute?	This is checked by the system if the attribute has no substitutes.

After you have finished configuring the similarity parameters, click **Run** to calculate the similarities. You see the results via the Similarity display table.

Transaction-Based Similarity Availability Per Category

The Transaction-Based Similarity table displays the following:

Table 3–14 Transaction-Based Similarity Availability Per Category

Field	Description
ID	An external code used to identify the category in other systems such as CMPO.
Name	The category name.
Description	A description that provides additional information about the category.
Available	A flag that indicates that a CDT version that contains data for this category has been made active.
Available As Of	Indicates the date that the CDT version was activated. This information can help you identify whether the CDT results are recent, or if they are potentially too old to use. For example, if the CDT data became available two years ago, you may consider that data to be out of date.

Similarity Display

The Similarity Display table shows the list of SKUs for which similarities have been calculated so that you can sort and analyze the results. You can search through the list of results by Category Name, Location, or Customer Segment.

Click the **Refresh** icon to update the fields and see the latest information in this table.

Figure 3–13 DT Similarity Display

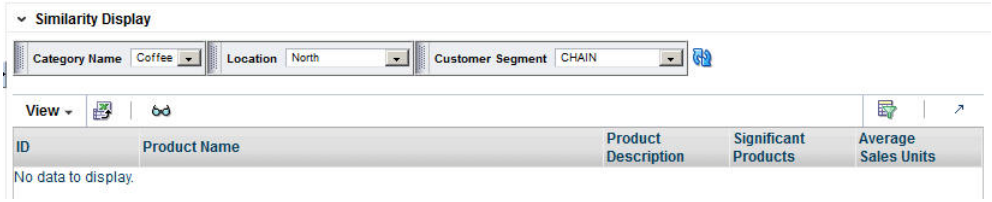


Table 3–15 Similarity Display

Field	Description
ID	An external code used to identify the category in other systems such as CMPO.
Product Name	The name identifying the product.
Product Description	A detailed description of the product.
Significant Products	A list of products that have High or Very High similarity. The threshold for how many products are considered significant can be configured in the database.
Average Sales Units	The average number of units used in the calculation.

Click the **See Similarities** icon to see detailed results for a specific set of SKUs.

The detailed results include the following fields:

Table 3–16 Similarity Display Results

Field Name	Field Description
ID	The product SKU.
Product Name	The name identifying the product.
Product Description	A detailed description of the product.
Similarity Strength	An indication of the similarity for the product: Very High, High, Medium, Low, Very Low.
Similarity Value	The calculated value for the similarity, from 0 to 1. A higher value indicates a higher degree of similarity.
Similarity Code	The numeric value associated with the similarity: 4 = Very High, 3 = High, 2 = Medium, 1 = Low, 0 = Very Low.

When you are satisfied with the Similarity results, click **Next** to go to the Elasticity Calculation stage.

Elasticity Calculation

During the Elasticity Calculation stage, the assortment elasticity is calculated. You do not configure any parameters. Click **Run** to initiate the calculation.

The assortment elasticity should not be a positive value because the transference model does not work properly if the value is positive. In addition, it should not be a null value because a null value indicates that the calculation of assortment elasticity failed and did not produce an assortment elasticity value. If an assortment elasticity value is positive, it must be replaced with a negative value. The replacement occurs during the escalation process.

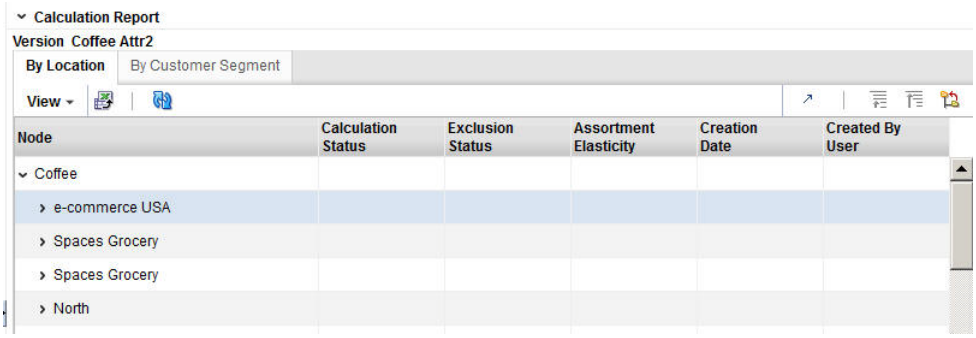
Process

The elasticity calculation is a background process. You use this stage to view the results. Note that the substitutable demand information is displayed and the percentages for the DT models are calculated after you set the time intervals within the Manage Models tab.

Calculation Report

The Calculation Report lists the status of the elasticity calculation and the exclusion, either by Location or by Customer Segment. An assortment elasticity is calculated for each category/location/segment combination selected during the Data Setup stage and the Calculation stage. The numerical result of the calculation, an output of DT generation, is used by Manage Models to calculate substitutable demand percentages.

Figure 3–14 DT Calculation Report



The Calculation Report has two tabs: By Location and By Customer Segment. Each tab has the following fields:

Table 3–17 Calculation Report Fields

Field	Description
Node	The node tree structure can be expanded in order to view its categories
Calculation Status	Indicates whether or not (Yes/No) the elasticity calculation has occurred.
Exclusion Status	Indicates whether or not (Yes/No) pruning has occurred.
Creation Date	The date when the version whose data is displayed was created.
Created By User	The login name of the person who created the version.
Assortment Elasticity	A number calculated by the application that is a parameter in the DT model. A larger magnitude indicates larger overall transference.

Pruning Report

The Pruning Report displays statistics about the results of data pruning.

Figure 3–15 DT Pruning Report



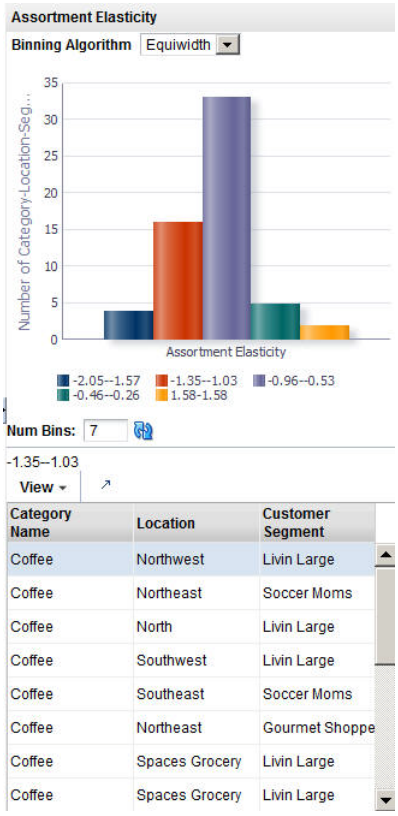
The Pruning Report contains the following fields:

Table 3–18 Pruning Report Fields

Field	Description
Number of Models Generated	The number of DT models produced by the calculation.
Number of Models Pruned	The number of DT models pruned by the calculation.
Categories with All Models Pruned	The names of the categories for which all DT models have been pruned.
Positive or Null Assortment Elasticity	The Assortment Elasticity (AE) value is a negative number used in calculating substitutable demand percentages and demand transference effects. A positive AE value may be produced as a result of missing or unreliable input data for some partitions. Such meaningless transference effects must be pruned. In addition, if a null AE value is generated, it must also be pruned.
Substitutable Demand	Substitutable demand is a measure of how much demand is retained by the rest of the assortment when an item that is removed. When the item is removed, a portion of its demand is transferred to the remainder of items in the assortment. These values are populated when you set up the time intervals in the Manage Models tab and run the calculation there.

Assortment Elasticity Histogram

Figure 3–16 DT Assortment Elasticity Histogram



Escalation

The Escalation stage is used to fill in the holes for partitions whose DT models were removed during pruning by setting up a search path through the customer segment hierarchy and the location hierarchy. The DT models used to fill in the holes are not used by Manage Models calculations.

Process

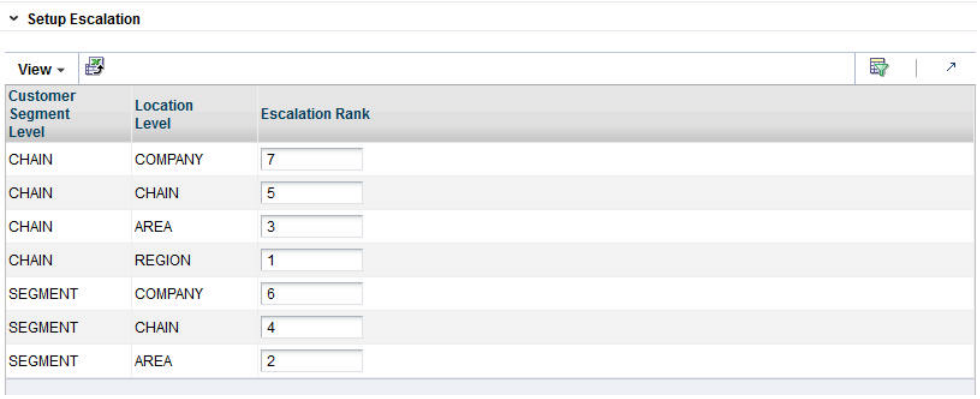
Here is the high-level process for setting up an escalation.

1. Enter a series of numbers to indicate the escalation rank, which determines the order in which the escalation occurs.
2. Click **Run** to start the escalation process.
3. Review the escalation results in the Escalation Report.
4. After reviewing the results, if necessary, make changes to the escalation ranks and re-run the stage.
5. When you are satisfied with the results, you can complete the version and make the version active so that it is available for other applications to use.

Setup Escalation

Escalation occurs along the customer segment hierarchy and the location hierarchy. Here is an example of an escalation path:

Figure 3–17 DT Setup Escalation



The following fields are required to set up the escalation.

Table 3–19 Setup Escalation

Field	Description
Customer Segment Level	Identifies the customer segment level in the escalation.
Location Level	Identifies the location level in the escalation.
Escalation Rank	Used to assign the ranks for the escalation, which determines the order in which the escalation occurs.

Here is an example of an escalation path.

Table 3–20 Example of Escalation Path

Segment Level	Location Level	Escalation Rank
Segment chain	Location chain	8
Segment chain	Region	7
Segment chain	Location area	6
Segment chain	Store cluster	5
Segments	Location chain	4
Segments	Region	3
Segments	Location area	2
Segments	Store cluster	1

You fill in the order of numbers. Every row must have an ordering number, and no ordering number can be re-used.

The escalation path is specific to the user and the current version that the user is working on.

The default ordering is to go up the location hierarchy first, and then up the segment hierarchy, as shown in the example above. The reason is that the segment hierarchy has only two levels, and so its top level is very general.

Escalation Report

The Escalation Report breaks down the numbers to provide counts for the number of positions filled with higher-level DT models and the number of partitions that have

not been changed by escalation. In addition, the fraction of DT models for each partition is displayed.

Figure 3–18 DT Escalation Report

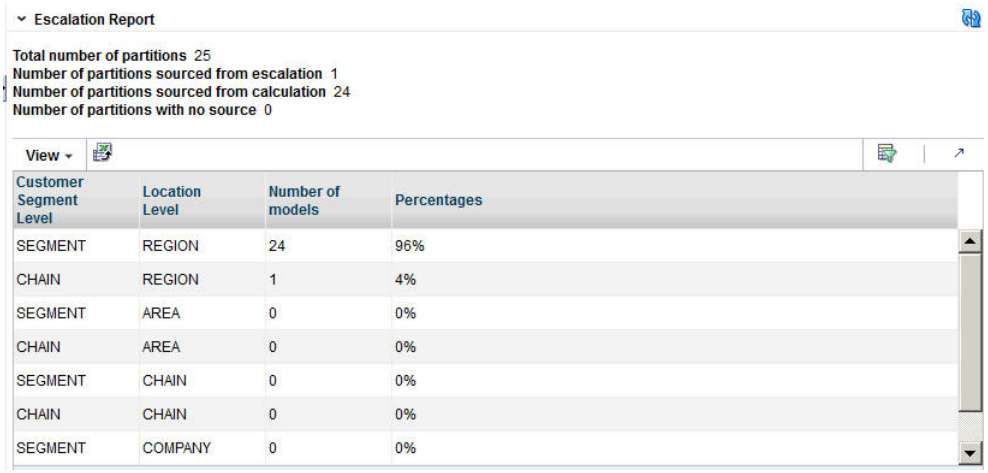


Table 3–21 Escalation Report

Field Name	Description
Total Number of Partitions	The number of partitions in the version.
Number of Partitions Sourced from Escalation	The number of partitions removed during escalation.
Number of Partitions Sourced from Calculation	The number of partitions removed during calculation.
Number of Partitions with No Source	A partition that does not have a model assigned to it because all models related to the partition have been pruned.
Customer Segment Level	Identifies the customer segment level.
Location Level	Identifies the location level.
Number of Models	The number of partitions that are trying to have a model assigned. This is generally the number of customer segments by the number of locations.
Percentages	The percentage of partitions that have been assigned a model from a given escalation level.

Completion of Process

When a version is complete, the results for the version are activated so that other applications can use the information. The similarity data that has been calculated during the generation process is also activated for use.

After the completion of this step, the intermediate results from each stage are removed from the database and can no longer be used.

Be aware that once a version is completed, it cannot be completed again unless a different version is completed first. Changes made to the version's data after completing it will not be copied to the relevant output tables.

Manage Models Tab

You can set up various time intervals to use in the evaluation of a version of a DT and configure the value for maximum substitutable demand and see how different maximum values affect the substitutable demand. This allows you to change the maximum value to one you find more suitable. You can see the percentage of demand of a SKU that is retained when the SKU is deleted from the stores where it is selling. This is the substitutable demand percentage. In this way you can evaluate the accuracy and usability of the elasticity calculation.

Substitutable demand is a measure of how much demand is retained by the rest of the assortment when an item is removed. When the item is removed, a portion of its demand is transferred to the remainder of items in the assortment. This portion is considered the retained demand. If the magnitude of the assortment elasticity is larger, then the amount retained will be higher. By examining the retained demand, you can evaluate the assortment elasticity value to see if its magnitude is too large. The key value to examine is the maximum substitutable demand percentage. For a given category, you may decide that this value is too large.

Process

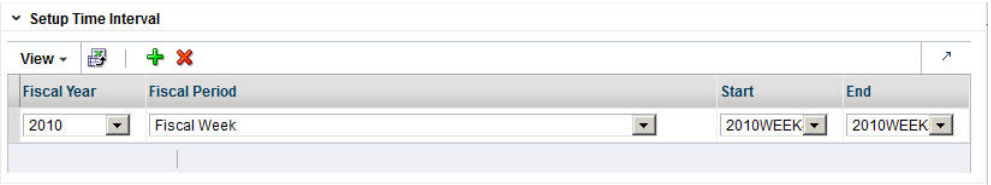
Here is the high-level process for determining suitable substitutable demand values and thus suitable DT models.

- Set up the time intervals you are interested in.
- Select the versions you want to evaluate.
- Click the **Calculate** icon to obtain an initial set of percentages.
- Enter various values for maximum substitutable demand and use the **Edit** icon to enter override values for Maximum Substitutable Demand Percent and the **Calculate** icon to determine the impact.
- You can use the **Revert** icon to restore the original percentage values.
- When you are satisfied with the results, you can choose which version to make active using the **Set Version as Complete** icon.

Time Interval Setup

The time interval defines the span of time for the sales history to be used to determine the amount of history that is retained when SKUs are dropped. A group of intervals can be defined. Gaps between intervals are permitted; however, intervals cannot overlap.

Figure 3–19 DT Setup Time Interval



You should select a time interval for which the historical assortments are reasonably representative of the assortments that will be used in the CMPO application. Because the time interval is used to calculate the substitutable demand information, selecting a representative interval provides substitutable demand information that is highly relevant to the actual application of demand transference in CMPO. Typically, the most

representative time period is a recent time interval, since that is generally when assortments are most similar to the current assortments. If you use a time period that is not recent, you run the risk of using assortments that are not as similar to the current ones. You should also make sure not to select an interval that is too large, because a large interval necessarily includes several assortment changes within that interval. An interval size of approximately four weeks is recommended.

The fields that define a time interval are:

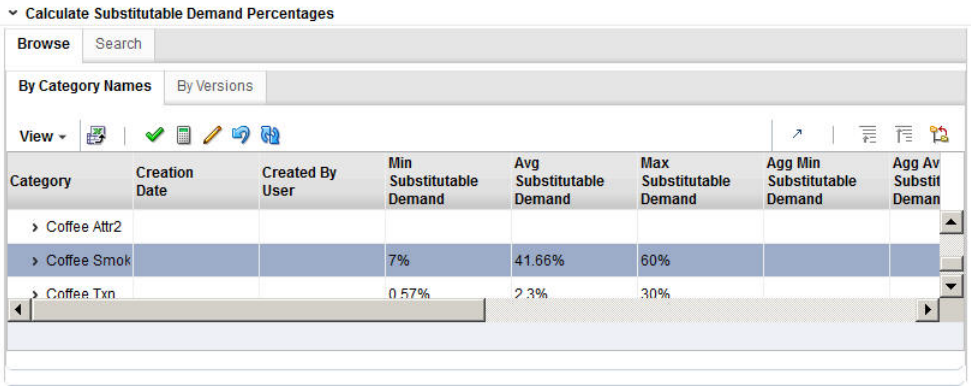
Table 3–22 Setup Time Interval: Fields

Field Name	Field Description
Fiscal Year	The fiscal year for the time interval.
Fiscal Period	The fiscal period within the fiscal year (Fiscal Quarter, Fiscal Period, or fiscal Week).
Start	The time unit when the time interval specified in Fiscal Period begins.
End	The time unit when the time interval specified in Fiscal Period ends.

Calculate Substitutable Demand Percentages

You can vary the value for the maximum substitutable demand percentage and see the impact on selected categories or versions.

Figure 3–20 DT Calculate Substitutable Demand Percentages



Browsing and Searching

You can browse by Category Names or by Versions. You can also search by name or by creation user. You see the following DT model data displayed:

Table 3–23 Browsing

Field Name	Description
Category/Version	The name of the category.
Creation Date	The date when the version was created.
Created By User	The user name of the person who created this version.
Min Substitutable Demand	The percentage value for the minimum substitutable demand.
Avg Substitutable Demand	The percentage value for the average substitutable demand.

Table 3–23 (Cont.) Browsing

Field Name	Description
Max Substitutable Demand	The percentage value for the maximum substitutable demand.
Agg Min Substitutable Demand	The aggregated value for the minimum substitutable demand.
Agg Avg Substitutable Demand	The aggregated value for the average substitutable demand.
Agg Max Substitutable Demand	The aggregated value for the maximum substitutable demand.
Calculation Status	Indication of whether or not (Yes/No) the calculation is complete.
Completion Status	Indication of whether or not (Yes/No) the status is active.
Completion Date	The date when the status became active.
Completion User	The user name of the person who activated the version.

Once you have made a selection, click the **Edit** icon and enter an override value between 0% and 100% for the Maximum Substitutable Demand Percent. Click the **Calculate** icon to initiate the calculation. To revert the calculation, click the **Revert** icon.

Once you have determined the substitutable demand value you want, you can click the **Complete** icon to make the version active.

Escalation Report

The Escalation Report breaks down the numbers to provide counts for the number of positions filled with higher-level DT models and the number of partitions that have not been changed by escalation. In addition, the fraction of DT models for each partition is displayed.

Figure 3–21 DT Escalation Report

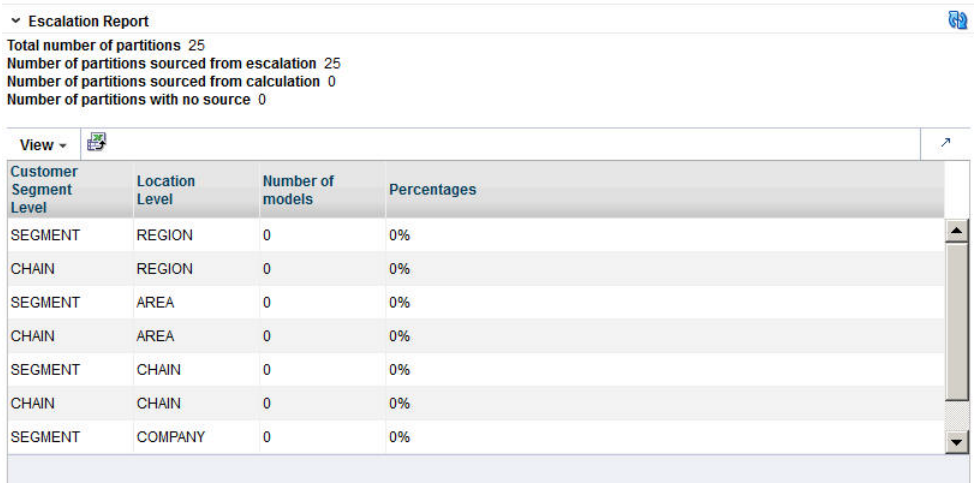


Table 3–24 Escalation Report

Field Name	Description
Total Number of Partitions	The number of partitions in the version.

Table 3–24 (Cont.) Escalation Report

Field Name	Description
Number of Partitions Sourced from Escalation	The number of partitions removed during escalation.
Number of Partitions Sourced from Calculation	The number of partitions removed during calculation.
Number of Partitions with No Source	A partition that does not have a model assigned to it because all models related to the partition have been pruned.
Customer Segment Level	Identifies the customer segment level.
Location Level	Identifies the location level.
Number of Models	The number of partitions that are trying to have a model assigned. This is generally the number of customer segments by the number of locations.
Percentages	The percentage of partitions that have been assigned a model from a given escalation level.

Histogram

The Substitutable Demand Distribution histogram displays the distribution of the substitutable demand values for SKU/stores.

4

Advanced Clustering

This chapter describes the Advanced Clustering Cloud Service module.

Introduction

Advanced Clustering is an enterprise-specific clustering solution that uses data mining to create store groupings at different product levels using a variety of inputs. These inputs include performance data (sales dollars, sales units, and gross profit), product attributes (brand, color, and size/fit), store attributes (climate, store format, size, and servicing distribution center), third-party data such as demographics (income, ethnicity, and population density), and customer segments.

The application's embedded science and automation helps you to identify unique patterns within your data that you can use to create the necessary customer-centric and targeted clusters. These can be used by the assortment planning, allocation and replenishment, pricing, and promotion processes.

It optimizes clusters in order to determine the minimum number of clusters that best describes the historical data used in the analysis and that best meets your business objectives, which you define during the design of your clusters.

You can use Advanced Clustering to execute localized or customer-centric assortments and for pricing. In addition, the application can help you when forecasting, for example, if you want to cluster stores based on similar seasonal patterns. You can also use the application for allocation, by clustering stores based on similar selling patterns.

Features

The key features of Advanced Clustering Cloud Service include:

- Scenario-based cluster generation, based on store or product attributes, customer segment profiles, or performance.
- Three-step cluster-generation process.
- What-if capabilities that can be used to create multiple clustering scenarios and then measure them against one another. This can help ensure that the most appropriate clusters are used by the applicable planning and execution processes.
- Automatic ranking of cluster scenarios to support what-if comparisons. Recommendations for the optimal cluster scenario and number of clusters are provided.
- Dynamic nesting of clusters, in which nested or mixed attribute clusters are created based on multiple attributes, performance data, and customer segments.

- Two types of algorithms are used.
  - Proprietary BaNG (Batch Neural Gas) algorithm for convergent cluster parameters
  - K-means approach for creating clusters in a hierarchical manner, which automatically determines the best attributes to split into an additional cluster.
- A variety of distance metrics that are suitable for real-value attributes, categorical attributes, profile-based measurements, and time-based performance.

Overview of Advanced Clustering Process

To use Advanced Clustering, follow this general process to create and manage clusters, working in the Generate Store Clusters tab and the Manage Store Clusters tab:

- Cluster Criteria. View all available clusters for specified merchandise, location, and calendar. Review cluster criteria or scenario details for each cluster. Use existing cluster as the basis for creating a new cluster.
- Explore Data. Examine the input data for the cluster. Review multiple data points and significant attributes using the contextual information.
- Cluster Setup. Define multiple what-if scenarios. Such scenarios can be compared with one another throughout the clustering process.
- Cluster Results. View the scenario results and compare scenarios.
- Cluster Insights. Gain an understanding about cluster results and cluster performance prior to approval by examining the contextual information.
- Manage Store Clusters. Manage existing cluster criteria. Perform manual overrides and approve clusters.

Cluster Criteria Overview Tab

The Cluster Criteria Overview tab displays a list of the most recently defined cluster criteria and provides the status, clustering criteria, applicable merchandise, location, and calendar nodes. You can click on the criteria name in order to access it within the Generate Store Clusters tab.

Figure 4–1 Cluster Criteria Overview Tab

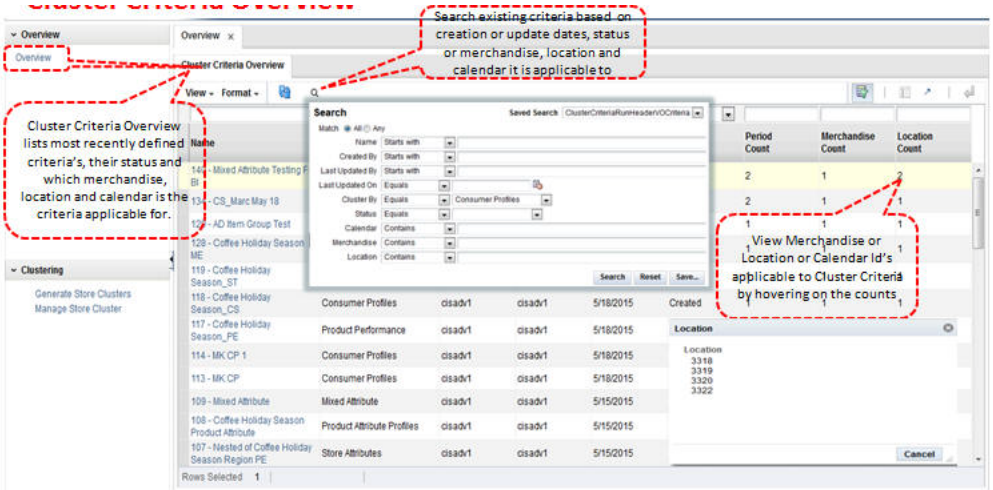


Table 4–1 Cluster Criteria Overview Tab

Field	Description
Name	The criteria ID and user-assigned name of the cluster.
Cluster By	A predefined group of attributes that include Consumer Profile, Product Performance, Store Attribute, Product Attribute, and Mixed Attribute. These criteria types are sets of attributes. For example, store attributes are the properties of a store. These properties can include ethnicity, store format, and store size.
Created By	The name of the user who created the cluster.
Last Updated By	The name of the user who most recently updated the cluster.
Last Updated On	The date when the cluster was most recently updated.
Status	Created, Ready for Approval, Completed with Errors, Approved, Rejected.
Period Count	The number of calendar nodes defined for the criteria. Hover over the count in order to see a list of the calendar keys associated with the criteria.
Merchandise Count	The number of merchandise nodes defined for the criteria. Hover over the count in order to see a list of the merchandise keys associated with the criteria.
Location Count	The number of location nodes defined for the criteria. Hover over the count in order to see a list of the location keys associated with the criteria.

Clustering Criteria

The following clustering criteria (which are also called "Cluster by") are the defaults:

Consumer Profile

Cluster stores based on the similarities in the customer profile mix whose members shop in the stores or trading areas. These clusters form the basis for additional analysis that can provide an understanding of which customers shop in which stores and how they shop. Information from market research firms such as the Nielsen Corporation can help retailers develop customer profiles. Such information can be provided via a data interface.

Location Attributes

Cluster stores based on how shopping behavior varies by store attribute. In combination with the profile mix, this provides an understanding of demographic details such as income level, ethnicity, education, household size, and family characteristics. Such knowledge can help the retailer to make assortment and pricing decisions. By analyzing cluster composition and studying business intelligence, the retailer can make informed decisions based on shopper demographics.

Product Attributes

Store share is generated based on product attributes. The store clusters produced can be used in an assortment. In this type of cluster, stores with a similar share of sales for one or more attributes are grouped together. For example, for the product coffee, stores can be differentiated by the sales patterns for premium, standard, and niche brands. The percentage of each store contribution is calculated using Sales Retail $ for each product attribute value to the total sales retail for the category or subcategory in a specified location. Product attributes can only be configured at the category or subcategory level.

Performance Criteria

Cluster stores based on the historical sales metrics by performance at various merchandise levels. Determine how shopping behavior varies by category. This information can be helpful in identifying low, medium, and high volume stores that all have similar sales patterns.

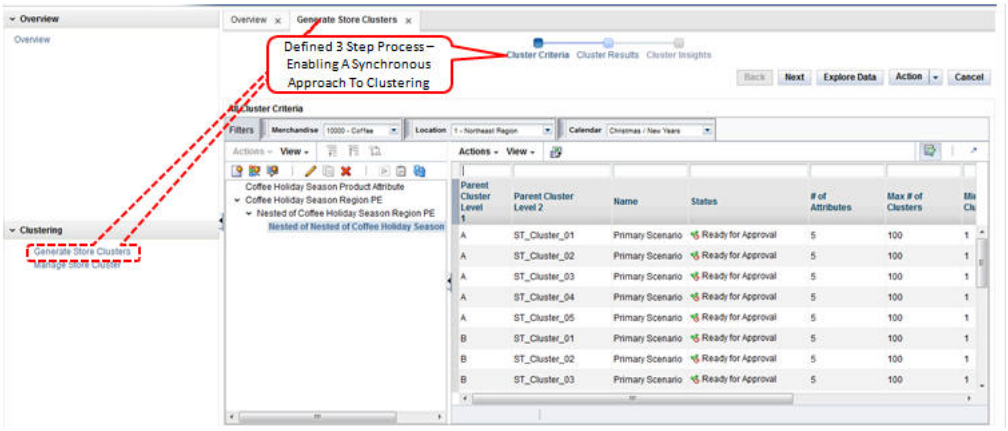
Mixed Criteria

Mixed criteria combines discrete and continuous attributes together. This allows a retailer to cluster stores using attributes from all of the first four listed cluster criteria at the same time.

Generate Store Clusters Tab

The Generate Store Clusters tab is used to create clusters and then model the clusters with various scenarios in order to determine the best clusters. It consists of three stages: Cluster Criteria, Cluster Results, and Cluster Insights.

Figure 4–2 Generate Store Clusters Tab



Cluster Criteria Stage

In this stage, you can view summary data about existing clusters and define the characteristics of new clusters.

Process

Here is the high-level process for defining a cluster.

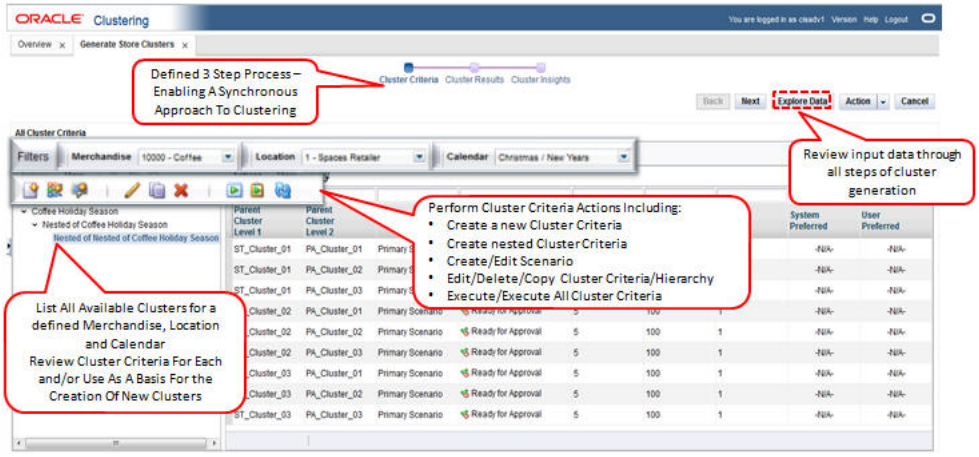
1. Provide a unique name for the cluster.
2. Define the type of data used to characterize the cluster.
3. Select merchandise and location nodes.
4. Define the time period for the cluster.
5. Define the historical time period for the data.

All Cluster Criteria

In this area of the page you can view information about existing clusters.

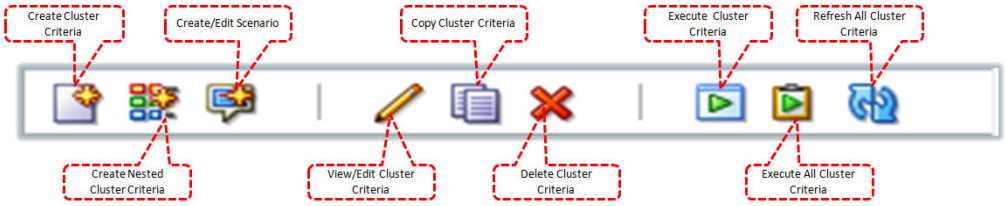
Use the View list to select existing cluster criteria. You can tailor your search for existing clusters by Merchandise, Location, and Calendar. Once you select a cluster, the defining details for that cluster are displayed in the Worksheet area.

Figure 4–3 All Cluster Criteria



In addition, you can use the toolbar buttons to:

Figure 4–4 All Cluster Criteria Toolbar



- Create a cluster criteria. You can define an initial cluster. The criteria include cluster name, the category for the cluster, the effective date for the cluster, and the history to use.
- Created a nested cluster criteria. In this way you can subdivide an existing cluster in order to analyze it further. Once you create a nested cluster, the name "Nested of <name of original cluster>" appears in the Cluster Criteria area. You can then define its characteristics in the same way you define any cluster.
- Create or edit a scenario. Modify the configured primary scenario or create another scenario to perform what-if analysis for a selected cluster criteria.
- View or edit a cluster criteria. The characteristics of the cluster are displayed in the Cluster Criteria pop-up.
- Copy a cluster criteria. Once you have copied it, you can modify it.
- Delete a cluster criteria. Delete the selected cluster criteria.
- Execute a cluster criteria. Execute all non-executed scenarios for the selected criteria.
- Execute all cluster criteria. Execute the entire cluster hierarchy criteria for the selected criteria at once.

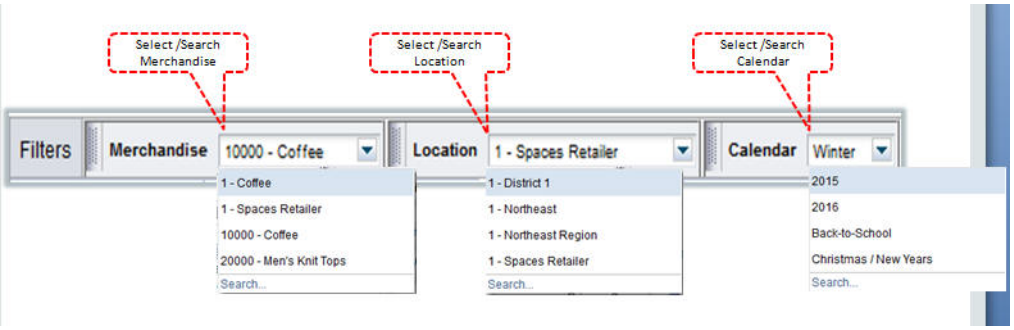
- Refresh all. Refresh all cluster criteria in order to view any updates to the existing cluster criteria.

All Cluster Criteria Filter

The All Cluster Criteria filter provides the following:

- Merchandise allows the user to filter the existing cluster criteria by searching for or selecting the merchandise for the supported hierarchies.
- Location allows the user to filter the existing cluster criteria by searching for or selecting the location for the supported hierarchies.
- Calendar allows the user to filter the existing cluster criteria by searching for or selecting the calendar for the supported hierarchies.

Figure 4–5 All Cluster Criteria Filter



All Cluster Criteria Summary

Once you have highlighted a cluster criteria to examine, details about that cluster criteria are displayed in a pop-up. The details include information about the cluster criteria and the scenarios created for that cluster criteria.

Table 4–2 Pop-Up Details

Field Name	Description
Cluster By	A predefined group of attributes that include Consumer Profile, Product Performance, Store Attribute, Product Attribute, and Mixed Attribute. These criteria types are sets of attributes. For example, store attributes are the properties of a store. These properties can include ethnicity, store format, and store size.
Shared Criteria	A check mark indicates that more than one merchandise or location node is used in the cluster criteria.
Merchandise Type	The merchandise type.
Scenario Created	The number of scenarios created for the cluster.
Scenario Executed	The number of scenarios executed for the cluster.
Location Type	The location type.
Parent Cluster Level	The name of the ancestor cluster that has been further clustered.

All Cluster Criteria Scenario List

This displays the scenarios for the selected cluster criteria in the All Cluster Criteria tree.

Table 4–3 Scenario List

Field Name	Description
Name	The name assigned to each scenario that has been created for the cluster.
Status	Created, Ready for Approval, Completed with Errors, Approved, Rejected.
User Preferred	Indicates whether or not the user prefers the cluster.
System Preferred	Indicates whether or not the system prefers the cluster.
# of Attributes	The number of attributes that were used in the cluster.
Max. # of Clusters	A user-provided value for the maximum clusters centers that the clustering process should consider.

Cluster Criteria

In this pop-up, you define the initial clustering parameters for the cluster criteria of a new cluster. Note that multiple hierarchies are supported in order to facilitate comparisons between clusters. For example, you can compare clusters for the market and retail location hierarchy.

Figure 4–6 illustrates how to use a simple approach to clustering by selecting attributes from a Cluster by. For example, you can select the performance Cluster by and generate clusters using store sales units or revenue.

Figure 4–6 Cluster Criteria

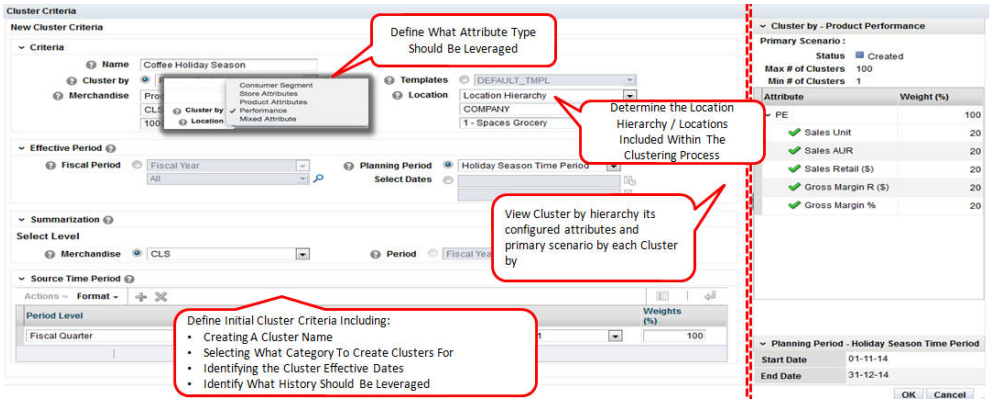
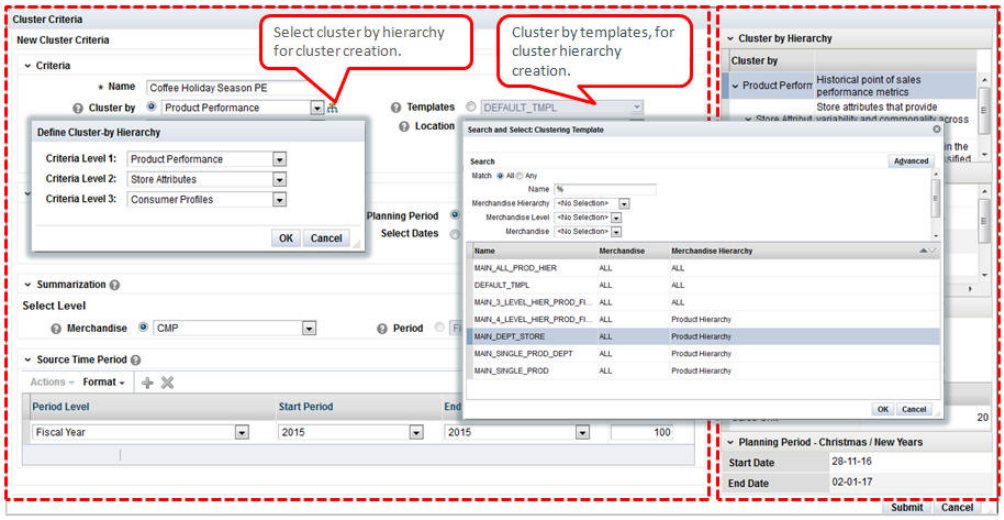


Figure 4–7 illustrates the use of a nested approach to clustering. Select a Cluster by hierarchy using the dynamic hierarchy pop-up or the pre-configured template hierarchies. For example, you can first create a cluster using the performance Cluster by and then further cluster using location attributes.

Figure 4–7 Cluster Criteria - Selection



The following information defines a cluster:

Table 4–4 New Cluster Definition

Field Name	Description
Name	A unique name to identify the cluster.
Cluster By	A predefined group of attributes that include Consumer Profile, Product Performance, Store Attribute, Product Attribute, and Mixed Attribute. These criteria types are sets of attributes. For example, store attributes are the properties of a store. These properties can include ethnicity, store format, and store size.
Merchandise	Once you choose the merchandise level for the cluster, you must select the hierarchy type, the hierarchy level, and the hierarchy node. These are specific to the merchandise level you select.
Location	Once you choose the location level for the cluster, you must select the hierarchy type, the hierarchy level, and the hierarchy node. These are specific to the location level you select.
Template	Select by name a predefined template that can be used to create a cluster hierarchy.

Effective Period

You can define a time interval for the cluster by either choosing a period from the list provided or by selecting a start date and an end date.

To define the Effective Period, you select either Planning Period, Fiscal Period, or Select Date:

Table 4–5 Effective Period

Option	Description
Fiscal Period	If you select this option, choose the period and the subdivisions of that period from the drop-down lists.
Planning Period	Select from the range of values provided for the period. Planning periods are user-defined buying periods for a season or a season subset.

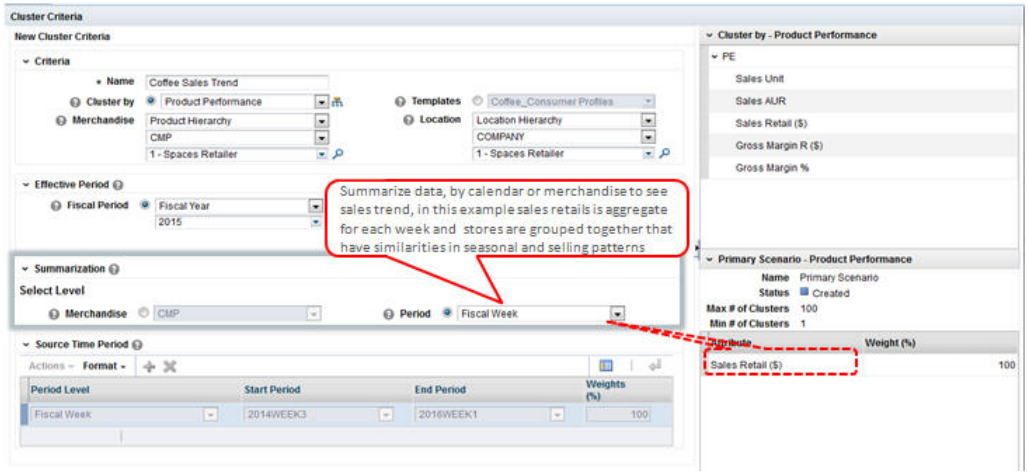
Table 4–5 (Cont.) Effective Period

Option	Description
Select Dates	If you select this option, choose the start and end dates using the calendar pop-up.

Summarization

Data summarization is available when you select the product performance Cluster by. Select the dimensions of the hierarchy (merchandise or calendar) to summarize the data and consider the dimension position in the clustering process. For example, when you use the category/week sales data to generate store clusters, you can select the week dimension in the calendar hierarchy summarization. The clustering process considers all weeks (week1, week2,...week52) as attributes and clusters stores based on their weekly sales patterns.

Figure 4–8 Summarization



Source Time Period

Select historical sales data for clustering and view contextual data to analyze cluster performance. You can specify more than one time period and assign different weights to different periods in order to place more or less emphasis on different periods.

Source time periods are available for all cluster criteria. With product and performance or mixed criteria, when performance metrics are used for clustering, these define the historical data used for the calculation. This time period also defines the historical data used to display BI when sales metrics are shown.

Table 4–6 Source Time Period

Field	Description
Period Level	Select from Fiscal Year, Fiscal Quarter, Fiscal Period, or Fiscal Week.
Start Period	Once you select the Period Level, you select the starting subdivision within that period.
End Period	Once you select the Period Level, you select the ending subdivision within that period.
Weight (%)	Used to define the weight given to the historical data from the defined time period.

Contextual Area

When you are creating a new cluster criteria, you can see details about the following parameters that can help you understand the cluster you are creating.

Cluster By Hierarchy The following information is displayed when you select a template or use the icon to select the Cluster by hierarchy.

Template

Table 4–7 Template Display

Property	Description
Template Name	Name of template configured during deployment.
Description	Description of template.
Cluster By	A predefined group of attributes that include Consumer Profile, Product Performance, Store Attribute, Product Attribute, and Mixed Attribute. These criteria types are sets of attributes. For example, store attributes are the properties of a store. These properties can include ethnicity, store format, and store size.

Hierarchy

A dynamic Cluster by hierarchy is displayed. For example, the template PE-ST-ST has a Cluster by hierarchy of performance/store attribute/store attribute.

Cluster By Primary Scenario You see this when you select Cluster by in the Criteria panel when you are setting the cluster parameters or when you select Cluster by in the contextual area for the hierarchy.

The attributes configured and the primary scenario properties defined for the selected cluster are displayed. The attributes listed are those that are significant for the clustering defined during deployment.

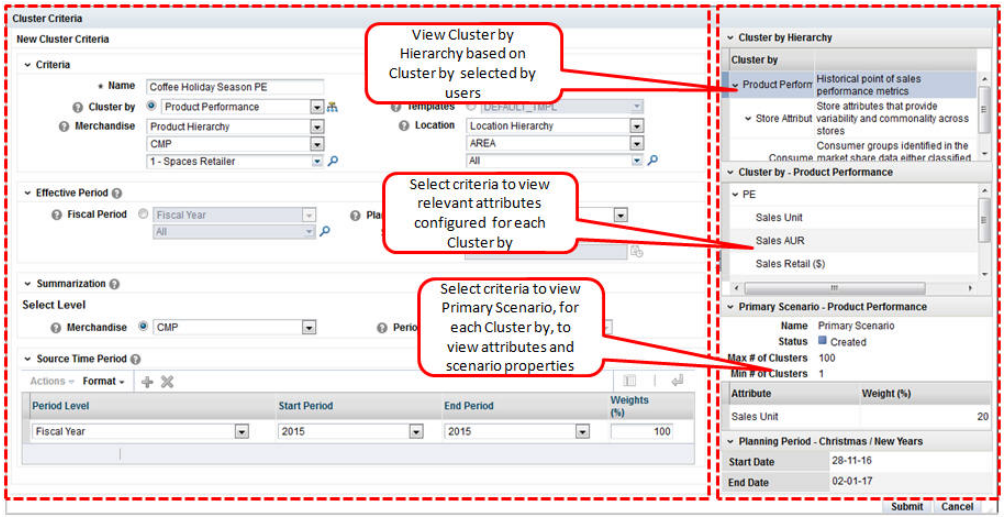
The primary scenario is the default scenario defined during deployment. The following information is displayed.

Table 4–8 Primary Scenario

Property	Description
Name	The name of the primary scenario.
Status	Created, Ready for Approval, Completed with Errors, Approved, Rejected.
Maximum # clusters	The maximum number of clusters. The default value is 100. This is used for analyzing the clusters.
Minimum # clusters	The minimum number of clusters. The default value is 1. This is used for analyzing the clusters.
Attribute	A list of the attributes configured during clustering.
Attribute weight	The weights associated with each attribute. This is used to calculate distance.

Planning Period This list displays the time period you selected for the cluster definition. This information is available only for planning periods, where it provides the start and end dates of the planning period. This content changes whenever planning period is selected in Effective Period when you are setting cluster parameters.

Figure 4–9 Contextual Information



Explore Data

Use the Explore Data pop-up to examine data for the cluster you defined. You can view the store that provides input into the clustering process.

Process

In this pop-up you can only view the data, so the only actions you can perform are drilling down through the data in the table and altering the arrangement of the table.

Summary

This area lists the criteria you initially selected to define the cluster.

Figure 4–10 Cluster Criteria Summary

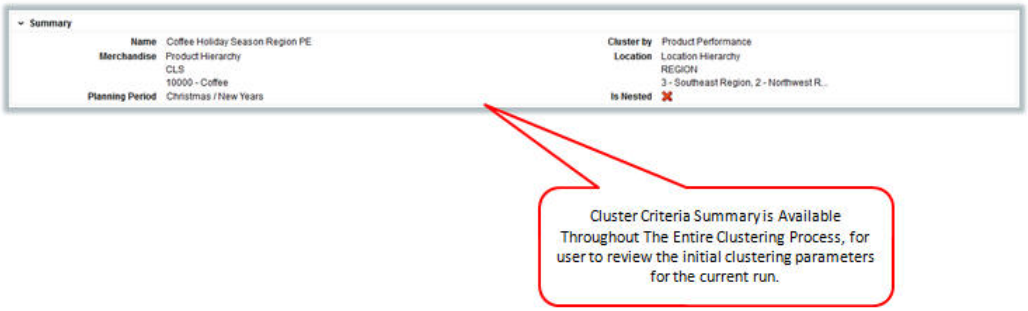


Table 4–9 Explore Data: Summary

Field	Description
Name	The name you provided for the cluster in the Cluster Criteria stage.
Cluster By	A predefined group of attributes that include Consumer Profile, Product Performance, Store Attribute, Product Attribute, and Mixed Attribute. These criteria types are sets of attributes. For example, store attributes are the properties of a store. These properties can include ethnicity, store format, and store size.

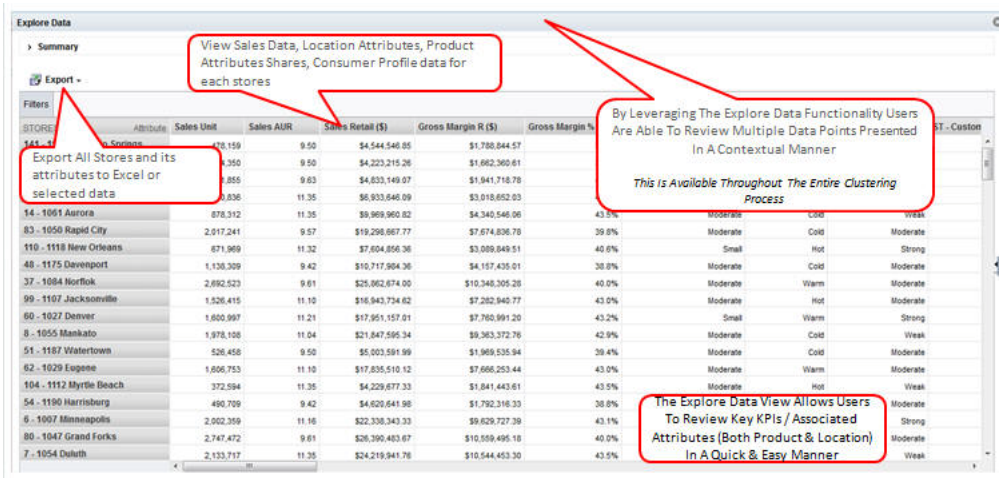
Table 4–9 (Cont.) Explore Data: Summary

Field	Description
Merchandise	The merchandise level and nodes for the cluster.
Location	The location level and nodes for the cluster.
Fiscal Period	The time period for the cluster.
Is Nested	Indicates whether or not the cluster is nested within another cluster.
Merchandise Hierarchy Type	Provides details about which type of hierarchy the cluster criteria have been created for.

View Stores

This area displays a nested list of the stores in the cluster you have defined and data for each store for each of the relevant attributes for the Cluster by option you selected to define the cluster. You can see data at the aggregated level as well as at the individual level. Filters are provided so that you can filter the display, for example, by category. You can see aggregated data at a higher level as well as at the store level.

Figure 4–11 Explore Data



Product Performance

This section displays stores, store sales metrics, location attributes, customer profiles, and product attribute profiles.

Table 4–10 Product Performance

Field	Description
Sales Unit	The sales units for the merchandise, location, and source time period selected when setting up clustering parameters.
Sales Average Unit Retail	The average unit retail sales for the merchandise, location, and source time period selected when setting up clustering parameters.
Sales Retail	The sales revenue for the selected merchandise, location, and source time period.

Table 4–10 (Cont.) Product Performance

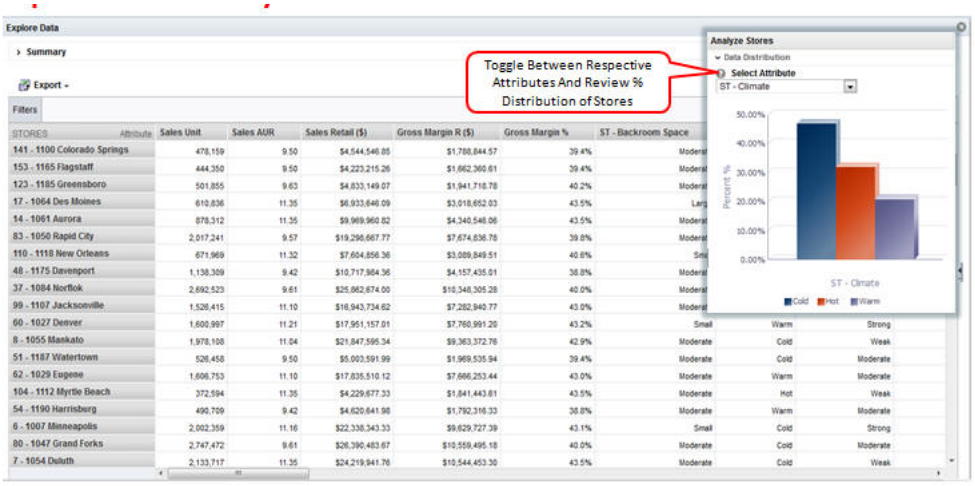
Field	Description
Gross Margin	The retail sales minus the cost of goods sold for the merchandise, location, and source time period selected when setting up clustering parameters.
Gross Margin Percent	The retail sales minus the cost of goods sold divided by the retail sales for the merchandise, location, and source time period selected when setting up clustering parameters.
Location Attribute	The retailer-configured location attributes.
Customer Profile	The retailer-configured customer profiles.
Product Attribute	The retailer-configured product attributes.

Contextual Area

This area provides a graphical illustration of the detailed data distribution about the cluster.

Analyze Stores In Explore Data, the BI displays the data distribution of the location by each participating attribute as well as other configured informational attributes. Advanced Clustering identifies the bins based on the underlying data and displays the histograms. It provides the percentage of stores that are present in a location. For example, a company may have 45 percent of stores in cold regions.

Figure 4–12 Clustering Analyze Stores



Category Variability By analyzing store variability, you can determine if it is worth creating store clusters for the selected categories in the selected location. Three sections are displayed.

A grid is displayed for the selected categories and the sales contribution for a selected location.

Table 4–11 Category Variability

Property	Description
Categories	A list of the selected categories that are used for store variability analysis.

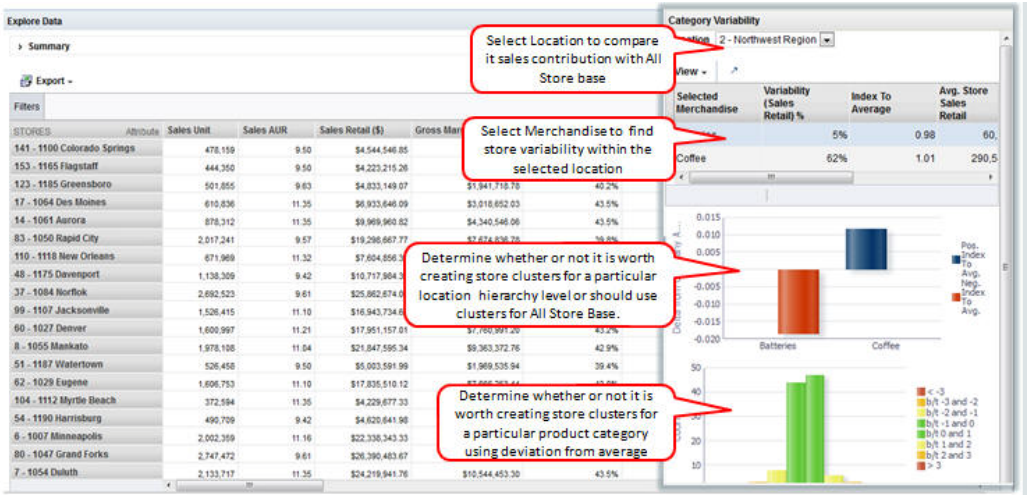
Table 4–11 (Cont.) Category Variability

Property	Description
Variability	The relative standard deviation of the stores in the category. A larger value for the standard deviation indicates greater store variability for the category. Such a category is a possible candidate for store clustering.
Index to average	For a selected location, an indication of how the store performs compared to the all store base. A value close to 1 indicates that the selected location is similar to the all store base. If the value is lower or higher, it indicates that the sales averages for the stores in the selected location are different from the all store base and that you should consider creating store clusters for the selected location.
Average store retail	Average store retail $ for the category for the selected location.
Average store unit	Average store units for the category for the selected location.
Positive/negative index to average	The difference in value for the index to average for the all store base to selected location. For example, a value of 1-index to average < 1 or a value of index to average -1 > 1.

A graph is displayed for the index to average. This shows how the selected location performs compared to the all store base if the average sales metric is below, above, or the same when compared to the all store base average. A red color indicates a value below the all store base average. A blue color indicates a value above the all store base average.

A graph for standard deviation is displayed. This shows the standard deviation for the selected category. If the store value is greater than two standard deviations, then store clustering should be considered for the selected merchandise because the stores sales variability is sufficient.

Figure 4–13 Category Variability

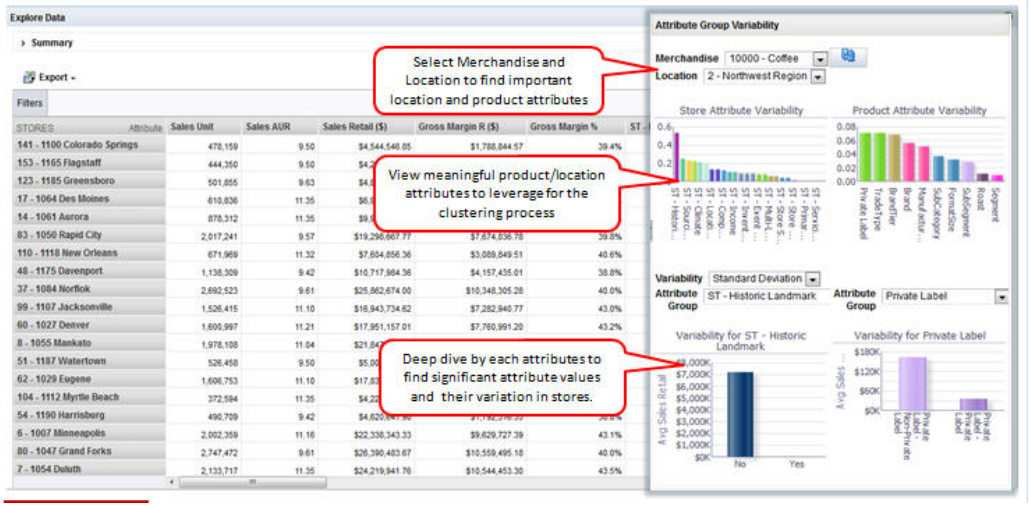


Attribute Group Variability This is used to analyze the most significant attributes (product or location) for a specific merchandise in a selected location.

The attribute group variability section shows attribute graphs for location and product, indicating the key attributes that are driving sales performance.

The attribute variability section, for each attribute group, shows the sales variability for each attribute value in the selected stores. It calculates the index for the attribute group, indicating attribute significance.

Figure 4–14 Attribute Group Variability



Cluster Setup Stage

You can use this stage to perform what-if analysis by defining one or more scenarios that are based on a specified number of clusters and attributes. You can select one or more attributes and assign different weights to the attributes. The attributes and weights you assign are then fed to the clustering analytics to calculate the weighted distance. Using these scenarios, you can experiment with different numbers of clusters, participating attributes, and their weights. You can either define the maximum number or the minimum number of clusters or alternatively define a specific number of clusters that you want to be generated. Once the scenarios are generated, different scenarios can be compared. You can also use other features in this stage to copy or delete scenarios.

Process

Here is the high-level process for setting up scenarios.

1. Either select the name of a scenario you want to modify or enter a name for the new cluster you want to create.
2. If you want the application to optimize the number of clusters, enter minimum and maximum values for the number of clusters.
3. If you want the application to generate a specific number of clusters, enter that value. In this case, the application generates the exact number of clusters and provides the optimal number of clusters as informational data.
4. Optionally, configure the weights assigned to the attributes. The total must add up to 100 percent. Use a value of 0 percent if you do not want a specific attribute to be part of the clustering process.
5. Click the **Execute** icon to execute the scenario. Once the processing is complete, you see the results in the Cluster Results stage.
6. To see a list of all scenarios and the status for each, go to the Scenario List tab.

7. To compare the defining characteristics of two different scenarios, go to the Scenario Compare tab.

Summary

This area lists the criteria you initially selected to define the cluster.

Figure 4–15 Cluster Criteria Summary

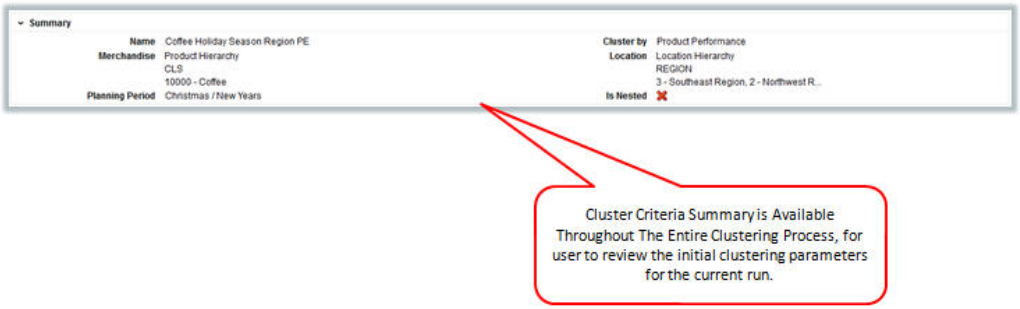


Table 4–12 Cluster Criteria Summary

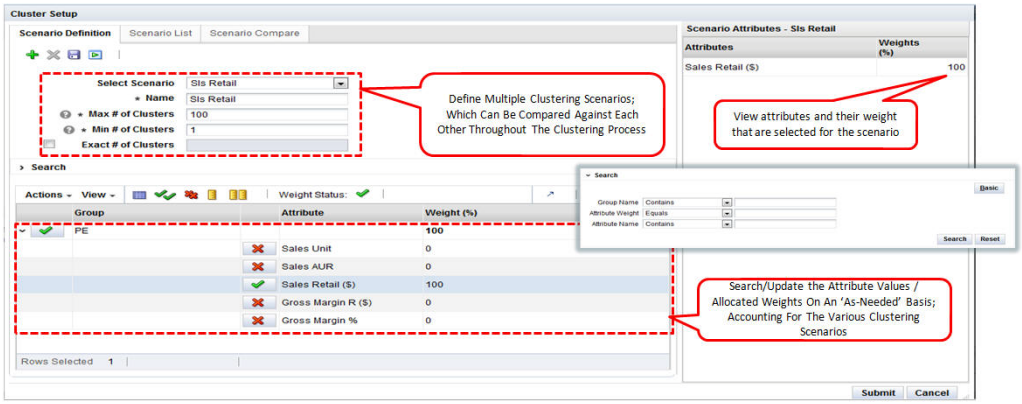
Field	Description
Name	The name you provided for the cluster in the Cluster Criteria stage.
Cluster By	A predefined group of attributes that include Consumer Profile, Product Performance, Store Attribute, Product Attribute, and Mixed Attribute. These criteria types are sets of attributes. For example, store attributes are the properties of a store. These properties can include ethnicity, store format, and store size.
Merchandise	The merchandise level and nodes for the cluster.
Location	The location level and nodes for the cluster.
Fiscal Period	The time period for the cluster.
Is Nested	Indicates whether or not the cluster is nested within another cluster.
Merchandise Hierarchy Type	Provides details about which type of hierarchy the cluster criteria have been created for.

Scenario Definition Section

This area has three tabs: Scenario Definition, Scenario List, and Scenario Compare.

Scenario Definition Tab

Figure 4–16 Clustering Scenario Definition



The following information is needed to define a scenario.

Table 4–13 Scenario Definition

Field Name	Description
Select Scenario	Select an existing scenario if you want to modify it.
Name	A unique name that identifies the scenario being defined.
Max. # of Clusters	Set the maximum number for the total number of clusters that can be generated. The application determines the optimal number of clusters during the generation process.
Min. # of Clusters	Set the minimum number for the total number of clusters that can be generated. The application determines the optimal number of clusters during the generation process.
Exact # of Clusters	Indicates that the exact number of clusters should be generated. The application does not determine the optimal number of clusters.

Attributes

The Attributes table is used to define which attributes are included in the cluster criteria and the weights that should be assigned to each participating attribute. You can

- search by attribute, attribute value, and attribute weight
- assign equal weight to the selected attribute
- assign weight to the selected attribute or attribute value
- reset weights to default values

The following information defines the attributes that are participating or non-participating.

Table 4–14 Attributes

Field Name	Description
Participating	A check in this column indicates that the attributes participates in the cluster criteria.
Groups	Identifies the group.

Table 4–14 (Cont.) Attributes

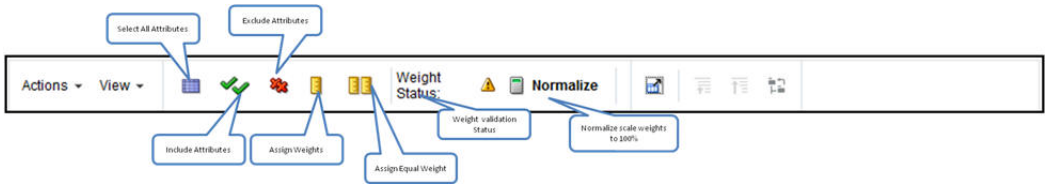
Field Name	Description
Attributes	A description of the attribute.
Weights	The weight assigned to the attribute. All participating attributes can have the same weight or each participating attribute can have a unique weight. The total of all the weights must add up to 100 percent.

The Attributes toolbar, shown in Figure 4–17, includes the following functionality:

Table 4–15 Attribute Toolbar

Function	Description
Action menu	Resets the weights to the default value overrides that the user provided during configuration.
Weight status	Provides the weight validation status. If the weights do not add up to 100 percent, then a warning is displayed and the scenario cannot be executed.
Include or exclude attributes	Any attribute with a weight equal to zero is not included in the clustering process.
Normalize	Scaling attribute weights to ensure that weights are valid. User-provided weights are normalized by applying a weighting average that adds up to 100 percent.

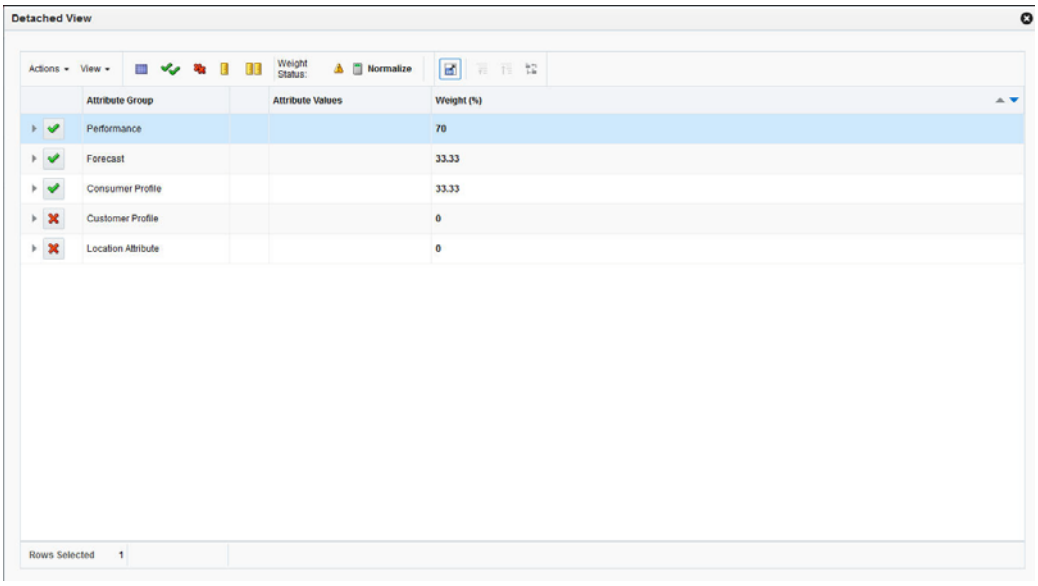
Figure 4–17 Attribute Toolbar



Contextual Area

The contextual business intelligence lists a set of attributes and weights that the current scenario includes as the participating attributes for the clustering process.

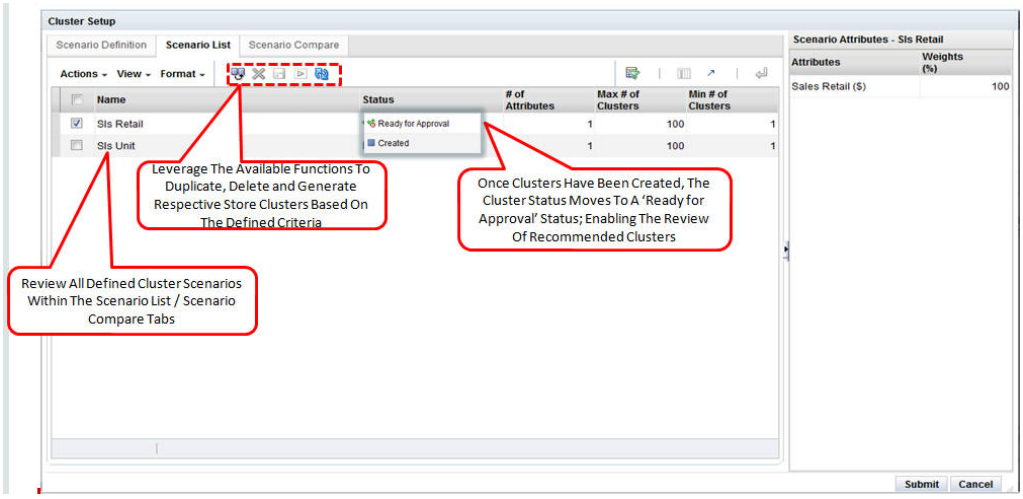
Figure 4–18 Scenario Attributes



Scenario List Tab

The Scenario List summarizes the characteristics for each scenario.

Figure 4–19 Scenario List



You can make a copy of a specific scenario in order to modify it in some way, delete a specific scenario, execute a specific scenario, or save a specific scenario. You can also refresh the scenario list in order to view the updated scenario status.

Table 4–16 Scenario List

Field Name	Description
Name	The unique name that identifies the scenario.
Status	Created, Ready for Approval, Completed with Errors, Approved, Rejected.
# of Attributes	The number of attributes is defined by the Cluster by option you select and the weights you optionally assign.

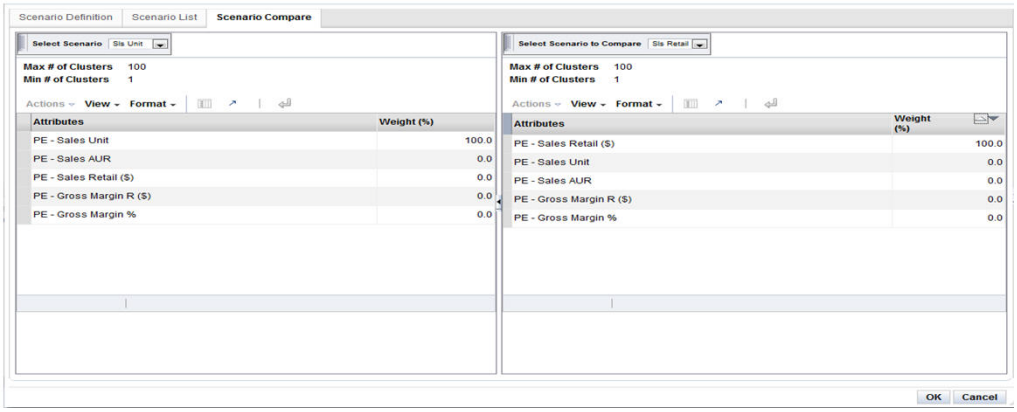
Table 4–16 (Cont.) Scenario List

Field Name	Description
Max. # of Clusters	If you provided a value for this in the scenario definition, that number is displayed here.
Min. # of Clusters	If you provided a value for this in the scenario definition, that number is displayed here.

Scenario Compare Tab

You can select two scenarios from the list to compare. The scenarios you select from the Scenario list are shown side-by-side to facilitate the comparison.

Figure 4–20 Scenario Compare



Cluster Results Stage

After you select a scenario and execute it, you can see the results in this stage. The application uses the data and the parameters you defined in order to group stores together that are most similar according to the characteristics you selected and to separate stores that are most dissimilar. You can also use this stage to rename a cluster.

Process

You use this stage to review clusters, their composition, and the cluster hierarchy, using the grid view and the graph view. This includes

- Reviewing a cluster to see the goodness of fit by using the scores. Determine if any clusters are outliers that warrant further analysis.
- Rank the scenarios (cluster sets) to see how well they are separated and how compact the stores are within each cluster.
- View the optimality of the clusters recommended by the application to determine if increasing the number of clusters beyond the optimal number is significant.
- Rename the cluster after analyzing the centroids and before the cluster is approved.

Summary

This lists the criteria you initially selected to define the cluster.

Figure 4–21 Cluster Criteria Summary

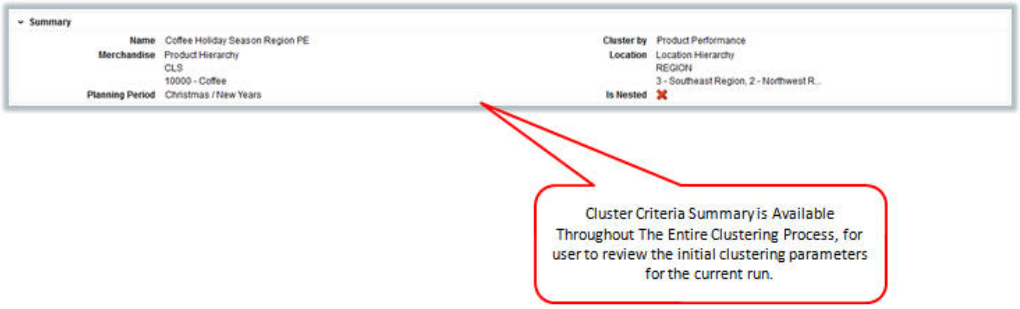


Table 4–17 Cluster Criteria Summary

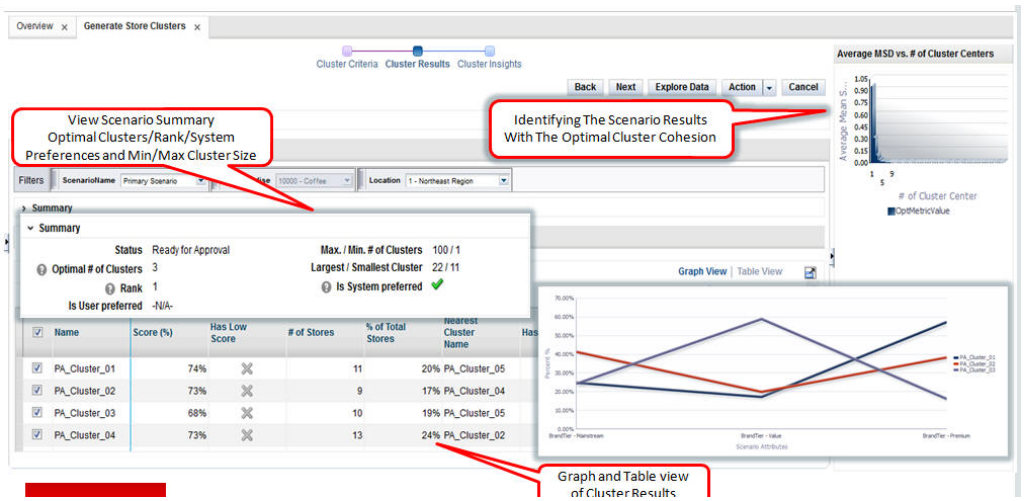
Field	Description
Name	The name you provided for the cluster in the Cluster Criteria stage.
Cluster By	A predefined group of attributes that include Consumer Profile, Product Performance, Store Attribute, Product Attribute, and Mixed Attribute. These criteria types are sets of attributes. For example, store attributes are the properties of a store. These properties can include ethnicity, store format, and store size.
Merchandise	The merchandise level and nodes for the cluster.
Location	The location level and nodes for the cluster.
Fiscal Period	The time period for the cluster.
Is Nested	Indicates whether or not the cluster is nested within another cluster.
Merchandise Hierarchy Type	Provides details about which type of hierarchy the cluster criteria have been created for.

Scenario Results Section

The Scenario Results section displays the following:

- The Scenario Summary, which provides key cluster set attributes for the executed scenario as well as its status.
- The Scenario Results, which has three tabs: Clusters, Cluster Composition, and Cluster Hierarchy.

Figure 4–22 Clustering Scenario Results



The Summary section provides an overview of the characteristics of the clusters.

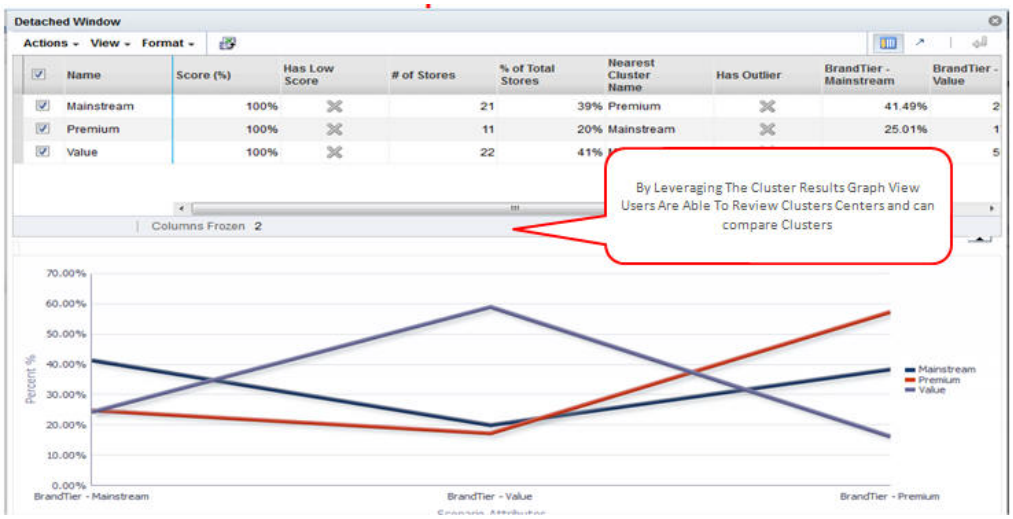
Table 4–18 Scenario Results Summary

Field	Description
Status	Created, Ready for Approval, Completed with Errors, Approved, Rejected.
Optimal # of Clusters	The optimal number of clusters determined by the optimization.
Rank	The application compares executed scenarios and ranks them. A value of 1 indicates the best scenario.
Max./Min # of Clusters	The number you provided for the maximum and minimum number of clusters to calculated.
Largest/Smallest Cluster	Provides the sizes of the largest cluster and the smallest cluster in order to show the range of values.
Is System Preferred	Indicates whether or not the system prefers the scenario.
Is User Preferred	Indicates whether or not the user prefers the scenario.

The Clusters section provides the cluster results for each individual cluster in the scenario in either a Graph View or a Table View. The attributes displayed depend on the Cluster by option chosen in the Cluster Criteria stage.

The Graph View shows the percentage for each attribute in the cluster.

Figure 4–23 Cluster Results Graph View



The Table View provides details that can help you analyze the cluster.

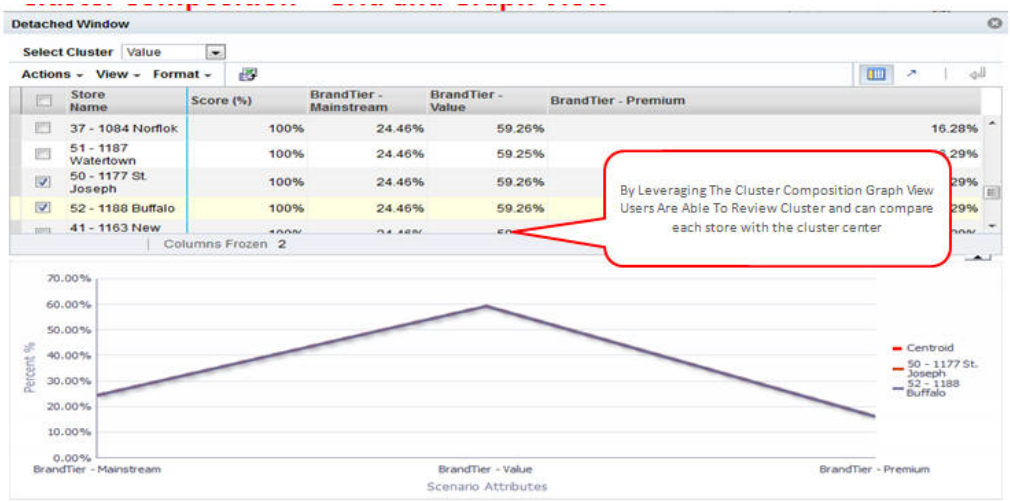
Table 4–19 Scenario Results - Clusters: Table View

Field	Description
Name	The name you assigned to the cluster.
# of Stores	The number of stores in the cluster.
% of Total Stores	The percentage of the total stores that the number of stores represents.
Nearest Cluster Name	The name of the cluster that is most similar to this cluster.
Score %	This value is calculated at the level of store and then averaged to the cluster. The probability, expressed as a percentage, of a store being present in this cluster rather than any of the other clusters.
Has Outlier	Indicates a cluster with the number of stores below a threshold. For example, the number of stores are below certain percentage of the number of stores in a cluster.
Has Low Score	The score threshold can be defined in two ways: The default threshold is calculated as a probability that a store exists in any one of the clusters. The user can further override this threshold at the time of deployment by each Cluster by.
Attributes and their %	For each attribute specific to the Cluster by option, the value indicates the percentage that attribute represents within the total cluster.

The Clusters Composition sub-tab breaks down the cluster into its component parts and shows the percentages for each attribute.

The Table View, shown in Figure 4–24, shows attributes and the score percent.

Figure 4–24 Cluster Composition Table View

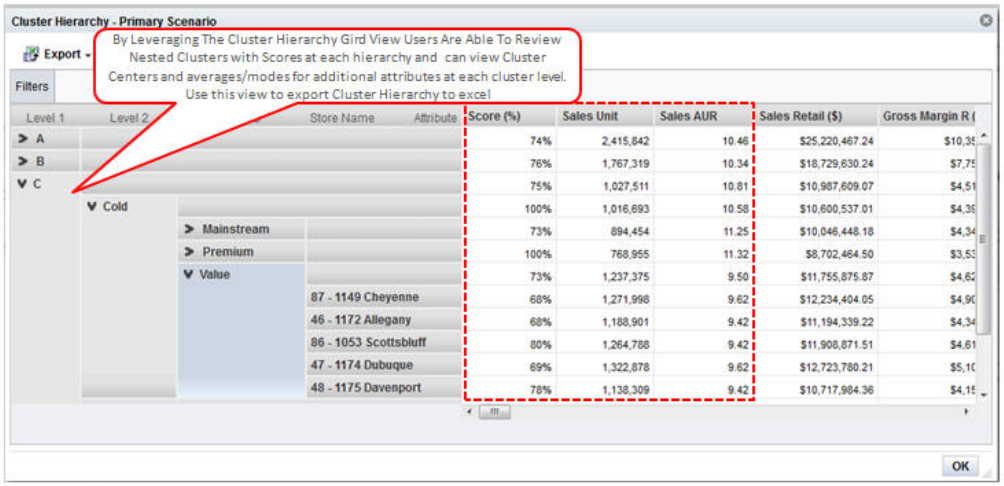


The Graph View shows the centroid of a cluster and allows the user to compare each store with the cluster centroid.

Cluster Hierarchy The cluster hierarchy shows the parent-to-child cluster relationship with the cluster centers. Any additional store attributes are rolled up as averages or modes for each cluster level. You can view the cluster hierarchy for each scenario.

The cluster hierarchy is displayed for each scenario. In the case of nested hierarchies, the cluster hierarchy is displayed in tree format. You can select the attributes to view and export the selected attributes with a cluster hierarchy tree in Excel format.

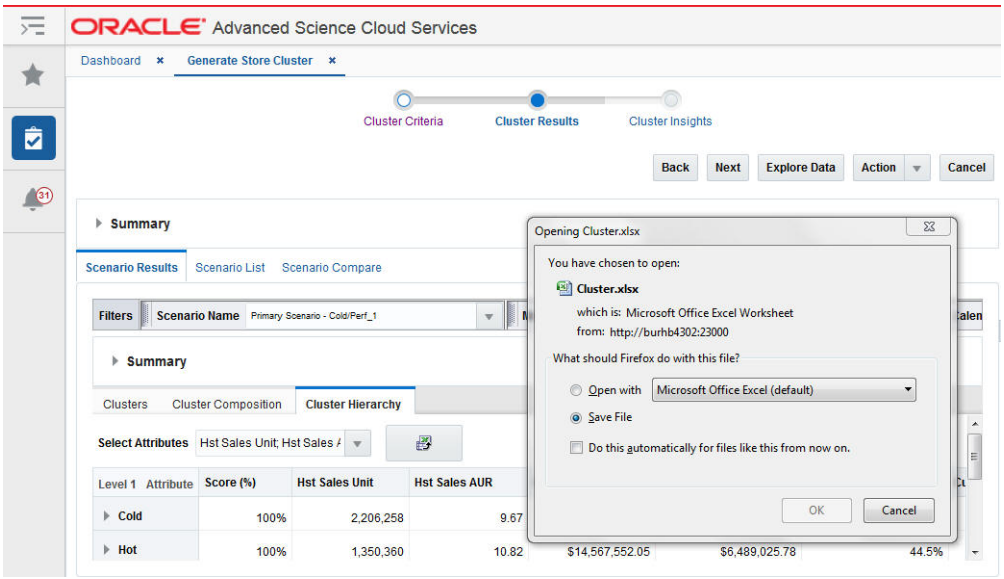
Figure 4–25 Cluster Hierarchy Grid View



Export to Excel

You can select attributes to view data in Pivot and Excel to Export. On selecting icon export to excel user can open or save excel sheet. See excel format below to review cluster hierarchy and their aggregate.

Figure 4–26 Save as Excel File



The Excel-generated file can be then used by external visualization tools to generate graphs not supported in the application.

Figure 4–27 Excel File

	A	B	C	D	E	F	G	H	I	J
1	LEVEL1	LEVEL2	LEVEL3	Store Name	Score (%)	Hst Sales Uni	Hst Sales AUR	Hst Sales Retail	st Gross Margin	t Gross Margin
2	Cold				100%	2,206,258	9.67	$20,994,088.41	$8,488,147.56	40.6%
3	Cold	Cold/Perf_3			81%	571,643	10.13	$5,781,374.03	$2,490,862.27	42.8%
4	Cold	Cold/Perf_4	CP_Cluster_01		100%	571,643	10.13	$5,781,374.03	$2,490,862.27	42.8%
5	Cold	Cold/Perf_5	CP_Cluster_02	68 - 1035 Spokane	100%	299,051	11.01	$3,293,860.63	$1,555,842.29	47.2%
6	Cold	Cold/Perf_6	CP_Cluster_03	22 - 1069 Detroit	100%	323,853	11.07	$3,583,657.79	$1,692,901.38	47.2%
7	Cold	Cold/Perf_7	CP_Cluster_04	51 - 1187 Watertown	100%	606,559	8.71	$5,285,405.82	$2,002,515.98	37.9%
8	Cold	Cold/Perf_8	CP_Cluster_05	76 - 1043 Casper	100%	342,003	10.19	$3,484,737.64	$1,361,867.45	39.1%
9	Cold	Cold/Perf_9	CP_Cluster_06	89 - 1151 Vail	100%	592,923	8.83	$5,237,635.10	$2,019,730.73	38.6%
10	Cold	Cold/Perf_10	CP_Cluster_07	66 - 1033 Great Falls	100%	813,199	10.89	$8,854,491.61	$4,056,034.53	45.8%
11	Cold	Cold/Perf_11	CP_Cluster_08	18 - 1065 Omaha	100%	589,169	11.31	$6,661,449.26	$3,174,290.53	47.7%
12	Cold	Cold/Perf_12	CP_Cluster_09	21 - 1068 Lansing	100%	305,039	10.98	$3,348,926.53	$1,587,049.30	47.4%
13	Cold	Cold/Perf_13	CP_Cluster_10	20 - 1067 South Bend	100%	275,063	10.96	$3,014,701.44	$1,419,485.37	47.1%
14	Cold	Cold/Perf_14	CP_Cluster_11	16 - 1063 Cedar Rapids	100%	709,508	11.33	$8,035,512.33	$3,837,476.49	47.8%
15	Cold	Cold/Perf_15	CP_Cluster_12	91 - 1197 Pierre	100%	607,093	8.75	$5,313,445.67	$2,025,722.48	38.1%
16	Cold	Cold/Perf_16	CP_Cluster_13	14 - 1061 Aurora	100%	701,834	11.32	$7,943,662.79	$3,796,386.44	47.8%
17	Cold	Cold/Perf_17	CP_Cluster_14	52 - 1188 Buffalo	100%	572,589	8.76	$5,015,910.56	$1,906,445.48	38.0%
18	Cold	Cold/Perf_18	CP_Cluster_15	17 - 1064 Des Moines	100%	496,015	11.22	$5,563,602.56	$2,652,263.81	47.7%

Scenario List

The Scenario List section contains one table with details about each cluster. Here, you can approve or reject a cluster.

After a cluster is approved, it is available for other applications.

Figure 4–28 Clustering Scenario List

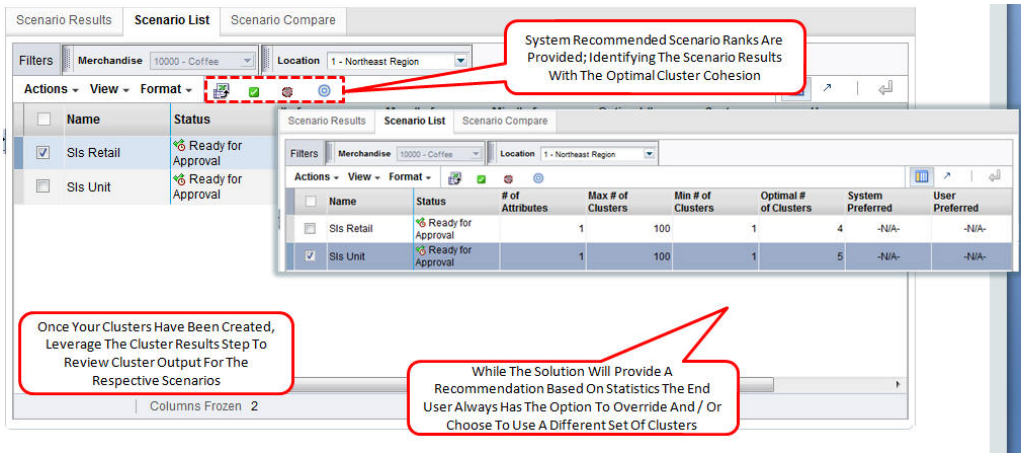


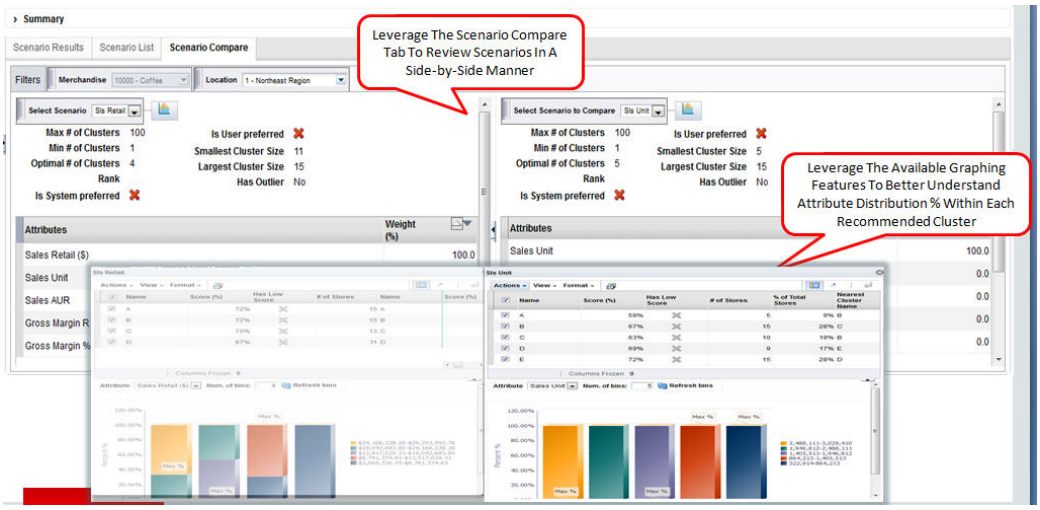
Table 4–20 Scenario List

Field	Description
Name	Name assigned to the scenario.
Status	Created, Ready for Approval, Completed with Errors, Approved, Rejected.
Parent Cluster Level #	The parent cluster name to which the scenario results apply if nested cluster criteria is selected.
# of Attributes	The total number of attributes used, as determined by the weight assigned to each attribute.
Max. # of Clusters	The value used for the maximum in the scenario execution, if this option used.
Min. # of Clusters	The value used for the minimum in the scenario execution, if this option used.
Optimal # of Clusters	The value used for the optimal number of clusters in the scenario execution, if this option used.
System Preferred	Indicates whether the scenario is the one the application prefers.
User Preferred	Indicates whether the scenario is one the user prefers.
Rank Sequence	Indicates the ranking the scenario is given by the application.

Scenario Compare

The Scenario Compare section shows two clusters of your choosing side by side so that you can compare the main results of each, using the same characteristics used in Scenario Results and Scenario List.

Figure 4–29 Compare Scenarios



The information displayed includes:

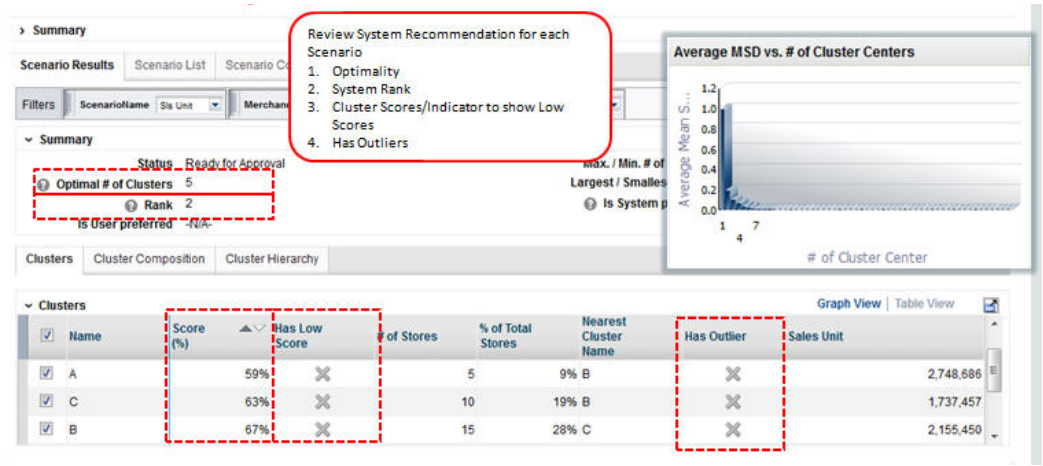
Table 4–21 Scenario Compare

Field Name	Description
Max. # of Clusters	The value used for the maximum in the scenario execution, if this option used.
Min. # of Clusters	The value used for the minimum in the scenario execution, if this option used.
Optimal # of Clusters	The value used for the optimal number of clusters in the scenario execution, if this option used.
Rank	The value for the rank.
Is System Preferred	Indicates whether the scenario is the one the application prefers.
Is User Preferred	Indicates whether the scenario is the one the user prefers.
Smallest Cluster Size	The size of the smallest cluster.
Largest Cluster Size	The size of the largest cluster.
Has Outlier	Indicates a cluster with the number of stores below a threshold. For example, the number of stores are below certain percentage of the number of stores in a cluster.
Attributes	A list of relevant attributes.

Scenario System Recommendations

The application provides the following recommendations at the scenario (cluster set), cluster, and store levels.

Figure 4–30 Scenario Recommendations



Scenario Optimality

This graph indicates how the system identifies the best number of clusters for a given data set. It starts with a small number of cluster centers and searches for the number beyond which there is little improvement in the mean squared distance (MSD). At this point, increasing the number of cluster centers any more only decreases the MSD by a small amount, and the marginal improvement is small.

Scenario Rank

You can see the ranking of all scenarios in the Cluster Results step. The scenario with the highest rank is designated as System Preferred. The ranking is based on the following:

- How many similar stores are contained in the cluster.
- How well separated the clusters are from each other.

Outlier Indicator

This provides an outlier indicator in the cluster list if the cluster has an outlier store. Two outlier rules are supported. The distance from the centroid indicates that if a store is beyond a certain limit for the configuration threshold from the centroid, then the cluster to which the store belongs is marked as an outlier. When the size of a cluster is compared with the total stores, if the number of stores in the cluster is below a certain configured percentage of the total stores, then the cluster is marked as an outlier.

Cluster Scores

The application provides scores for clusters, based on the calculated threshold score. The score is based on the assumption that each store has an equal chance of being a member of the cluster. A high score indicates that the store is close to the centroid. A low score indicates that the store is an outlier.

New Stores and Stores with a Poor History

Advanced Clustering supports post-processing rules in order to allocate stores that are new or that have a poor history. These rules can be configured for each criterion and can be changed during deployment.

- Like Stores. This rule allocates new stores or stores with a poor history to the same clusters that the like location belongs to. It requires data to be provided to

Advanced Clustering that defines the mapping between the location and like locations. This mapping can be configured by merchandise, and one location can be mapped to multiple locations with different weights. For example, a like location can be used to correct a store with poor history or to allocate a new store to a valid performance cluster.

- Largest Clusters. This rule allocates new stores or stores with a poor history to the largest cluster identified by Advanced Clustering. Stores can be allocated to a bigger group of stores. For example, a store that has not yet formed a customer base can be allocated to the largest cluster.
- Cohesive Clusters. This rule allocates new stores or stores with a poor history to the most compact cluster identified by Advanced Clustering. Stores can be allocated to a compact group of stores. For example, stores can be assigned to a cluster that has not been affected because of outliers.

Insights Stage

Use the Insights stage to analyze a scenario, its clusters, and its hierarchy, based on performance and attribute contributions, prior to the approval of the scenario. This stage includes the following tasks:

- Approve a cluster scenario
- Create a new cluster within a scenario
- Rename clusters within a scenario
- Rank scenarios, if not completed earlier
- Flag a cluster scenario as "system preferred"
- Review a cluster hierarchy in a nested cluster

Select from the following views in this stage:

- Criteria view. Displays the parent cluster, if it exists, or the scenario for root-level clusters.
- Parent cluster. Displays any child clusters.
- Cluster. Displays the stores under the selected cluster.

Each cluster is identified by the following information:

Table 4–22 Insights Stage

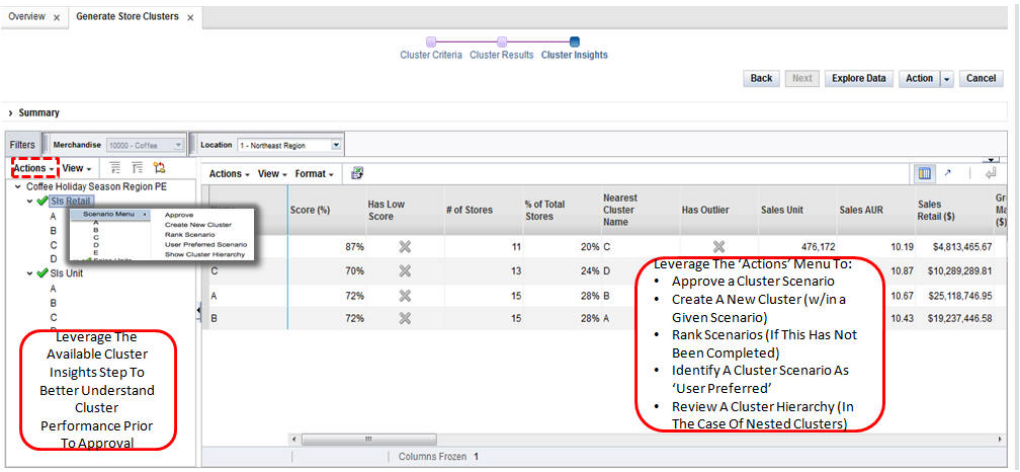
Field Name	Description
Name	The name assigned to the cluster.
Nearest Cluster Name	The name of the cluster that is most similar to the named cluster.
# of Stores	The number of stores in the cluster.
Is Outlier	Whether or not the cluster is considered an outlier. If it is an outlier, you may want to review that store.
All Sales Metric	Financial information about the store.
Score %	This value is calculated at the level of store and then averaged to the cluster. The probability, expressed as a percent, of a store being present in this cluster rather than any of the other clusters.
Has Low Store	Indicates a cluster that falls below a defined threshold.

Table 4–22 (Cont.) Insights Stage

Field Name	Description
% Total Stores	The percentage of the total stores that the number of stores represents.
Sales Retail	The sales revenue for the cluster or store. The merchandise and the source time period are selected when the cluster parameters are set up.
Sales Unit	The sales units for the cluster or store. The merchandise and the source time period are selected when the cluster parameters are set up.
Sales Average Unit Retail	The average unit retail sales for the cluster or store. The merchandise and the source time period are selected when the cluster parameters are set up.
Gross Margin Retail	The retail sales minus the cost of good sold for the cluster or store. The merchandise and the source time period are selected when the cluster parameters are set up.
Gross Margin Percent	The retail sales minus the cost of goods sold divided by the retail sales for the cluster or store. The merchandise and the source time period are selected when the cluster parameters are set up.
Cluster Centers	Additional attributes for each cluster.

You use this stage to view all scenarios, examine and compare sales metrics for the various clusters, and manage clusters by approving, rejecting, merging, or deleting clusters.

Figure 4–31 Clustering Insights



Summary

This lists the criteria you initially selected to define the cluster.

Figure 4–32 Cluster Criteria Summary

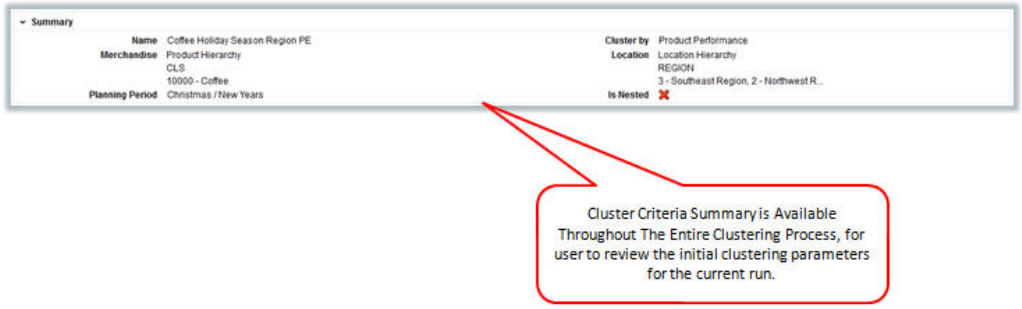


Table 4–23 Cluster Criteria Summary

Field	Description
Name	The name you provided for the cluster in the Cluster Criteria stage.
Cluster By	A predefined group of attributes that include Consumer Profile, Product Performance, Store Attribute, Product Attribute, and Mixed Attribute. These criteria types are sets of attributes. For example, store attributes are the properties of a store. These properties can include ethnicity, store format, and store size.
Merchandise	The merchandise level and nodes for the cluster.
Location	The location level and nodes for the cluster.
Fiscal Period	The time period for the cluster.
Is Nested	Indicates whether or not the cluster is nested within another cluster.
Merchandise Hierarchy Type	Provides details about which type of hierarchy the cluster criteria have been created for.

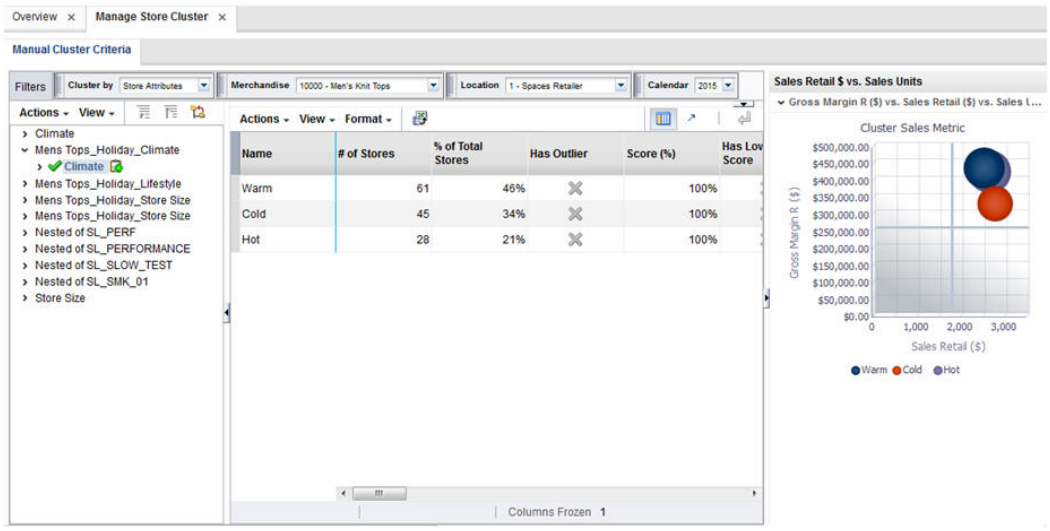
Contextual Information

The graphs displayed in the Clustering Insights stage are also displayed in the Manage Store Clusters tab. For a discussion, see Contextual Information.

Manage Store Clusters Tab

You can use the Manage Store Clusters tab to view a list of the already executed cluster criteria and associated summary details. You can create manual clusters in a scenario, rank/approve/reject scenario, mark a scenario as user preferred, rename clusters, and override any cluster composition within a scenario.

Figure 4–33 Clustering Manage Cluster Criteria



The following information is displayed.

Table 4–24 Manage Store Clusters

Field	Description
Is Nested	Indicates whether or not the cluster is nested within another cluster.
Is Deployed	Indicates that the cluster has been deployed.
Is Shared	A check mark indicates more than one merchandise or location node is used in the cluster.
Scenario Created	The number of scenarios that were created.
Scenario Executed	The number of scenarios that were executed.
Name	The name of the scenario.
Status	Created, Ready for Approval, Completed with Errors, Approved, Rejected.
# Attributes	The total number of attributes used, as determined by the weight assigned to each attribute.
Max # Clusters	The value used for the maximum number of cluster centers in the scenario execution, if this option used.
Min # Clusters	The value used for the minimum number of cluster centers in the scenario execution, if this option used.
Optimal # Clusters	The value used for the optimal number of clusters in the scenario execution, if this option used.
System Preferred	Indicates whether the scenario is the one the application prefers.
Rank Sequence	Indicates the ranking the scenario is given by the application.

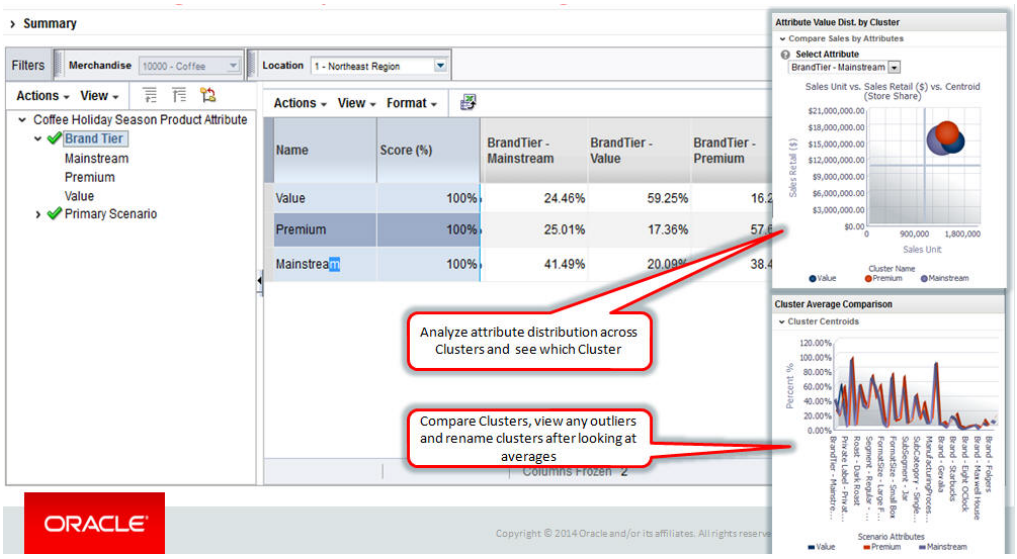
Contextual Information

The following charts are available.

Attribute Value Dist By Cluster

The chart shows a comparison of sales with the average value distribution by clusters. The *y* axis is sales retail and the *x* axis is sales units. The *z* bubble size shows the average value for the selected attribute for the cluster.

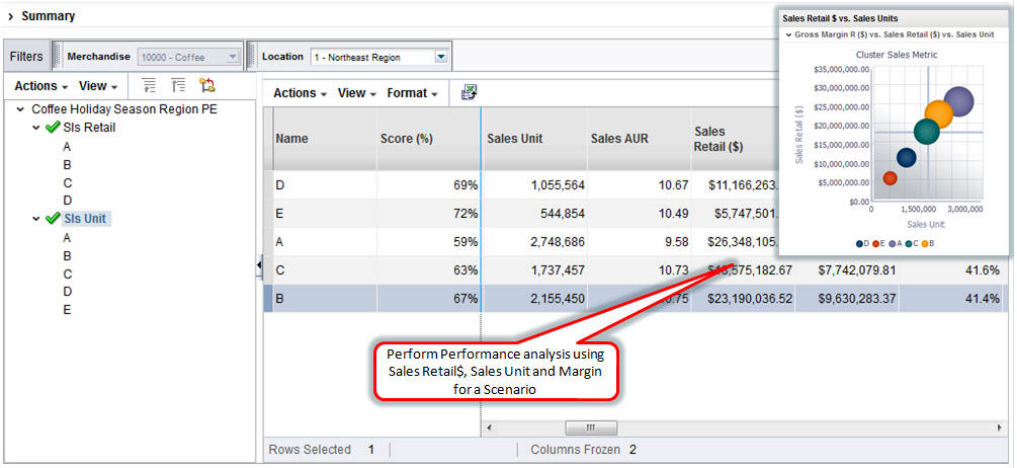
Figure 4–34 Attribute Value Dist. by Cluster



Sales Retail vs. Sales Units

The chart shows a comparison of sales retail and sales unit. The *y* axis is sales retail and the *x* axis is sales unit. The *z* bubble size shows the gross margin retail for a cluster.

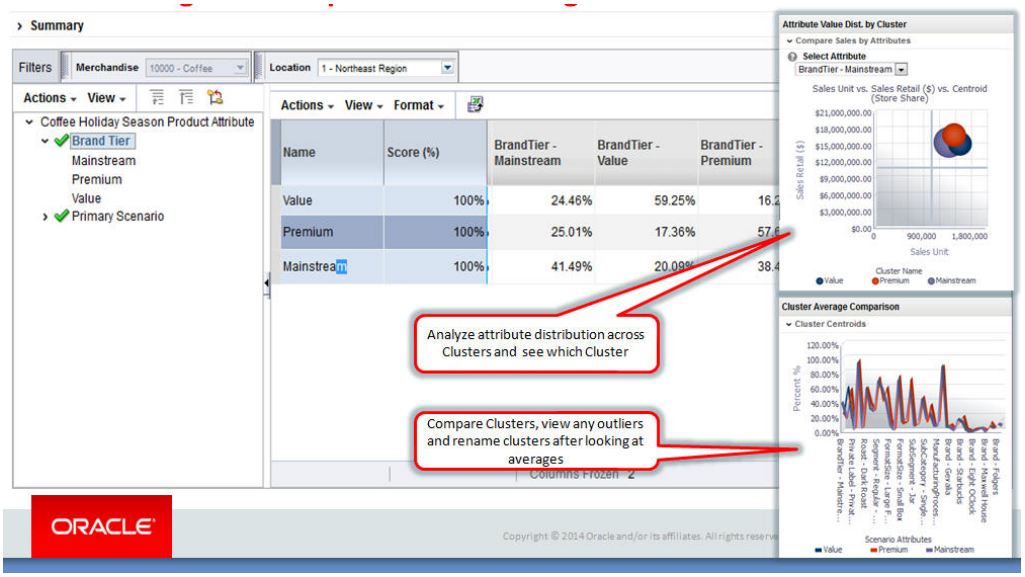
Figure 4–35 Sales Retail vs. Sales Units



Cluster Average Comparison

The chart overlays the centroid of the cluster and provides a comparison of cluster averages so that you can rename the clusters based on the information in the graph.

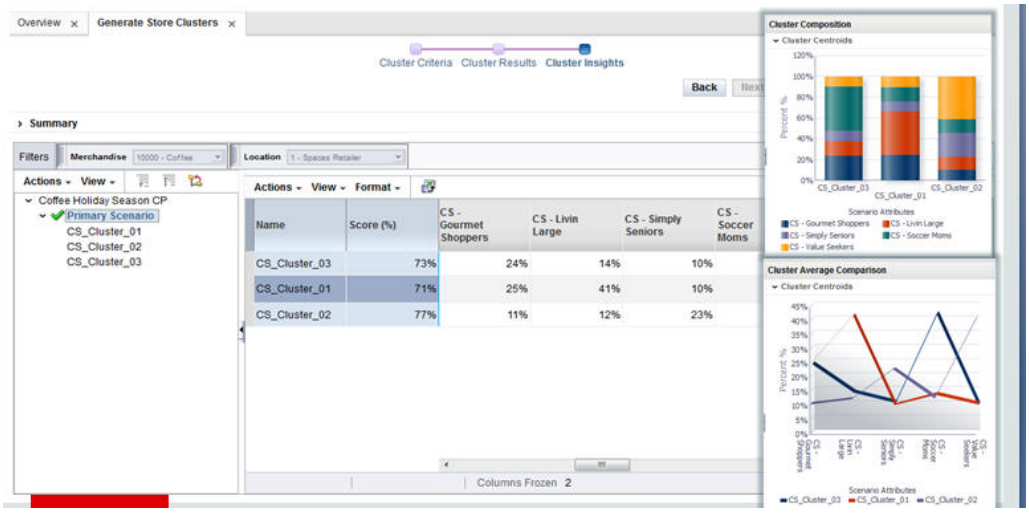
Figure 4–36 Cluster Average Comparison



Cluster Comparison

The chart shows the stacked contribution of each attribute by percentage for each cluster. You can use this chart to determine which cluster contributes the most for an attribute by viewing the clusters side by side.

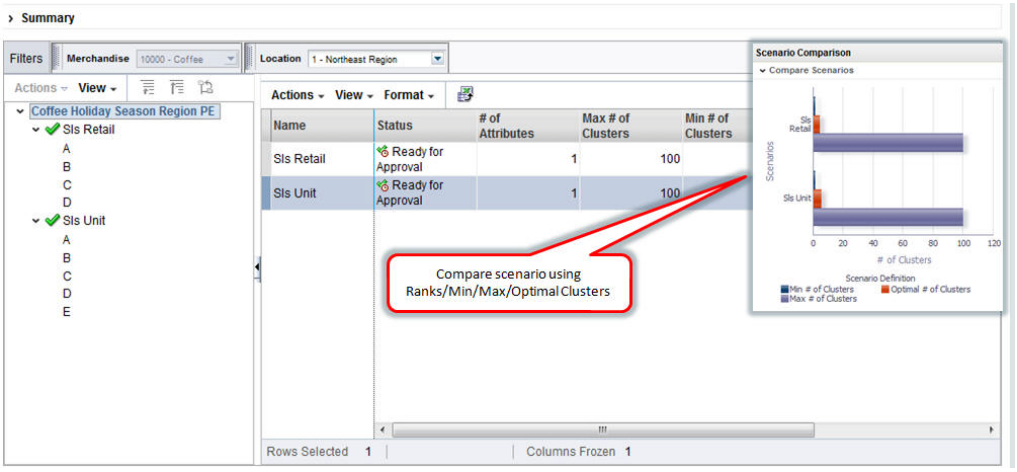
Figure 4–37 Cluster Comparison



Scenario Compare

Multiple scenarios are compared using optimality and scenario rank. This chart is displayed when multiple scenarios for the criteria are available. The following properties are used by the chart: minimum number of clusters, maximum number of clusters, optimal number of clusters, and the rank of the scenarios, from top to bottom.

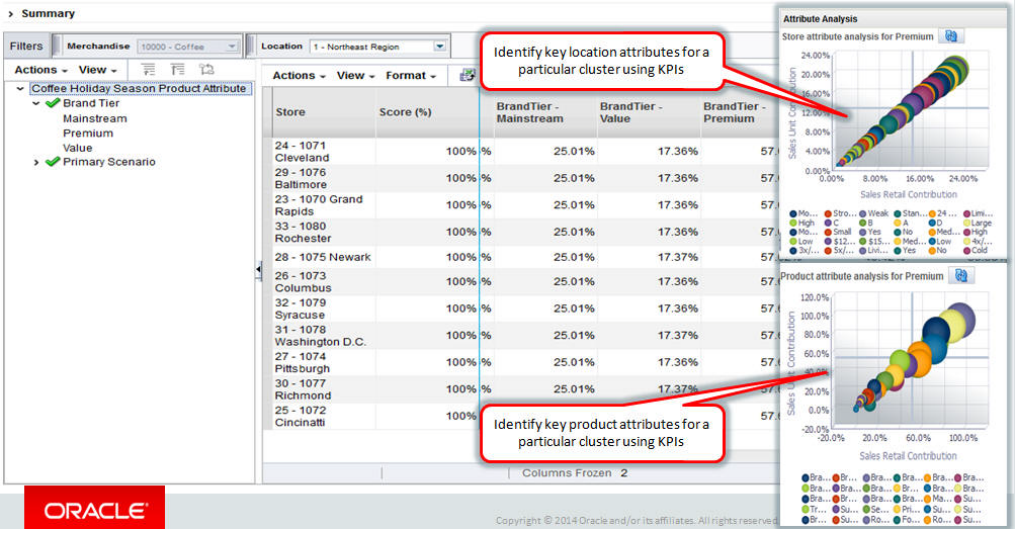
Figure 4–38 Scenario Compare



Attribute Analysis

Stores and product attributes are analyzed and compared to identify the most prominent selling attributes within a cluster. This indicates how the store and product attributes are correlated and what each attribute in the specified cluster contributes to sales. You can make inferences about which attributes in the cluster contribute to significant sales and the potential attributes that should be considered for assortment planning in order to improve sales even more. With both store and product attribute graphs, you can see which location attributes drive the product attribute sales.

Figure 4–39 Attribute Analysis



The following properties are displayed by the graph. Note that the store and product attributes are only displayed when they are configured as part of the Cluster by process.

Table 4–25 Properties

Axis	Description
x-axis	Sales Retail $ contribution, calculated using the sales revenue share of each attribute in the cluster with respect to the total cluster sales revenue.
y-axis	Sales Unit $ contribution, calculated using the sales unit share of each attribute in the cluster with respect to the total cluster sales units.
z-axis	Total Sales Retail $ of each attribute in the selected cluster, indicating, via the bubble, the magnitude of the sales contribution.

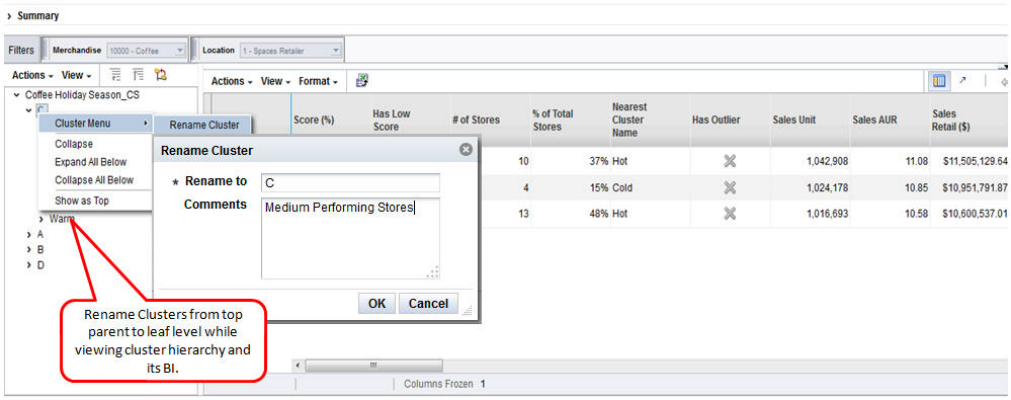
Cluster Overrides

Both Cluster Insights and Manage Store Cluster can be used to override clusters manually. The following tasks, to be completed before the approval process, are available.

Rename Cluster

To complete this on a nested cluster, select the parent or child clusters in the tree. With the contextual information provided, you have the necessary details to understand the cluster across the hierarchy. This task can be performed using the Cluster Contextual menu.

Figure 4–40 Rename Cluster



Create New Manual Cluster

You can create a new manual cluster, describe it, and tag the cluster as Inactive, Flagship, or Manual. This task can be performed using the Scenario Contextual menu.

Inactive

Inactive stores are allocated to these types of clusters. Such stores are either closed, tagged as invalid, or in construction for a specific effective period.

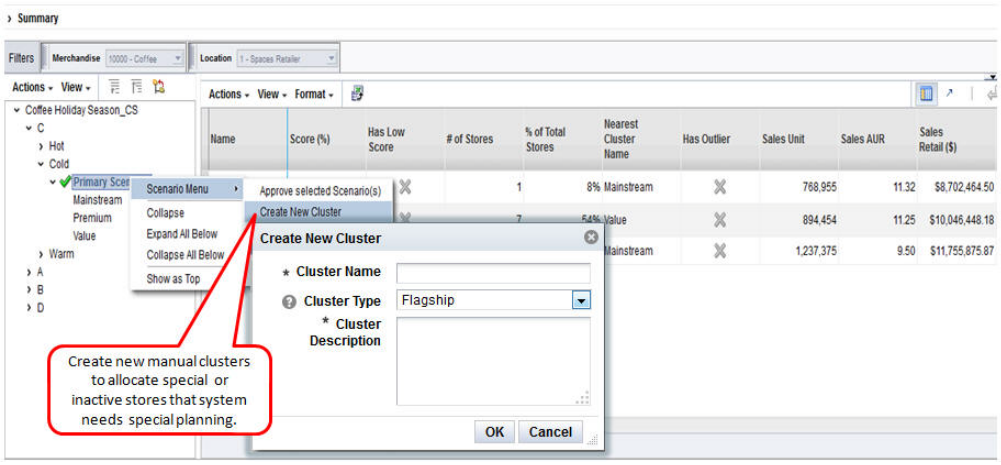
Flagship

Special clusters in which only certain stores reside.

Manually Created

A user manually creates an empty cluster and allocates stores to the cluster using the drag and drop feature. The application automatically re-calculates the cluster centers after the cluster composition changes.

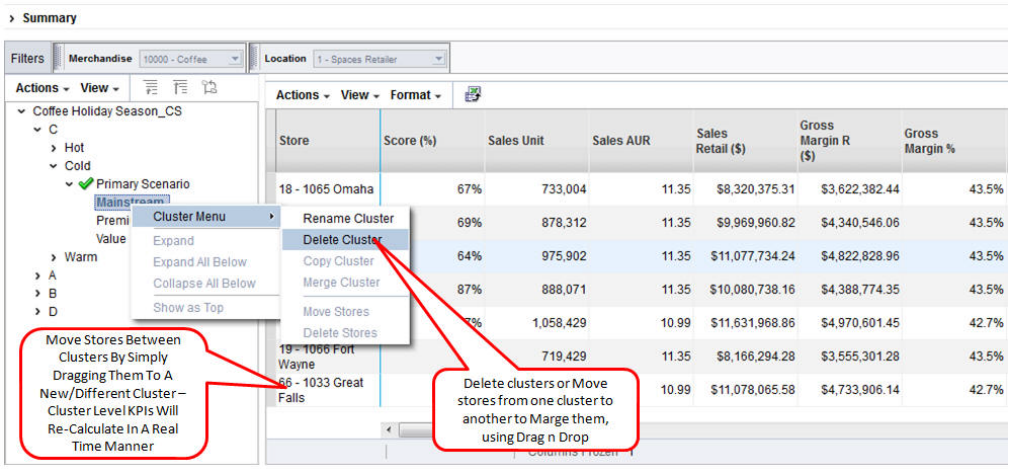
Figure 4–41 Create Manual Cluster



Delete Cluster

You can delete an empty cluster. Prior to deleting the cluster you must move the already allocated stores to another cluster. This help you to merge clusters at the same level. The application automatically re-calculates the cluster centers after the cluster composition changes. This task can be performed using the Cluster Contextual menu.

Figure 4–42 Delete Cluster



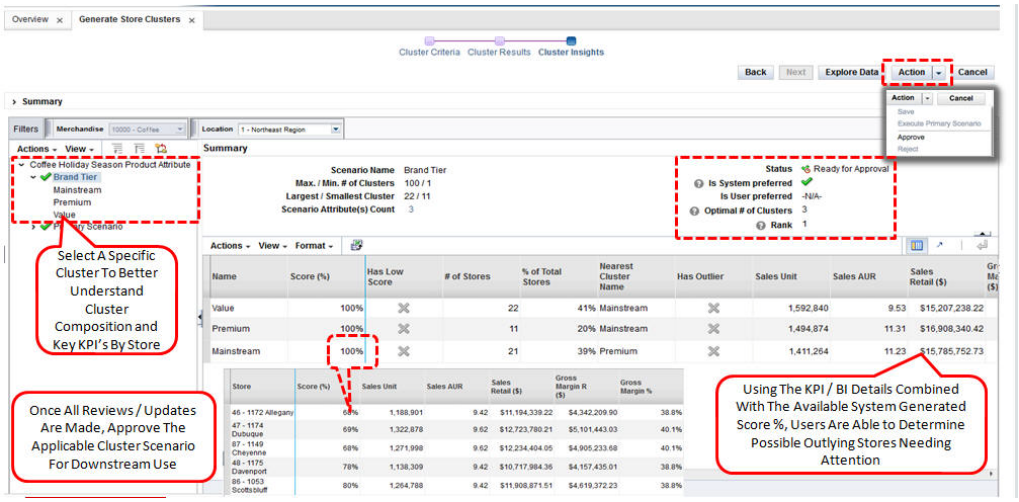
Move Stores to Clusters

Stores can be moved from one cluster to another using drag and drop. The application automatically re-calculates the cluster centers after a cluster composition changes.

Cluster Review, Approve, and Adjust

The approval step is the last step after the review and any manual overrides for the scenario are completed. The scenario results are a set of clusters that are effective for a merchandise, location, and calendar combination.The batch export process selects the last updated approved cluster to deploy to the subscribing applications using an interface file. You can reject an already approved cluster and deploy it if the selected scenario results are better. You are notified of any manual overrides in which approved clusters are modified. They are automatically redeployed.

Figure 4–43 Cluster Review, Approve, and Adjust



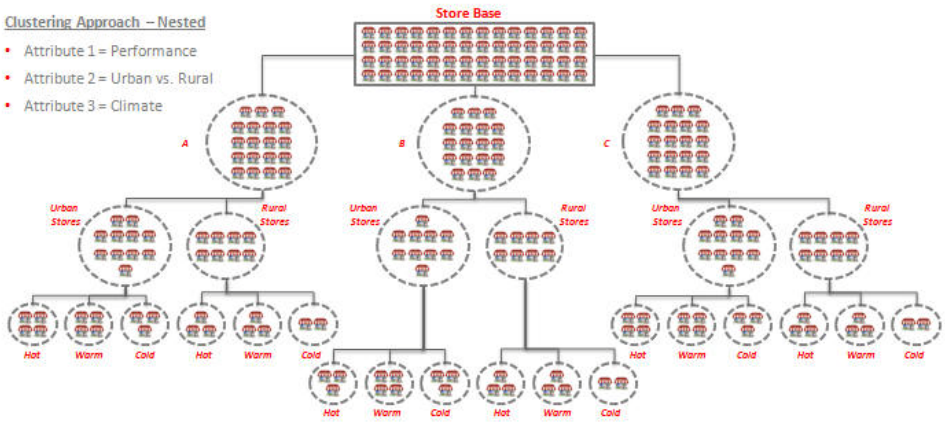
Nested vs. Mixed Attribute

This section describes nested attributes as compared to mixed attributes.

Nested

By default, all Cluster by except mixed attributes can have nested hierarchies. Performance attributes can be further clustered by location attributes, which can be further clustered by location attributes. This approach facilitates dynamic hierarchies in clusters. Nesting can be configured to be enabled or disabled.

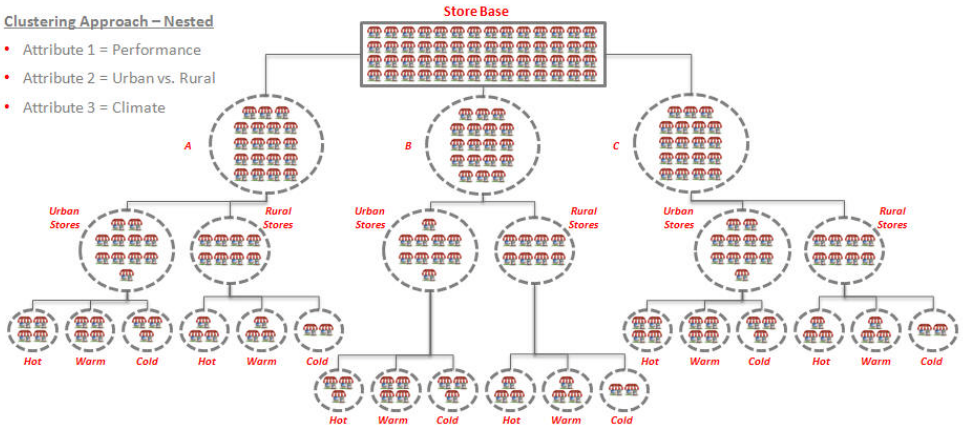
Figure 4–44 Nested Hierarchies



Nested clusters can be created by

- Dynamic hierarchy. Select the Cluster by for each level of a hierarchy while creating a cluster criteria.
- Templates. Select a predefined Cluster by hierarchy while creating a cluster criteria.
- Manual nesting. Create a single cluster criteria, review the results, and then determine whether or not to further cluster. The number of clusters using this approach is granular. The cluster results are hierarchical.

Figure 4–45 Nested Clusters



Mixed Attributes

The following mixed attributes are supported by default: performance, customer segment, location attributes, and product attributes. You can combine attributes from different Cluster by. For example, you can combine attributes from customer segment and performance Cluster by and generate a cluster using sales revenue and customer segment distributions. The number of clusters that are generated using mixed attributes are usually limited, as compared to nested clusters. This approach generates flat clusters with no hierarchy, as attributes are scaled based on the weights you provide.

Figure 4–46 Mixed Attribute Clusters



Figure 4–47 Create Cluster Using Mixed Attributes

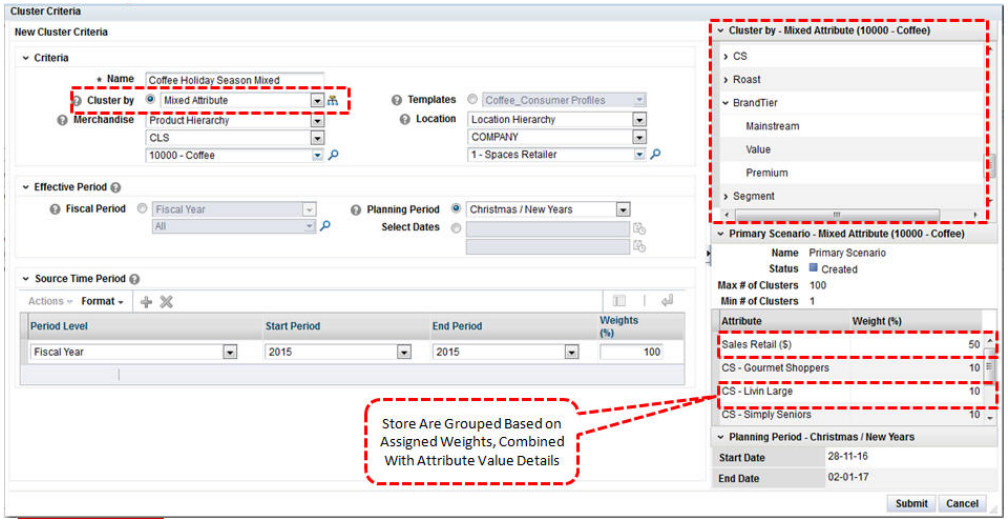
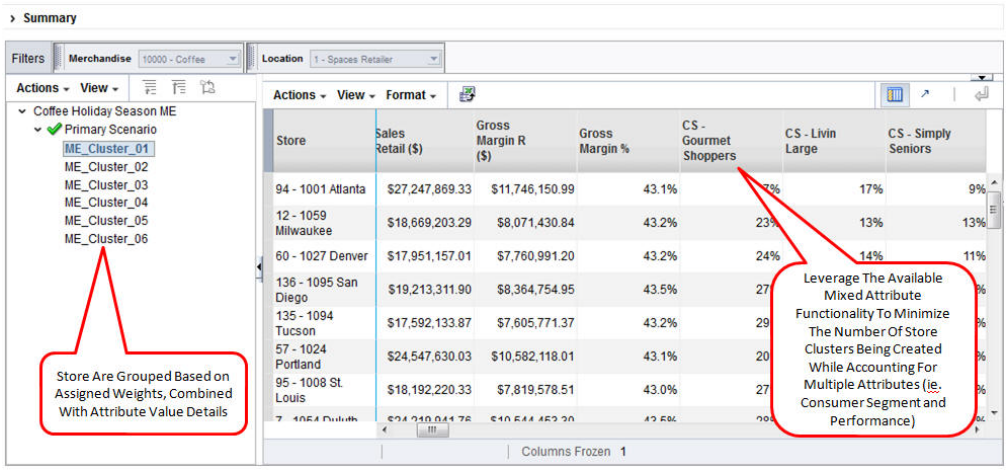


Figure 4–48 Renaming Mixed Attribute Cluster Results



5

Customer Segmentation

This chapter describes the Customer Segmentation module.

Introduction

Customer segmentation is an enterprise-specific solution that uses data mining to group customers based on customer attributes and customer transactions. The retailer can use this information to describe and predict customer behavior. It provides the retailer with a vehicle to target customers with offers, pricing, assortment, and experience.

Retailers understand that shoppers are heterogeneous in nature, that they possess different wants and needs, and that it is impossible to satisfy them all. Retailers can differentiate themselves from their competitors by specializing and offering goods and services that are tailored to one or more market segments.

Customer Segmentation can be used to group customers and to discover hidden customer segments based on the contents of customer shopping baskets and the number of shopping trips they make. Loyalty card data is used to determine if these customer segments differ in terms of socio-demographic or lifestyle characteristics and whether these characteristics can be used to target different customer segments with more relevant product offers.

Retailers can create localized assortments and use customer insights to determine which products to offer by location or channel. This provides insights into the importance of a product to key customer segments and helps when making drop or keep decisions about products within an assortment.

This science-driven tool helps to automate segmentation in a repeatable process and bridges the gap between targeted marketing to targeted assortments.

Features

The key features of Customer Segmentation include:

- Scenario-based segment generation, based on customer attributes, customer behavior, and transactions.
- Attribute importance and correlation mining to identify significant attributes and their associations.
- The ability to generate granular customer segments via the UI. These customer segments include departments such as men, women, and children, health and beauty, or groceries.
- A three-step segment-generation process.

- What-if capabilities that can be used to create multiple segmentation scenarios and then measure them against one another. This can help ensure that the most appropriate segments are used by the applicable planning and execution processes.
- Automatic ranking of segment scenarios to support what-if comparisons.
- Recommendations for the optimal segment scenario and number of segments.
- Descriptions of the main characteristics of the data assigned to each segment using segment rules.
- The use of Oracle data mining hierarchical segmentation and an enhanced version of the k-means algorithm.

Overview of the Customer Segmentation Process

To use Customer Segmentation, follow this general process to create segments, working in the Generate Customer Segment tab:

- Versions – Perform preprocessing actions such as filter and sample a customer's shopping-specific categories or departments. Also identify important attributes that are specific to merchandise and location for customer segmentation.

  Segment Criteria – View all available segments for the specified merchandise, location, and calendar. Review the segment criteria or scenario details for each segment. Use an existing segment as the basis for creating a new segment.
- Explore Data – Examine attributes and their summaries for the segment. Review multiple significant attributes and their correlations using the contextual information.
- Segment Setup – Define multiple what-if scenarios. Such scenarios can be compared with one another throughout the segmentation process.
- Segment Results – View the scenario results and compare scenarios.
- Segment Insights – Gain an understanding about segment results and segment performance prior to approval by examining the contextual information.

Customer Segmentation Overview Tab

The Customer Segmentation Overview tab displays a list of the most recently defined segment criteria and provides the status, segmentation criteria, and applicable merchandise, location, and calendar nodes. You can click on the criteria names in order to access them within the Generate Customer Segments tab.

Figure 5–1 Customer Segmentation Overview Tab

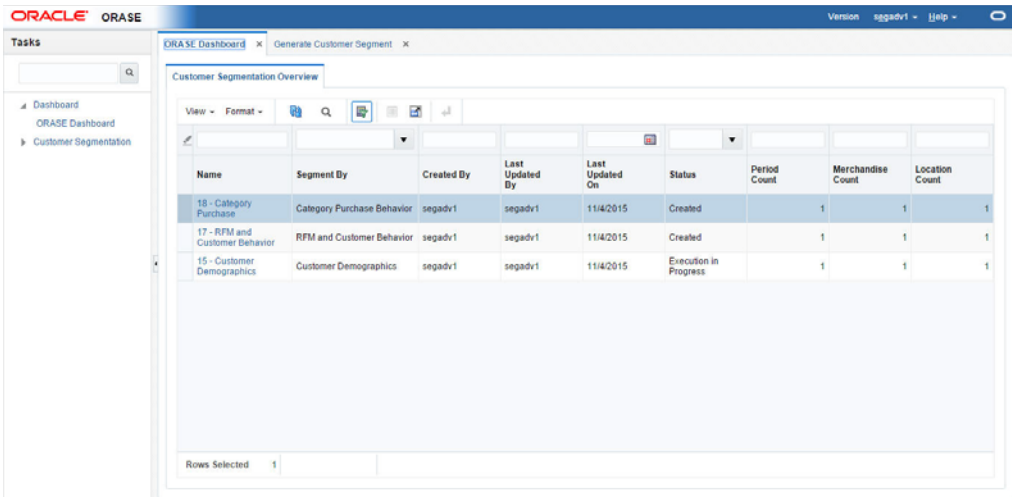


Table 5–1 Cluster Criteria Overview Tab

Field	Description
Name	The criteria ID and user-assigned name of the segment.
Segment By	The Segment By option used for the segment.
Created By	The name of the user who created the segment.
Last Updated By	The name of the user who most recently updated the segment.
Last Updated On	The date when the segment criteria were most recently updated.
Status	The most recent, up-to-date status across the scenarios for the segment criteria. Value include Created, Ready for Preview, Ready for Approval, Approved, and Rejected.
Period Count	The number of calendar nodes defined for the criteria. Hover over the count in order to see a list of the calendar keys associated with the criteria.
Merchandise Count	The number of merchandise nodes defined for the criteria. Hover over the count in order to see a list of the merchandise keys associated with the criteria.
Location Count	The number of location nodes defined for the criteria. Hover over the count in order to see a list of the location keys associated with the criteria.

Segmentation Criteria

The following segmentation criteria are supported by default:

Customer Demographics

This descriptive segmentation technique leverages customer loyalty programs and demographic information (such as residence, profession, age, gender, ethnicity, marital status, and education) about customers to generate demographics based on customer segments.

RFM and Customer Behavior

Segmentation based on purchase behavior aims at discovering groups of customers who exhibit similar purchasing behavior. However, the definition of behavior in this

context includes many factors. For example, retailers may want to distinguish between light and heavy users, regular stock-up shoppers versus emergency top-up shoppers, lunchtime shoppers versus evening shoppers, home and daytime shoppers versus work and weekend shoppers, or fast-checkout customers versus regular checkout customers. The two important behavioral dimensions for understanding customer motivations are visit behavior (identified by the time of day and the day of week that the visits take place) and shopping behavior (identified by the customer's spend dispersion across categories purchased during the trip).

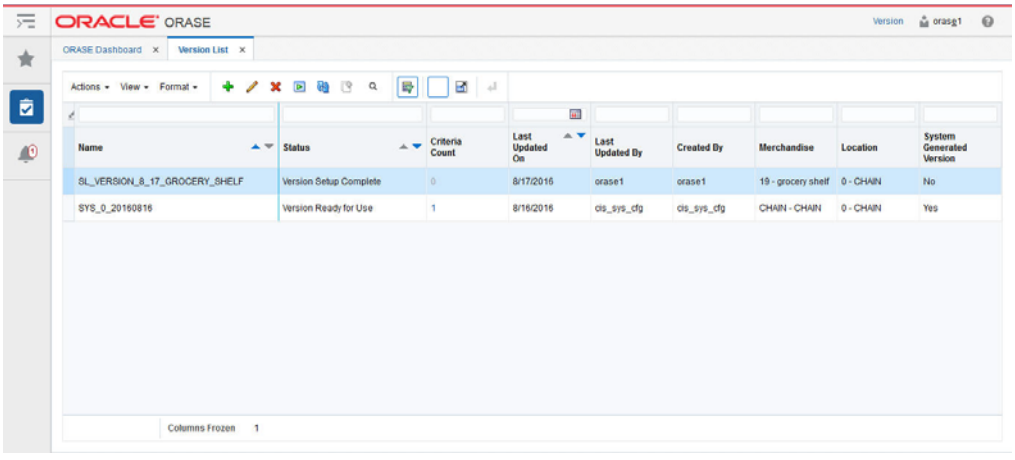
Category Purchase Behavior

Another type of behavior segmentation aims at segmenting the customers who seek similar benefits when evaluating and choosing or purchasing products. These benefits can be measures such as economical price, bulk products, durability, or free shipping. Here, the segment process considers factors that capture customer sensitivity to price and promotions for each category. This can help retailers to segment customers by distinguishing predicted customer responses to the targeted or general promotion of products.

Version List Tab

The Version List tab is used to review, define, and execute preprocessing criteria for customer segments. Version criteria allow users to define versions for merchandise and location combinations. These preprocessed versions can be used to generate customer segments at global or granular levels for merchandise and location. The user can then perform version actions such as creating, updating, and deleting versions and can use selected versions to define new segment criteria.

Figure 5–2 Version List Tab



The Version List, which, by default, lists the global segment version, is defined as part of the batch process. This version is defined for all merchandise and all locations. The version has three key components: filtering, sampling, and attribute mining. It helps when performing the preprocessing of granular customer segments for departments such as men/women/children or segments such as heath and beauty, because customers who are shopping for such categories can differ from customers who are shopping for groceries. These segments can be used by applications such as CMPO for planning at the global level or for promotion planning for granular level segments. The version list is sorted by last update, in descending order. The version status is

sorted in descending order. The name is sorted in ascending order. Searching, using QBE and a search panel option, is provided for both basic and advanced searches.

Version Actions

The following actions are available.

Figure 5–3 Actions



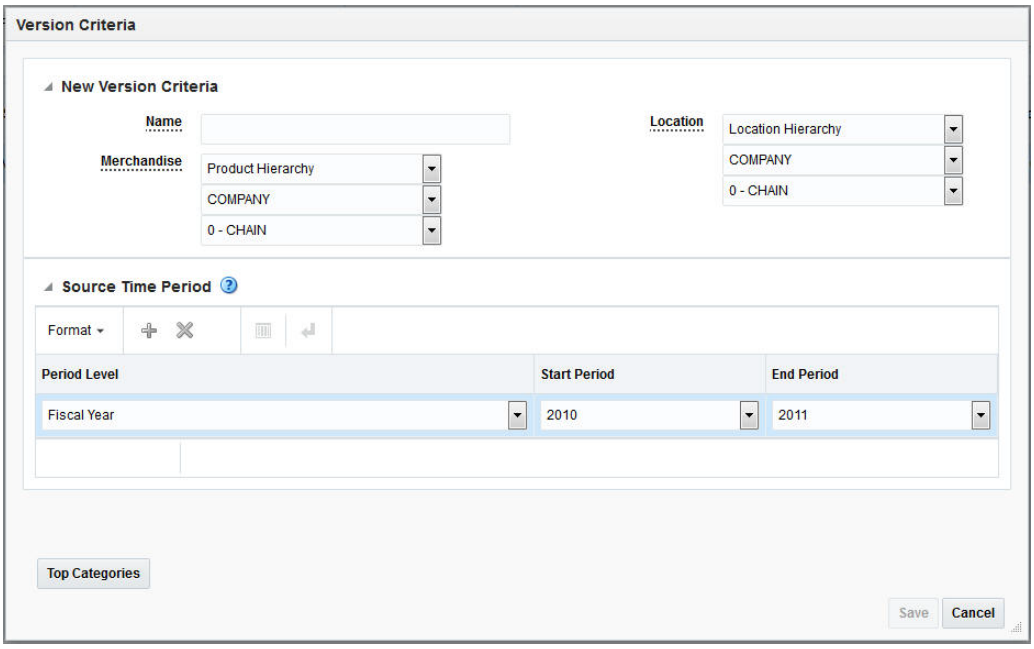
- Create, edit, and delete a version. If a version has a status of Ready for Use and at least one segment is associated with the version, then edit and delete are not available.
- Execute version. This is available when the setup is complete.
- Create a customer segment using an existing version. This opens the Generate Customer Segment tab. The New Segment Criteria pop-up is displayed, and the default values are filled in. (See the defaults that are selected on selection of the version in segment criteria.) This option is only available if the version has a status of Ready for Use.

Table 5–2 Version Details

Name	Description
Name	The name of the version.
Criteria Count	The number of criteria associated with the version. When you click the link, a pop-up is displayed that lists the segment criteria details that are associated with the version.
Last Updated By	The name of the user who updated the version most recently.
Last Updated On	The date when the version was last updated.
Created By	The name of the user who created the version.
Status	The current status of the version. Values include Version Setup Complete, Filtering Completed, Sampling Completed Successfully, Attribute Mining Completed Successfully, Version Ready for Use.
Merchandise	The merchandise node defined for the version.
Location	The location node defined for the version.
System Generated Version	This column provides a flag indicating whether or not the version has been created using a batch process and is set up with a default configuration.

Version Criteria Pop-up

Figure 5–4 Version Criteria Pop-Up



Merchandise

Once you choose the merchandise level for the version, you must select the hierarchy node. The merchandise defined as part of the version is then used for creating segments.

Location

Once you choose the location level for the version, you must select the hierarchy node. These are specific to the location level you select. The location defined as part of the version is then used for creating segments.

Source Time Period

Select historical sales data for the version. You can specify more than one time period. Source time periods selected as part of the version are then applicable for the segments.

Table 5–3 Source Time Period

Time Period	Description
Period Level	Select from Fiscal Year, Fiscal Quarter, Fiscal Period, or Fiscal Week.
Start Period	Once you select the Period Level, select the starting subdivision within that period.
End Period	Once you select the Start Period, select the End Period.

Top Categories

Use this pop-up to select top categories by reviewing category KPIs, such as Sales Share, Revenue, and Margin. The system selects the top categories by each

department. You can override these top categories further while defining versions for merchandise. By default, the top categories in the selected merchandise are driven based on the sales share for each category. You can remove categories to reduce processing time and storage, using the user interface and adding emerging to the top category list.

Version Status

The version status has one of the following values.

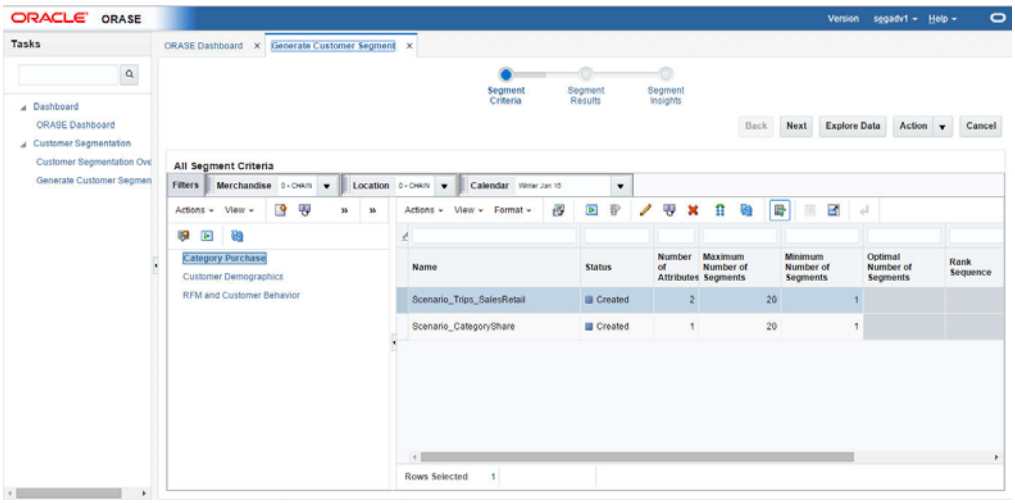
Table 5–4 Version Status

Status	Description
Version Setup Complete	Indicates that version setup is complete.
Filtering Completed with Errors	Indicates that version filtering has completed with errors.
Filtering Execution in Progress	Indicates that version filtering is in progress.
Filtering Completed Successfully	Indicates that version filtering has completed successfully.
Sampling Completed with Errors	Indicates that sampling has completed with errors.
Sampling Execution in Progress	Indicates that sampling execution is in progress.
Sampling Completed Successfully	Indicates that sampling has completed successfully.
Attribute Mining Completed with Errors	Indicates that attribute mining has completed with errors.
Attribute Mining Execution in Progress	Indicates that attribute mining execution is in progress.
Attribute Mining Completed Successfully	Indicates that attribute mining has completed successfully.
Version Ready for Use	Indicates that all three execution phases (filtering, sampling, and attribute mining) have completed successfully. Once the version is ready for use, it is ready for the creation of segment criteria.

Generate Customer Segments Tab

The Generate Customer Segments tab is used to create segments and then model the segments with various scenarios in order to determine the best set of segments. It consists of three stages: Segment Criteria, Segment Results, and Segment Insights.

Figure 5–5 Generate Customer Segments Tab



Segment Criteria Stage

Use the Segment Criteria stage to view summary data about existing segments and define the characteristics of new segments.

Process

Here is the high-level process for defining a segment.

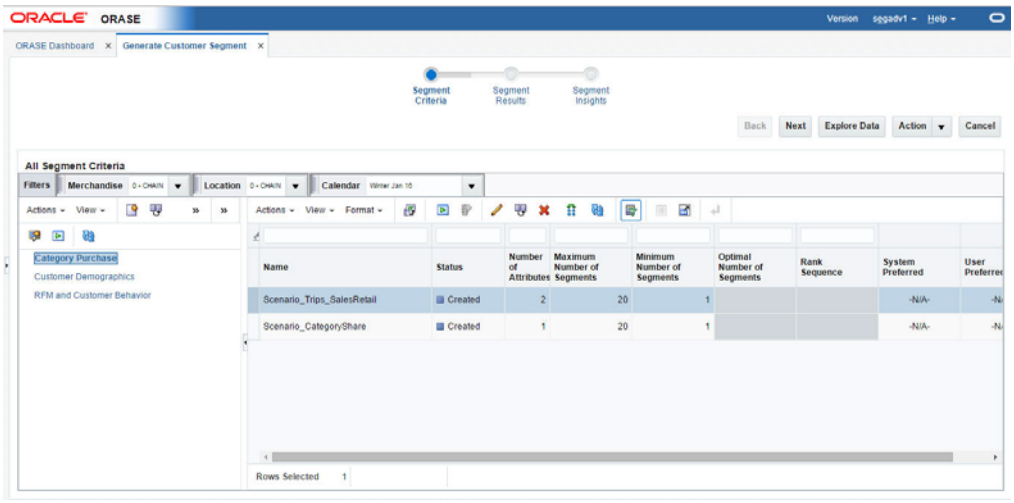
1. Provide a unique name for the segment.
2. Define the type of data used to characterize the segment.
3. Select the merchandise and location nodes.
4. Define the time period for the segment.
5. Define the historical time period for the data.

All Segment Criteria

In this area of the page you can view information about existing segments.

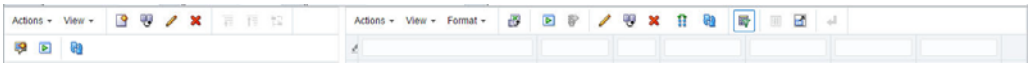
Use the View list to select existing segment criteria. You can tailor your search for existing segment criteria by Merchandise, Location, and Calendar. Once you select the segment criteria, the defining details for the segment criteria are displayed in the Worksheet area.

Figure 5–6 All Segment Criteria



In addition, you can use the toolbar buttons to:

Figure 5–7 Toolbar



- Create segment criteria. You can define an initial segment. The criteria include segment name, the merchandise for the segment, the effective date for the segment, and the history to use.
- View or edit segment criteria. The characteristics of the segment are displayed in the Segment Criteria pop-up.
- Copy segment criteria. Once you have copied the criteria, you can modify them.
- Delete segment criteria. Delete the selected segment criteria.
- Execute segment criteria. Execute all non-executed scenarios for the selected criteria.
- Evaluate KPIs after the execution of the scenario is complete.
- Refresh all. Refresh all segment criteria in order to view any updates to the existing segment criteria.

All Segment Criteria Filter

Once you have highlighted the segment criteria to examine, details about the segment criteria are displayed in a pop-up.

Figure 5–8 All Segment Criteria Filter



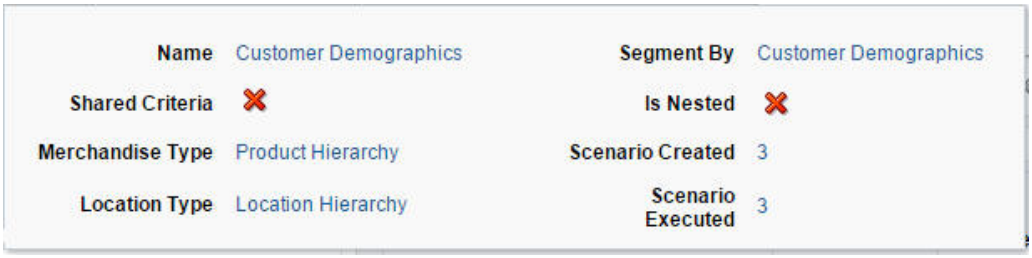
The details include information about the segment criteria and the scenarios created for segment criteria.

- Merchandise allows the user to filter the existing segment criteria by searching for or selecting the merchandise for the supported hierarchies.
- Location allows the user to filter the existing segment criteria by searching for or selecting the location for the supported hierarchies.
- Calendar allows the user to filter the existing segment criteria by searching for or selecting the calendar for the supported hierarchies.

All Segment Criteria Summary

Once you have highlighted the segment criteria to examine, details about the segment criteria are displayed in a pop-up.

Figure 5–9 Segment Criteria Details



The details include information about the segment criteria and the scenarios created for the segment criteria.

Table 5–5 Pop-Up Details

Field Name	Description
Name	The name of the segment criteria.
Shared Criteria	A check mark indicates that more than one merchandise or location node is used in the segment criteria.
Merchandise Type	The merchandise type.
Location Type	The location type.
Segment By	A set of attributes to be used in the creation of the segment criteria: Customer Demographics, RFM and Customer Behavior, Category Purchase Driven.
Scenario Created	The number of scenarios created for the segment.
Scenario Executed	The number of scenarios executed for the segment.

All Segment Criteria Scenario List

This displays the scenarios for the selected segment criteria in the All Segment Criteria tree.

Table 5–6 Scenario List

Field Name	Description
Name	The name assigned to each scenario that has been created for the segment.
Status	Created, Ready for Preview, Ready for Approval, Completed with Errors, Approved, Rejected.
Rank	The system-calculated rank for the segment.

Table 5–6 (Cont.) Scenario List

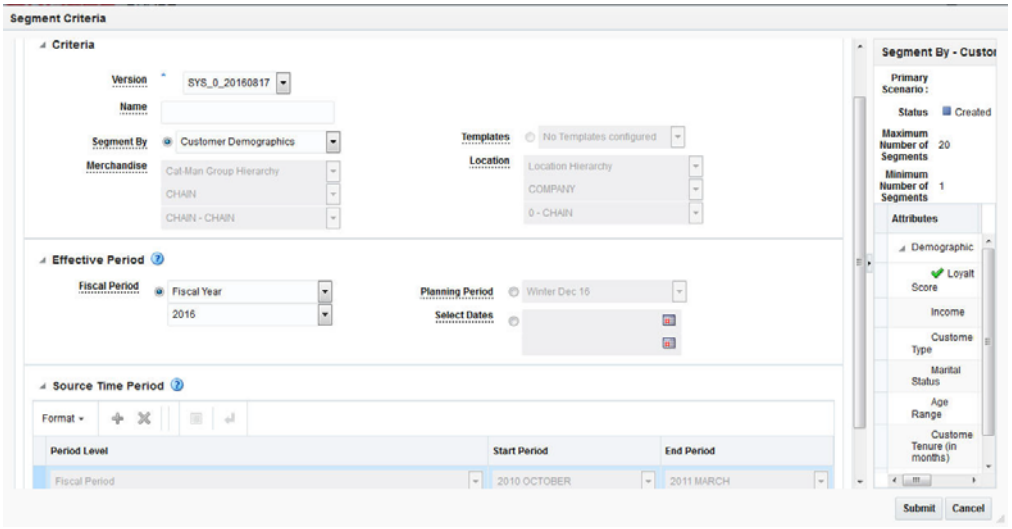
Field Name	Description
Optimal # of Segments	The system-calculated optimal number of segment centers.
User Preferred	Indicates whether or not the user prefers the segment.
System Preferred	Indicates whether or not the system prefers the segment.
# of Attributes	The number of attributes that were used in the segment.
Max. # of Segments	A user-provided value for the maximum segment centers that the segmenting process should consider.
Min. # of Segments	A user-provided value for the minimum segments centers that the segmenting process should consider.

Segment Criteria

In this pop-up, you define the initial segmenting parameters for the segment criteria of a new segment. Note that the hierarchy type supported can be configured at the time of deployment.

Figure 5–10 illustrates how to use a simple approach to segmenting by selecting attributes from a Segment by. For example, you can select the RFM and Customer Behavior Segment by and generate segments using the total number of trips and the amount spent by customer.

Figure 5–10 New Segment Criteria Pop-Up



The following information defines a segment:

Table 5–7 New Segment Definition

Field Name	Description
Name	A unique name to identify the segment.
Version	Select a predefined version to use for creating segment criteria. When the version is selected, the merchandise and location are selected, based on the version.

Table 5–7 (Cont.) New Segment Definition

Field Name	Description
Segment By	A predefined group of attributes that include Customer Demographics, RFM and Customer Behavior, Category Purchase Driven. These criteria types are sets of attributes. For example, customer demographics are the properties of a customer. These properties can include ethnicity, income, and age.
Merchandise	Merchandise is selected based on the version selected for the segment criteria.
Location	Location is selected based on the version selected for the segment criteria.
Template	Template is selected based on the version selected for the segment criteria.

Effective Period

You can define a time interval for the segment by either choosing a period from the list provided or by selecting a start date and an end date.

To define the Effective Period, you select either Planning Period, Fiscal Period, or Select Date.

Table 5–8 Effective Period

Option	Description
Fiscal Period	If you select this option, choose the period and the subdivisions of that period from the drop-down lists.
Planning Period	Select from the range of values provided for the period. Planning periods are user-defined buying periods for a season or a season subset.
Select Dates	If you select this option, choose the start and end dates using the calendar pop-up.

Summarization

Data summarization is available by default and set to either Category or Sub Category. It is applied to Category Purchase Driven Segment by. The segmentation process considers the top selected categories and their attributes and groups customers based on their sales patterns.

Source Time Period

Source time periods are selected based on the version selected for the segment criteria.

Table 5–9 Source Time Period

Field	Description
Period Level	Select from Fiscal Year, Fiscal Quarter, Fiscal Period, or Fiscal Week.
Start Period	Once you select the Period Level, you select the starting subdivision within that period.
End Period	Once you select the Period Level, you select the ending subdivision within that period.

Contextual Area

When you are creating new segment criteria, you can see details about the following parameters that can help you understand the segment you are creating.

Segment By Hierarchy The following information is displayed when you select a template or use the icon to select the Segment by.

Table 5–10 Template Display

Property	Description
Template Name	Name of template configured during deployment.
Description	Description of template.
Segment By	A predefined group of attributes that include Customer Demographics, RFM and Customer Behavior, Category Purchase Driven. These criteria types are sets of attributes. For example, customer demographics are the properties of a customer. These properties can include ethnicity, income, and age.

Segment By Primary Scenario You see this when you select Segment by in the Criteria panel when you are setting the segment criteria parameters or when you select Segment by in the contextual area for the hierarchy.

The system displays the primary scenario, its preconfigured properties, and the significant attributes identified during the attribute importance process for each segment by.

The following information is displayed.

Table 5–11 Primary Scenario

Property	Description
Name	The name of the primary scenario.
Status	Created, Ready for Preview, Ready for Approval, Completed with Errors, Approved, Rejected.
Maximum # segments	The maximum number of segments. The default value is 20. This is used for analyzing the segments.
Minimum # segments	The minimum number of segments. The default value is 1. This is used for analyzing the segments.
Attribute	A list of the attributes configured during segmentation.
Attribute importance	The weighted average across attributes and importance index for each attribute.

Planning Period This list displays the time period you selected for the segment criteria. This information is available only for planning periods, where it provides the start and end dates of the planning period. This content changes whenever the planning period is selected in Effective Period when you are setting segment criteria parameters.

Explore Data

Use the Explore Data pop-up to examine data for the segment criteria you defined. You can view the customer and attribute summary that provides input into the segmentation process.

Process

In this pop-up you can view criteria and attribute summaries as well as their related contextual BIs.

Summary

The Summary lists the criteria you initially selected to define the segment.

Table 5–12 Explore Data: Summary

Field	Description
Name	The name you provided for the segment in the Segment Criteria stage.
Segment By	A predefined group of attributes that include Customer Demographics, RFM and Customer Behavior, Category Purchase Driven. These criteria types are sets of attributes. For example, customer demographics are the properties of a customer. These properties can include ethnicity, income, and age.
Merchandise	The merchandise level and nodes for the segment.
Location	The location level and nodes for the segment.
Fiscal Period	The time period for the segment.
Merchandise Hierarchy Type	Details about which type of hierarchy the segment criteria have been created for.

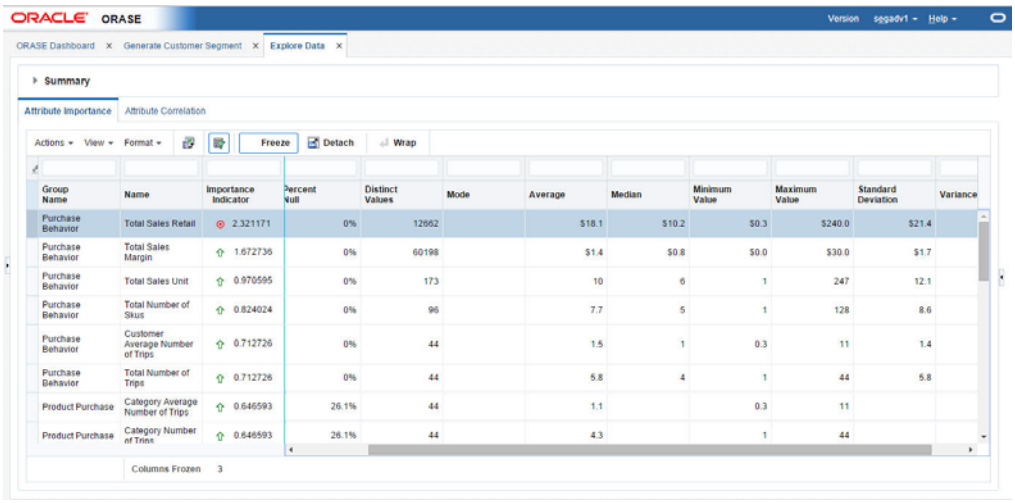
Attribute Mining

This screen provides you with insights about the attribute mining process, which lists the attributes' significance and their correlations. The attributes required for the customer segmentation process come from different dimensions such as customers, their households, demographics, and purchasing behaviors. This process helps to eliminate redundant attributes and to identify the attributes that may have the most influence on generating customer segment. The attribute summary displays information about data availability and data quality by providing the attributes' distinct values, percentage of nulls, and statistics summaries, such as mean, median, and standard deviation.

Attribute Importance

Along with attribute summaries, the system generates an attribute importance index that sums up data quality, data distribution, and its representation of each attribute in the data. See Figure 5–11.

Figure 5–11 Attribute Importance



The system indicates the relative importance of attributes, based on upper and lower threshold values.

Table 5–13 Attribute Importance Values

Category	Range
Low Importance	Minimum value for the lower threshold, indicated in red.
Average Importance	Values between the lower threshold and the upper threshold, indicated in yellow.
High Importance	Values above the upper threshold are indicated in green and are the best candidates for mining attributes.

The following attributes are displayed in Attribute Importance.

Table 5–14 Attributes

Attribute Name	Description
Group Name	The attribute group name (attributes such as demographics, purchase behavior, or product profile).
Name	The name of the attribute (for example, income, ethnicity, or total sales retail)
Histogram	The spark chart for displaying the attribute distribution. This corresponds to the current data distribution graph.
Importance Indicator	The attribute importance index with an image to indicate if attribute has high, average, or low importance.
Percent Null	The percentage of data that is null for the attribute.
Distinct Value	The distinct value applicable to discrete attributes. If no value is available, this is empty.
Mode	The most common value of the discrete attributes. If no value is available, this is empty.
Average	The mean value of the numeric attributes. If no value is available, this is empty.
Median	The median value of the numeric attributes. If no value is available, this is empty.

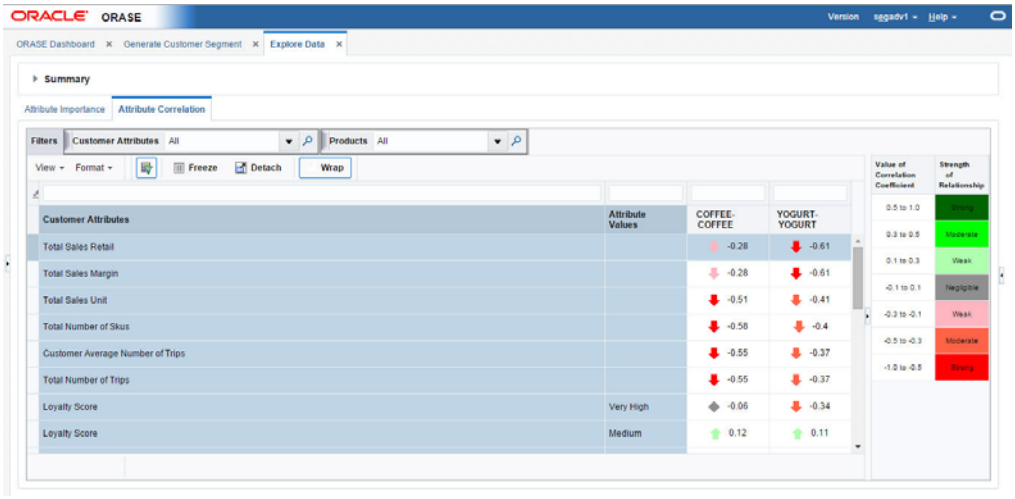
Table 5–14 (Cont.) Attributes

Attribute Name	Description
Minimum Value	The minimum value of the numeric attributes. If no value is available, this is empty.
Maximum Value	The maximum value of the numeric attributes.If no value is available, this is empty.
Standard Deviation	The standard deviation indicates the deviation from the average for the numeric attributes. If no value is available, this is empty.
Variance	The variance indicates the dispersion from the average for the numeric attributes.If no value is available, this is empty.

Attribute Correlation

The attribute mining process generates a correlation between customer attributes and products. This information can help you understand which combination of customer attributes may drive product purchase behavior.

Figure 5–12 Attribute Correlation



Filter

You can use customer and product filters to reduce the number of customer and product attributes that are displayed in the correlation matrix.

Figure 5–13 Customer and Product Filter



Correlation Matrix

The correlation matrix displays attributes, attributes values (in the case of discrete attributes), and correlations between customer and products. Each cell in the matrix provides a visual indication of how attributes are correlated, along with the correlation value. Various colors indicate the strength of the correlation among attributes, with -1 and 1 indicating strong association.

Table 5–15 shows the strength of the correlation for the different ranges of values for the correlation coefficient.

Table 5–15 Correlation Coefficient Values

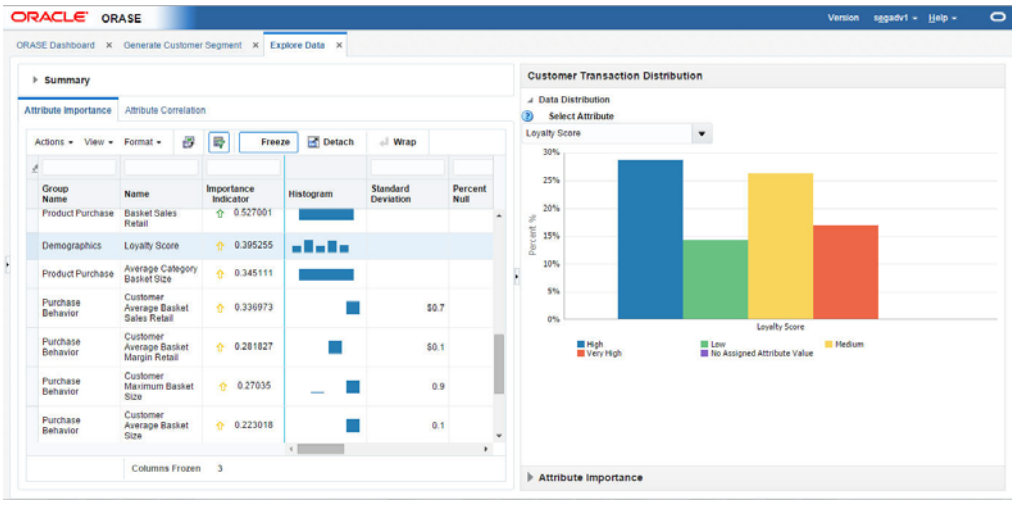
Correlation Coefficient Value	Strength of Relationship
1.0 to 0.5	Strong
0.3 to 0.5	Moderate
0.1 to 0.3	Weak
-0.1 to 0.1	None or very weak
-0.3 to -0.1	Weak
-0.5 to -0.3	Moderate
-1.0 to -0.5	Strong

Contextual Area

This area provides graphical illustrations of the detailed data distribution about the customers and their attribute importance.

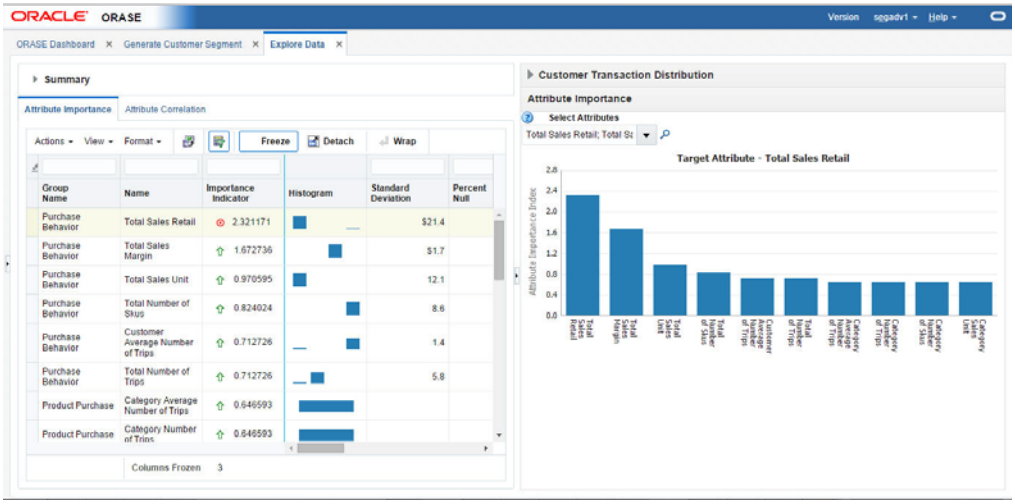
Analyze Customers In Explore Data, the BI displays the data distribution of the customers by each participating attribute as well as other configured informational attributes. Customer Segmentation identifies the bins based on the underlying sample data and displays the histograms. It provides the percentage of customers that are present in a selected location. For example, a company may have ten percent of premium customers who are high spenders and who shop frequently.

Figure 5–14 Analyze Customers



Attribute Importance The attributes are ranked in order of their significance in predicting a target (the default is set to Customer Total Spent) and by analyzing the quality and variability of the data. The results in Figure 5–15 shows that loyalty score and income have the most effect on whether or not the customer spends more shopping at the retailer.

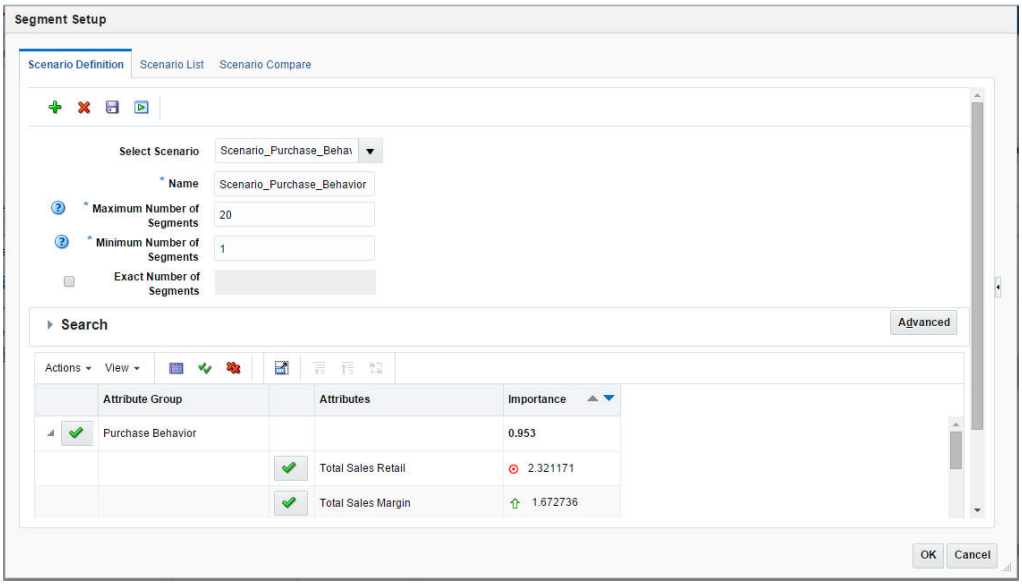
Figure 5–15 Attribute Importance



Segment Setup Stage

You can use this stage to perform what-if analysis by defining one or more scenarios that are based on a specified number of segments and attributes. You can select one or more attributes based on the attribute importance index. The attribute importance index guides you when you are defining the scenario setup. Using these scenarios, you can experiment with different numbers of segments and participating attributes. You can either define the maximum number or the minimum number of segments or alternatively define a specific number of segments that you want to be generated. Once the scenarios are generated, different scenarios can be compared. You can also use other features in this stage to copy or delete scenarios.

Figure 5–16 Segment Setup



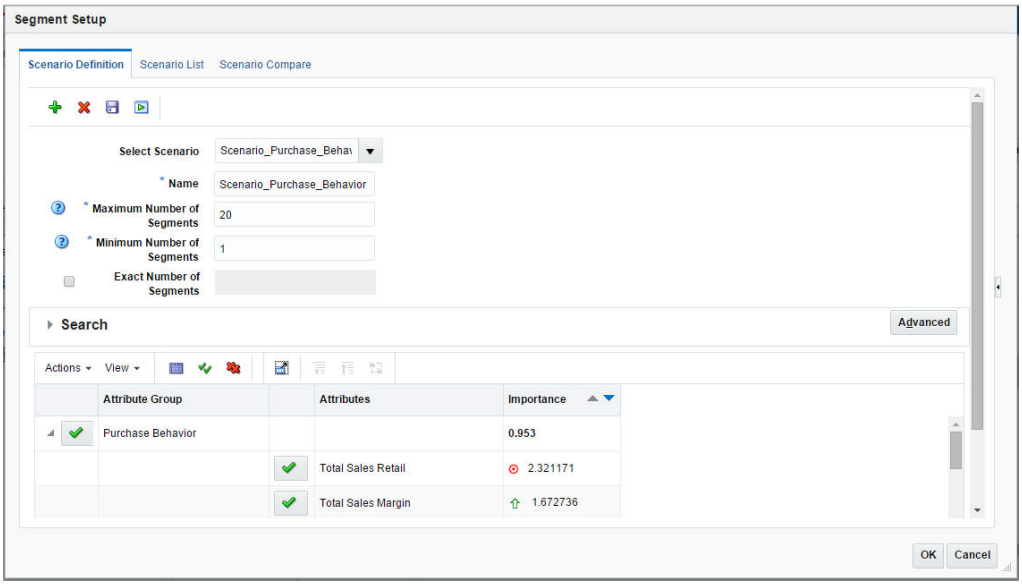
Process

Here is the high-level process for setting up scenarios.

1. Either select the name of a scenario you want to modify or enter the name for the new segment criteria you want to create.
2. If you want the application to optimize the number of segments, enter the minimum and maximum values for the number of segments.
3. If you want the application to generate a specific number of segments, enter that value. In this case, the application generates the exact number of segments and provides the optimal number of segments as informational data.
4. Optionally, review the attribute importance assigned to the attributes and select an attribute if you want a specific attribute to be part of the segmentation process.
5. Click the **Execute** icon to execute the scenario. Once the processing is complete, you see the results in the Segment Results stage.
6. To see a list of all scenarios and the status for each, go to the Scenario List tab.

Scenario Definition Tab

Figure 5–17 Scenario Definition Tab



The following information defines a scenario.

Table 5–16 Scenario Definition

Field Name	Description
Select Scenario	Select an existing scenario if you want to modify it.
Name	A unique name that identifies the scenario being defined.
Max. # of Segments	Set the maximum number for the total number of segments that can be generated. The application determines the optimal number of segments during the generation process.
Min. # of Segments	Set the minimum number for the total number of segments that can be generated. The application determines the optimal number of segments during the generation process.
Exact # of Segments	Indicates that the exact number of segments should be generated. The application does not determine the optimal number of segments.

Attributes

The Attributes table is used to define which attributes are included in the segment criteria. You can

- search by attribute, attribute value, or attribute importance
- select and deselect attributes

The following information defines the attributes that are participating or non-participating.

Table 5–17 Attributes

Field Name	Description
Participating	A check in this column indicates that the attribute participates in the segment criteria.
Attribute Group	A logical grouping of attributes such as demographics or purchase behavior.
Attributes	Attributes that are potential candidates for generating segments.
Importance	System-generated attribute importance index that indicates the significance of each attribute.

The Attributes toolbar includes the following functionality:

Figure 5–18 Attribute Toolbar



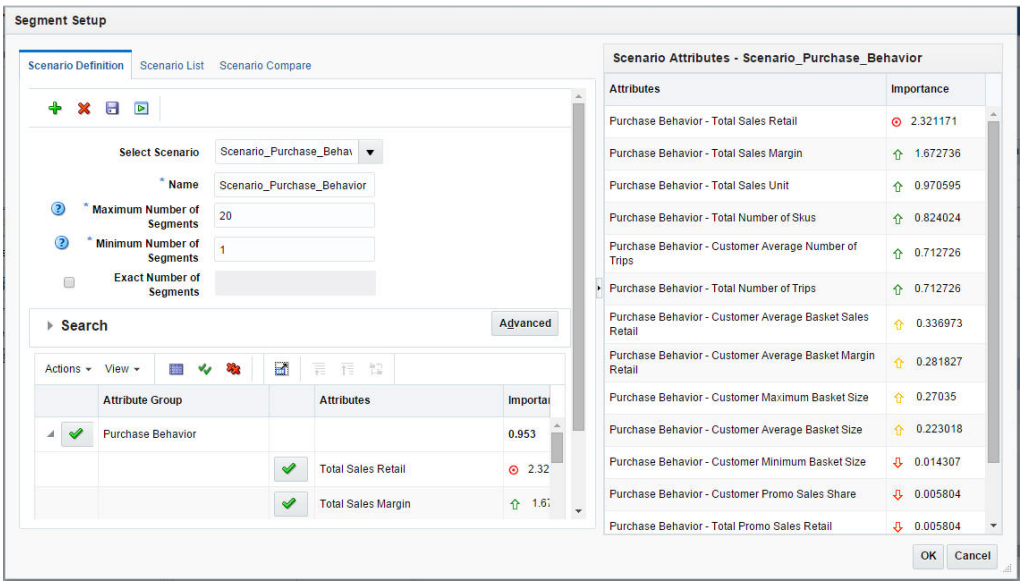
Table 5–18 Attribute Toolbar

Function	Description
Action menu	Resets the attribute selection to the default selection that system identified using attribute importance thresholds.
Include or exclude attributes	Any attribute beyond a certain threshold is not included in the segmentation process.

Contextual Area

The contextual business intelligence lists a set of attributes that the current scenario includes as the participating attributes for the segmentation process.

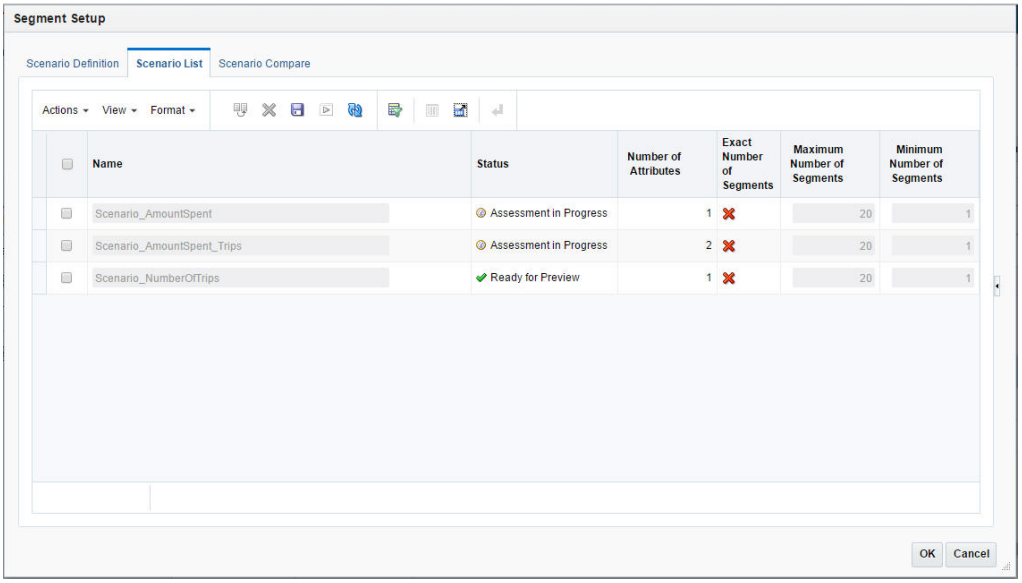
Figure 5–19 Scenario Definition: Contextual Area



Scenario List Tab

The Scenario List summarizes the characteristics for each scenario.

Figure 5–20 Scenario List Tab



You can make a copy of a specific scenario in order to modify it in some way, delete a specific scenario, execute a specific scenario, or save a specific scenario. You can also refresh the scenario list in order to view the updated scenario status.

Table 5–19 Scenario List

Field Name	Description
Name	The unique name that identifies the scenario.
Status	These include Created, Ready for Preview, Ready for Approval, Completed with Errors, Approved, Rejected.

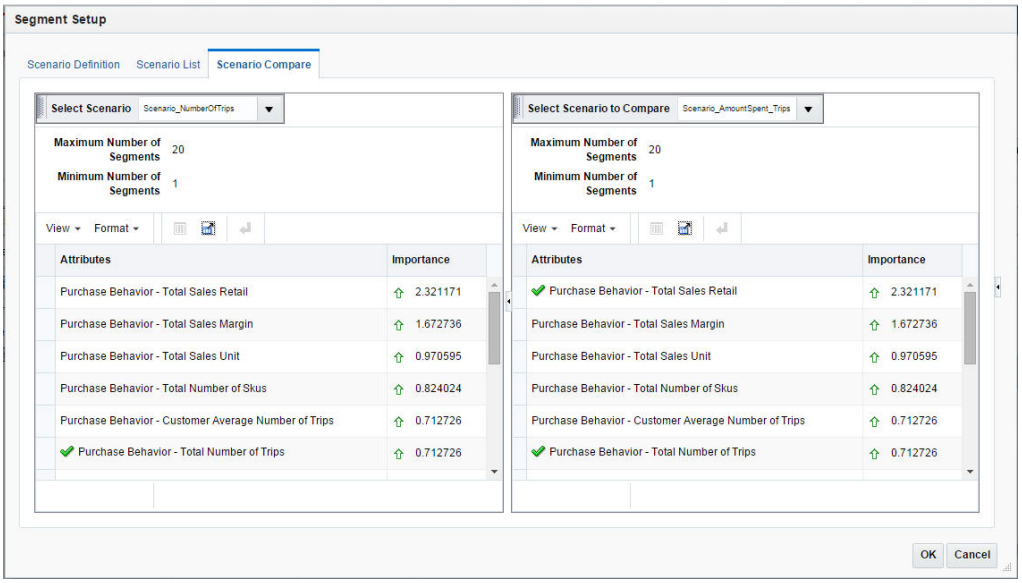
Table 5–19 (Cont.) Scenario List

Field Name	Description
# of Attributes	The number of attributes is defined by the Segment by option you select.
Max. # of Segments	If you provided a value for this in the scenario definition, that number is displayed here.
Min. # of Segments	If you provided a value for this in the scenario definition, that number is displayed here.

Scenario Compare Tab

You can select two scenarios from the list to compare. The scenarios you select from the Scenario list are shown side-by-side to facilitate this.

Figure 5–21 Scenario Compare Tab



Segment Results Stage

After you select a scenario and execute it, you can see the results in this stage. The application uses the data and the parameters you defined in order to group customers together that are most similar according to the characteristics you selected and to separate customers that are most dissimilar. You can also use this stage to rename a segment.

Process

You use this stage to review segments and their composition, using the grid view and the graph view. This includes

- Review a segment to see the goodness of fit by using the scores. Determine if any segments are outliers that warrant further analysis.
- Rank the scenarios (segment sets) to see how well they are separated and how compact the customers are within each segment.

- View the optimality of the segments recommended by the application to determine if increasing the number of segments beyond the optimal number is significant.
- Rename the segment after analyzing the centroids and before the segment is approved.

Summary

This lists the criteria you initially selected to define the segment.

Figure 5–22 Results Summary



Table 5–20 Segment Criteria Summary

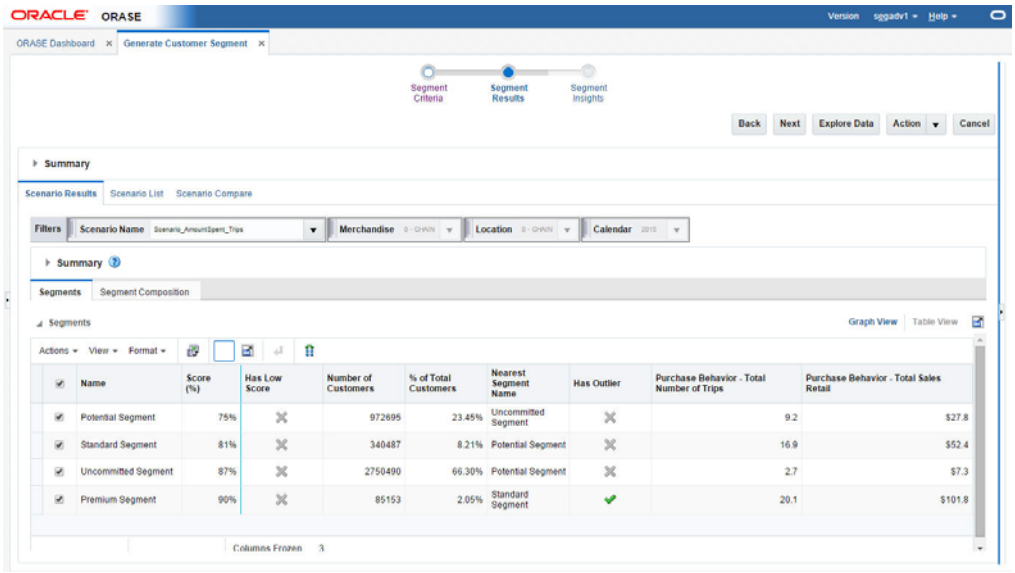
Field	Description
Name	The name you provided for the segment in the Segment Criteria stage.
Segment By	A predefined group of attributes that include Customer Demographics, RFM and Customer Behavior, and Category Purchase Driven. These criteria types are sets of attributes. For example, customer demographics are the properties of a customer. These properties can include ethnicity, income, and age.
Merchandise	The merchandise level and nodes for the segment.
Location	The location level and nodes for the segment.
Fiscal Period	The time period for the segment.
Merchandise Hierarchy Type	Provides details about which type of hierarchy the segment criteria have been created for.

Scenario Results Section

The Scenario Results section displays the following:

- The Scenario Summary, which provides key segment set attributes for the executed scenario as well as its status.
- The Scenario Results, which has two tabs: Segments and Segment Composition.

Figure 5–23 Customer Segmentation Scenario Results



The Summary section provides an overview of the characteristics of the scenario.

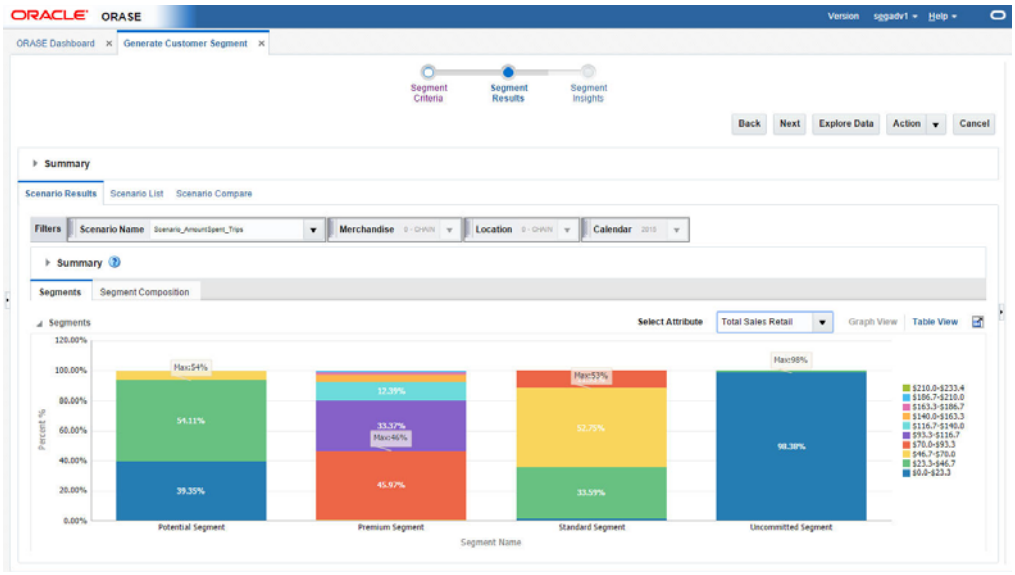
Table 5–21 Scenario Results Summary

Field	Description
Status	Ready for Preview, Ready for Approval, Approved, Rejected.
Optimal # of Segments	The optimal number of segments determined by the optimization.
Rank	The application compares executed scenarios and ranks them. A value of 1 indicates the best scenario.
Max./Min # of Segments	The number you provided for the maximum and minimum number of segments to calculated.
Largest/Smallest Segment	Provides the sizes of the largest segment and the smallest segment in order to show the range of values.
Is System Preferred	Indicates whether or not the system prefers the scenario.
Is User Preferred	Indicates whether or not the user prefers the scenario.

The Segments section provides the segment results for each individual segment in the scenario in either a Graph View or a Table View. The attributes displayed depend on the Segment by option chosen in the Segment Criteria stage.

The Graph View shows the percentage for each attribute in the segment.

Figure 5–24 Segments Graph View



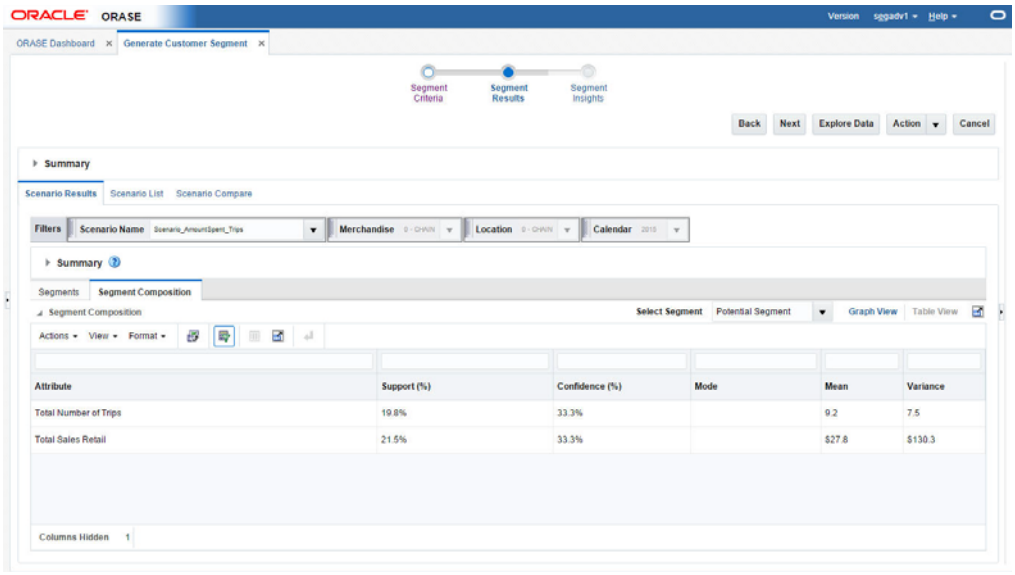
The Table View provides details that can help you analyze the segment.

Table 5–22 Scenario Results - Segments: Table View

Field	Description
Name	The name you assigned to the segment.
# of Customers	The number of customers in the segment.
% of Total Customers	The percentage of the total customers that the number of customers represents.
Nearest Segment Name	The name of the segment that is most similar to this segment.
Score %	This value is calculated at the level of customer and then averaged to the segment. The probability, expressed as a percent, of a customer being present in this segment rather than any of the other segments. See Segment Scores.
Has Outlier	Indicates a segment with the number of customers below a threshold. For example, the number of customers are below certain percentage of the number of customers in a segment. See Outlier Indicator.

The Segment Composition sub-tab breaks down the segment into its component parts and shows the percentages for each attribute.

Figure 5–25 Segment Composition



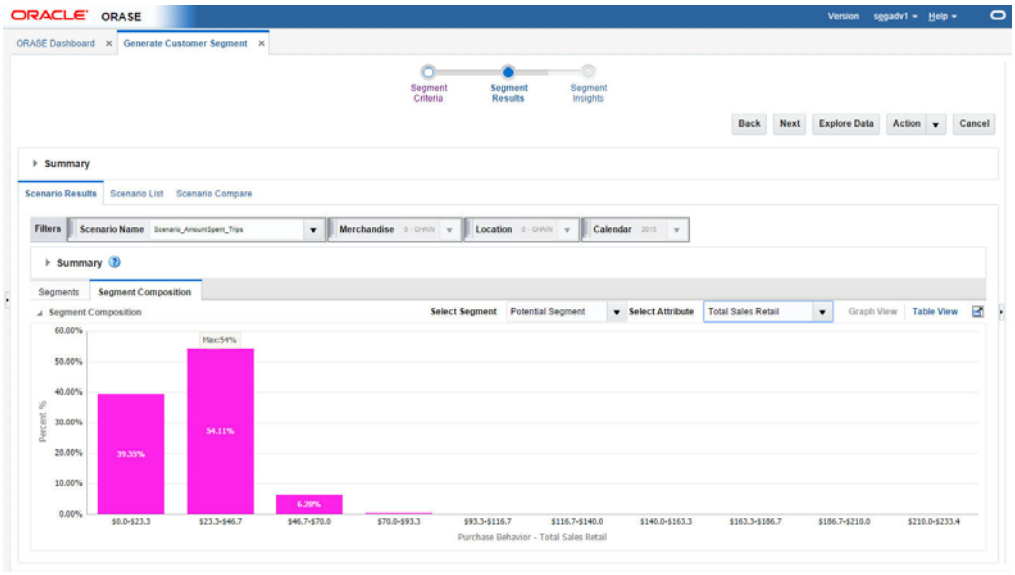
The Table View shows attributes, their confidence, and summary statistics within each segment.

Table 5–23 Segments Composition Table View

Field	Description
Attribute	The name of the attribute that was used in generating the segment.
Category	The category name, which is only available for Category Purchase Driven segment by.
Support	The percentage of the customers that are concentrated within the segment according to the attribute.
Confidence	Measures the goodness of match between the distribution of the data points for the attribute, indicating the power in recognizing a new pattern in data. An attribute with a higher confidence number plays a more important role in that segment.
Mean	The average value of the numeric attribute within the selected segment. If no value is available, this is empty.
Mode	The most common value of the discrete attributes within the selected segment. If no value is available, this is empty.
Variance	This explains the dispersion from the average for the numeric attributes within the selected segment. If no value is available, this is empty.

The Graph View shows the segment composition by attribute.

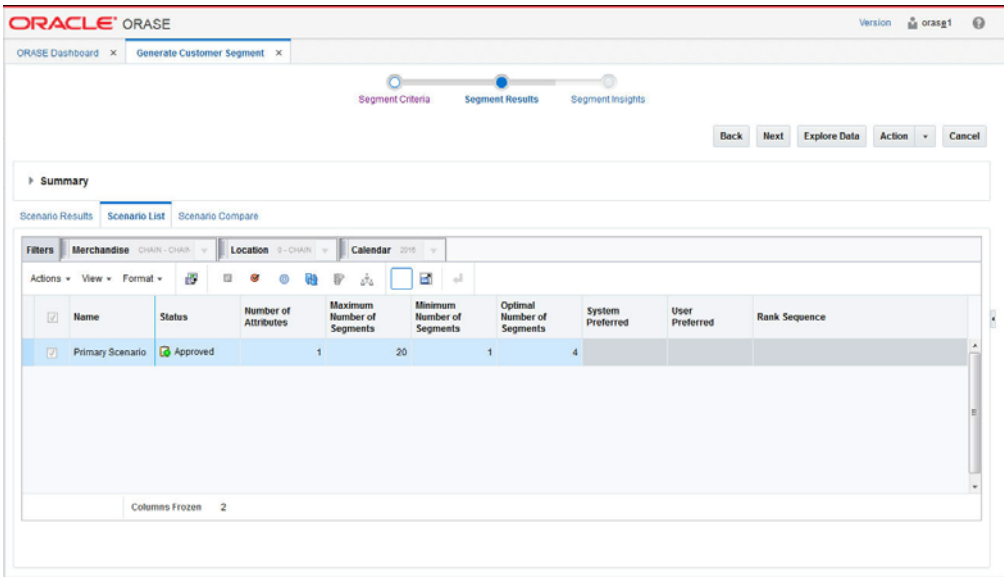
Figure 5–26 Segment Composition Graph View



Scenario List

The Scenario List section contains one table with details about each segment. Here, you can evaluate KPIs and approve or reject a segment.

Figure 5–27 Segment Results Scenario List



After a segment is approved, it is available for other applications.

Table 5–24 Scenario List

Field	Description
Name	Name assigned to the scenario.
Status	Ready for Preview, Ready for Approval, Approved, Rejected.
# of Attributes	The total number of attributes used in the segmentation process.

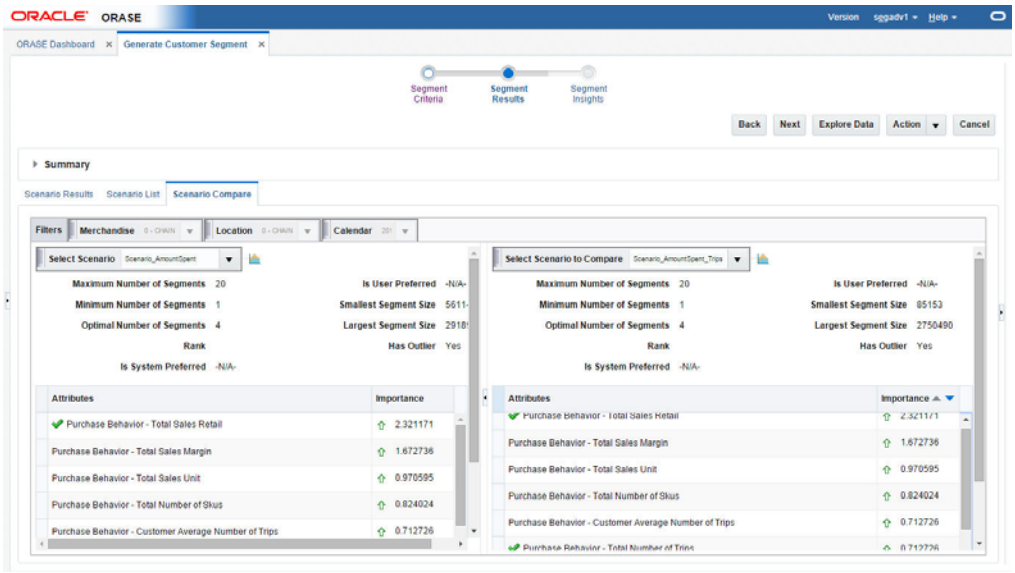
Table 5–24 (Cont.) Scenario List

Field	Description
Max. # of Segments	The value used for the maximum in the scenario execution, if this option used.
Min. # of Segments	The value used for the minimum in the scenario execution, if this option used.
Optimal # of Segments	The value used for the optimal number of segments in the scenario execution, if this option used.
System Preferred	Indicates whether the scenario is the one the application prefers.
User Preferred	Indicates whether the scenario is one the user prefers.
Rank Sequence	Indicates the ranking the scenario is given by the application.

Scenario Compare

The Scenario Compare section shows two clusters of your choosing side by side so that you can compare the segment results of each, using the same characteristics used in Scenario Results and Scenario List.

Figure 5–28 Segment Results Scenario Compare



The information displayed includes:

Table 5–25 Scenario Compare

Field Name	Description
Max. # of Segments	The value used for the maximum in the scenario execution, if this option used.
Min. # of Segments	The value used for the minimum in the scenario execution, if this option used.
Optimal # of Segments	The value used for the optimal number of segments in the scenario execution, if this option used.
Rank	The value for the rank.
Is System Preferred	Indicates whether the scenario is the one the application prefers.

Table 5–25 (Cont.) Scenario Compare

Field Name	Description
Is User Preferred	Indicates whether the scenario is the one the user prefers.
Smallest Segment Size	The size of the smallest segment.
Largest Segment Size	The size of the largest segment.
Has Outlier	Indicates a segment with the number of customers below a threshold. For example, the number of customers are below a certain percentage of the number of customers in a segment.
Attributes	A list of relevant attributes.

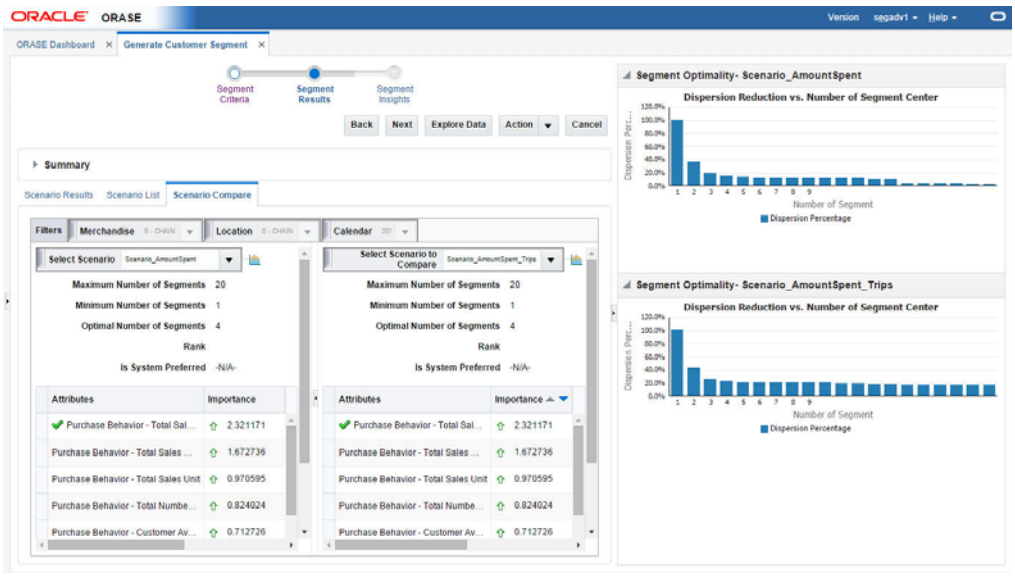
Scenario System Recommendations

The application provides the following recommendations at the scenario (segment set), segment, and customer levels.

Scenario Optimality

This graph indicates how the system identifies the best number of segments for a given data set. It starts with a small number of segment centers and searches for the number beyond which there is minimal dispersion. At this point, increasing the number of segment centers any more only reduces dispersion by a small amount, and the marginal improvement is small.

Figure 5–29 Scenario Optimality



Scenario Rank

You can see the ranking of all scenarios in the Segment Results step. The scenario with the highest rank is designated as System Preferred. The ranking is based on the following:

- How many similar customers are contained in the segment.
- How well separated the segments are from each other.

The rank for a scenario is calculated by taking the average of the probability of all the data points in that scenario. A higher probability value for a customer is an indication of the goodness of fit of the segment.

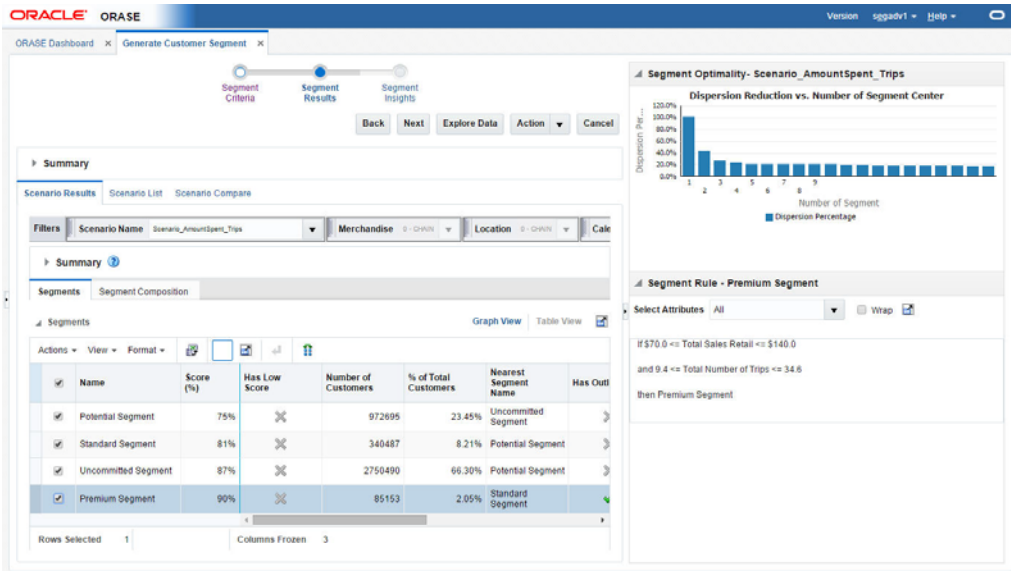
Outlier Indicator

This provides an outlier indicator in the segment list if the segment has an outlier customer. Two outlier rules are supported. The distance from the centroid indicates that if a customer is beyond a certain limit for the configuration threshold from the centroid, then the segment to which the customer belongs is marked as an outlier. When the size of a segment is compared with the total customers, if the number of customers in the segment is below a certain configured percentage of the total customers, then the segment is marked as an outlier. The distance from the centroid is the outlier rule that is configured by default.

Segment Rules

Segment rules describe the data in each segment where most of the data is concentrated. A rule is a conditional statement that captures the logic of how the system generates the hierarchy of segments and what conditions it uses for further splitting a segment. Each rule describes the conditions for a customer to be assigned with some probability to a segment. Figure 5–30 provides an example of rules.

Figure 5–30 Segment Rules



Segment Scores

The application provides scores for segments, based on the threshold score. The scoring measure the goodness of a segment. A score for a particular segment is calculated using the average of the probability scores of the customers in a segment. The score is based on the assumption that each customer has an equal chance of being a member of the segment. A high score indicates that the customer is close to the centroid. A low score indicates that the customer is an outlier. The low score indicator is turned on and off.

New Customers

Customer Segmentation supports periodic processing rules in order to allocate customers who are new. If a customer over time has built enough sales history, the system begins to consider customers and allocates them to segments based on the probability scores.

Generate Store Profile

This allows the user to calculate the sales share of customer segments for each store. These store profiles can be generated by the user from the user interface for the approved customer segments. The store profiles can then be consumed by RI or a subscribing application to be used in generating business reports. The profiles are also consumed in Advanced Store Clustering and used to generate customer-centric store clusters.

Segment Insights Stage

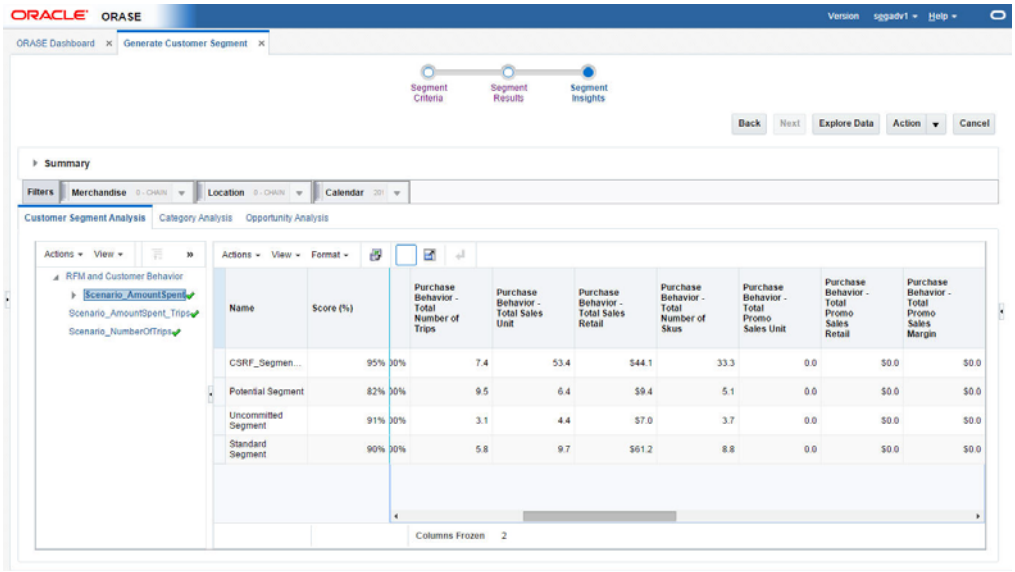
Use the Insights stage to analyze a scenario, its segments, and its hierarchy, based on performance and attribute contributions, prior to the approval of the scenario. This stage includes the following tabs:

- Customer Segment Analysis
- Category Analysis
- Opportunity Analysis

Customer Segment Analysis

The Customer Segment Analysis tab lists all the scenarios, the customer segments for each scenario, and the related KPIs. Aggregates (mean and mode) of the attributes are displayed for each segment. The attributes by default are included in the grids customer purchase behavior and customer demographics.

Figure 5–31 Customer Segment Analysis



This tab includes the following tasks:

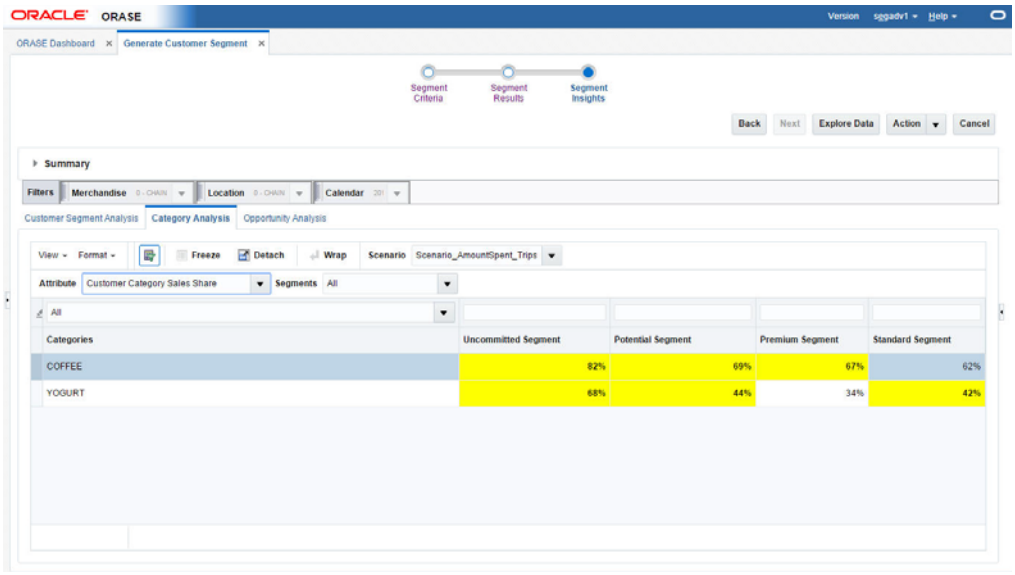
- Evaluate the KPIs for a scenario, it this has not been completed earlier

- Approve a scenario
- Rename segments within a scenario
- Rank scenarios, if not completed earlier
- Flag a segment scenario as system preferred
- Generate store profile for an approved scenario

Category Analysis

The Category Analysis tab lists the top categories vs. the customer segment matrix. You can analyze the segment performance in each category using different sets of attributes such as sales share or number of trips.

Figure 5–32 Category Analysis



The system highlights the top three performing segments for each category, and you can compare category performance across segments. You can filter categories or select a set of segments while comparing the results.

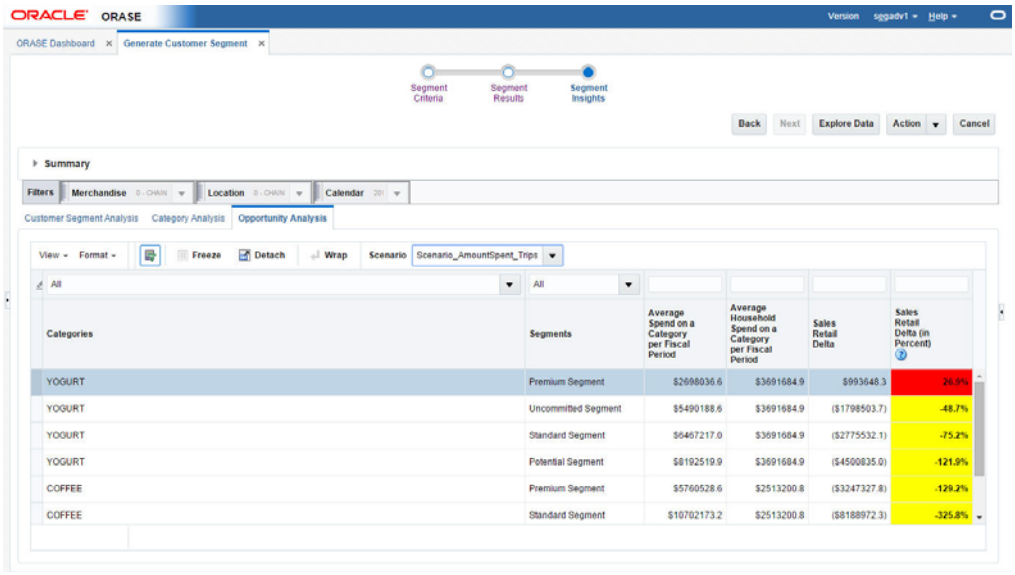
Table 5–26 Axis

Field	Description
Columns	The segments for the selected scenario.
Rows	The top identified categories in the preprocessing.
Cell	The KPI values based on the attribute selected.

Opportunity Analysis Tab

The Opportunity Analysis tab provides a comparison of the retailer to market data analysis for the category and provides an analysis about where there is an opportunity for the retailer to grow for each segment.

Figure 5–33 Opportunity Analysis



You can compare where the retailer is performing better by comparing the average sales retail spend per <configurable time period>.

Table 5–27 Opportunity Analysis

Field Name	Description
Segment	The segments for the selected scenario.
Category	The top identified categories in the preprocessing.
Average spend on a category per <configurable time period>	The average amount spent for a segment.
Average household spend on a category per <configurable time period>	The average household amount spent for a segment.
Sales retail delta	The difference in sales retail between retailer and market data.
Sales retail delta as a percent	The percentage difference in sales retail between retailer and market data.

The color indicators listed in the following table help you see if the difference in sales retail is above or below a certain threshold or within an acceptable range. This grid is sorted with the sales retail difference as negative or highlighted as red, so that you can review the values below threshold first.

Color	Value
Red	X percent below
Green	Y percent above
Yellow	In between X below and Y above

This grid provides a filter for customer segment and category, so that you can review the category for each segment or vice versa.

Summary

This lists the criteria you initially selected to define the segment.

Figure 5–34 Summary



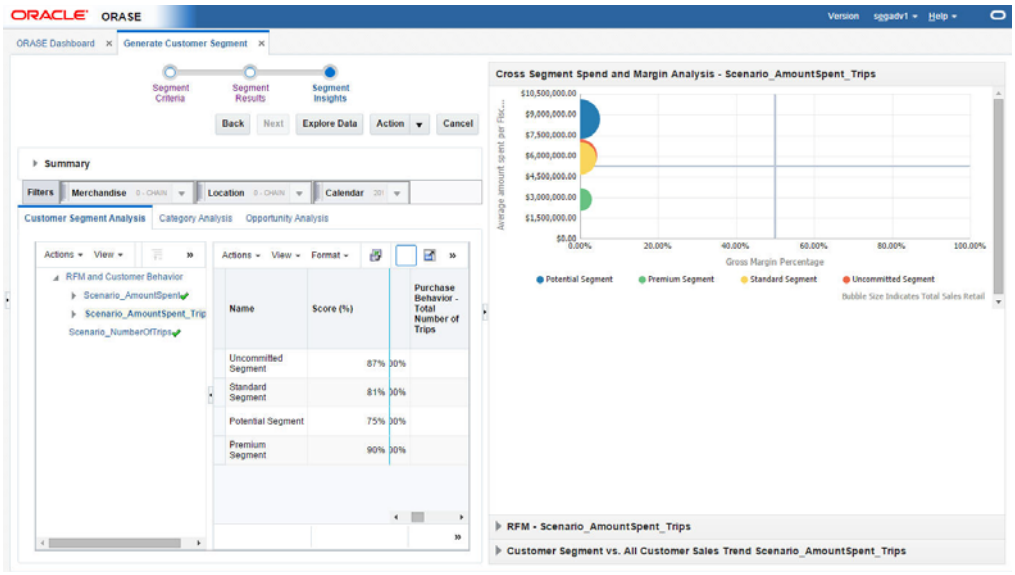
Table 5–28 Summary

Field Name	Description
Name	The name you provided for the segment in the Segment Criteria stage.
Segment by	A predefined group of attributes that include Customer Demographics, RFM and Customer Behavior, Category Purchase Driven. These criteria types are sets of attributes. For example, customer demographics are the properties of a customer. These properties can include ethnicity, income, and age.
Merchandise	The merchandise level and nodes for the segment.
Location	The location level and nodes for the segment.
Fiscal period	The time period for the segment.
Merchandise hierarchy type	This provides details about which type of hierarchy the segment criteria have been created for.

Cross Segment Spend and Margin Analysis BI

This BI compares how segments vary by Gross Margin% (*x* axis), Average Amount Spent <per configuration period> (*y* axis), and Total Sales Retail (*z* axis) for the customer segment. This can help you understand which segments are underperforming and have the potential for growth. The segments in the middle of the chart are the segments that have growth potential, as compared to the premium segments to the top right.

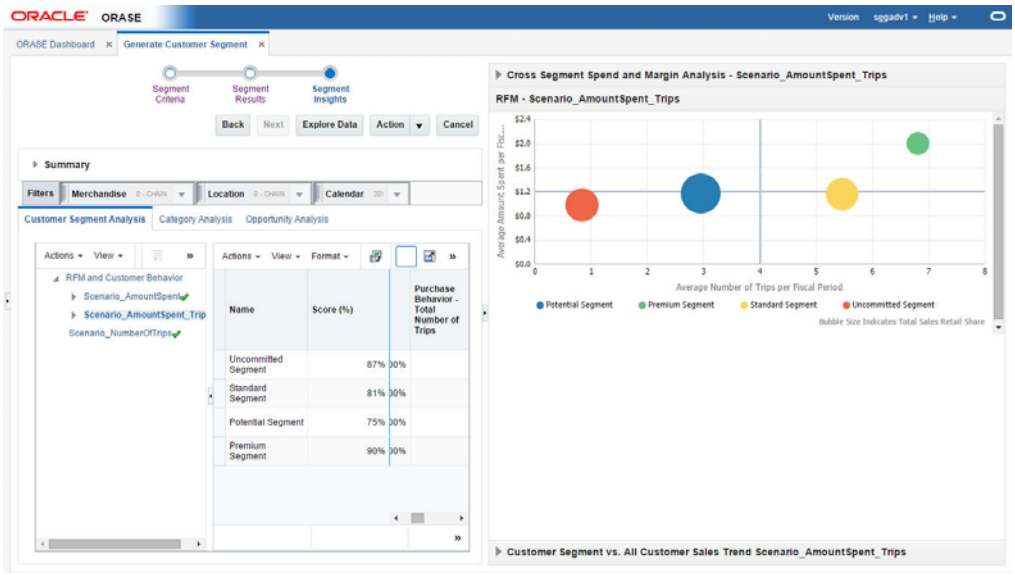
Figure 5–35 Cross Segment Spend Analysis



RFM Analysis BI

This BI provides analysis of segments based on frequency and amount spent. The more frequent customers tend to have a significantly better response and average spend than those who have shopped less recently or less frequently. This indicates customers that are high spender and frequent shoppers are premium segments, vs. low spender and infrequent shopper are uncommitted customers. Figure 5–36 shows how the premium customer segment is in the top right corner, the potential customer segments are close to the center, and the uncommitted shoppers are in the bottom left quadrant. This helps in planning how potential and uncommitted customers can be attracted by certain offers and how existing committed customers can be satisfied.

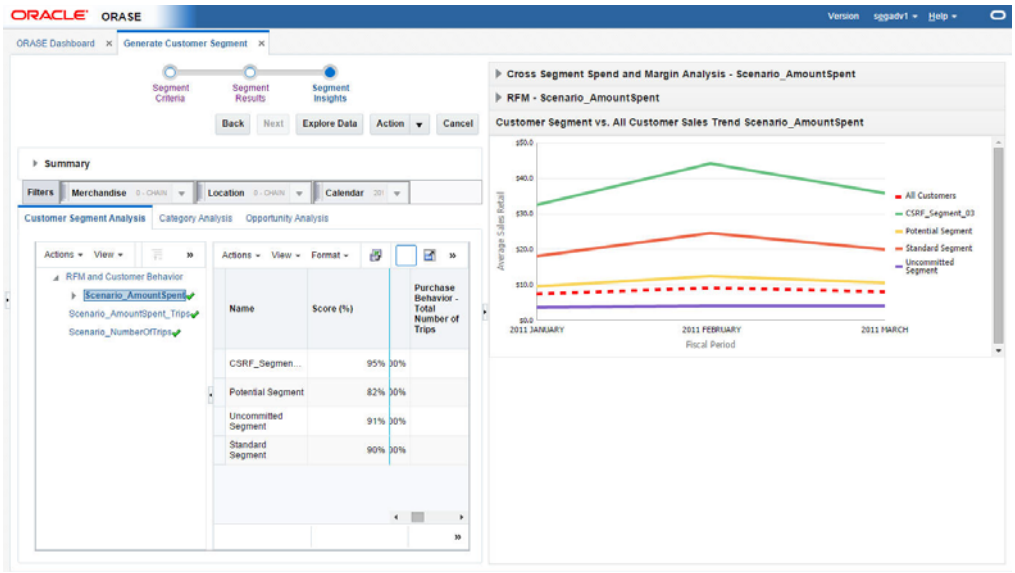
Figure 5–36 RFM Analysis



Customer Segment vs. All Customer Sales Trends BI

This BI shows a customer segment's sales trends by comparing Average Sales Retail per Segment vs. Average Sales Retail for all customers for the selected criteria time period. You can then see if sales retail has decreased for a specific customer segment and for which customer segment it has remained consistent for the defined quarters. This can help you determine if certain customers are migrating from one customer segment to another over the time period. Figure 5–37 shows how, for each segment, sales increased in 2011 February and the average sales of entire customer base, compared to the segments created.

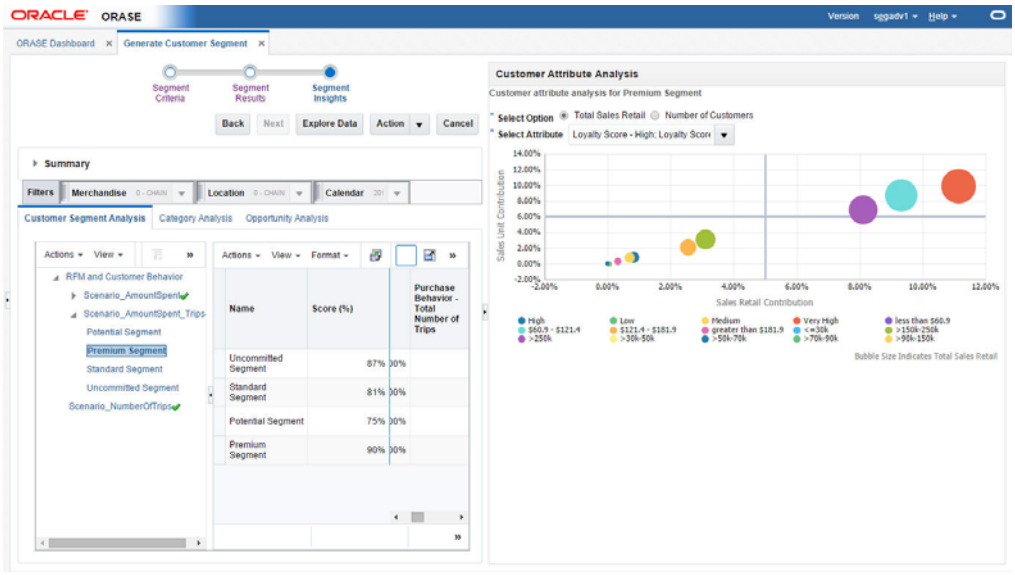
Figure 5–37 Customer Segment vs. All Customer Sales Trends



Customer Attribute Analysis BI

Customer attributes are analyzed and compared to identify the most prominent selling attributes within a segment. This indicates how the customer attributes are correlated and what each attribute in the specified segment contributes to sales. You can make inferences about which attributes in the segment contribute to significant sales and the potential attributes that should be considered for assortment planning in order to improve sales even more. This chart also provides you with the option to see the number of customers in each segment instead of total sales retail, thus providing insights about how many customers with certain attributes and attribute values are present in certain segments. Figure 5–38 shows how customers with high income and high loyalty score contribute to most of the sales in the segment.

Figure 5–38 Customer Attribute Analysis



The following properties are displayed by the graph. Note that the customer attributes are only displayed when they are configured as part of the Cluster by process.

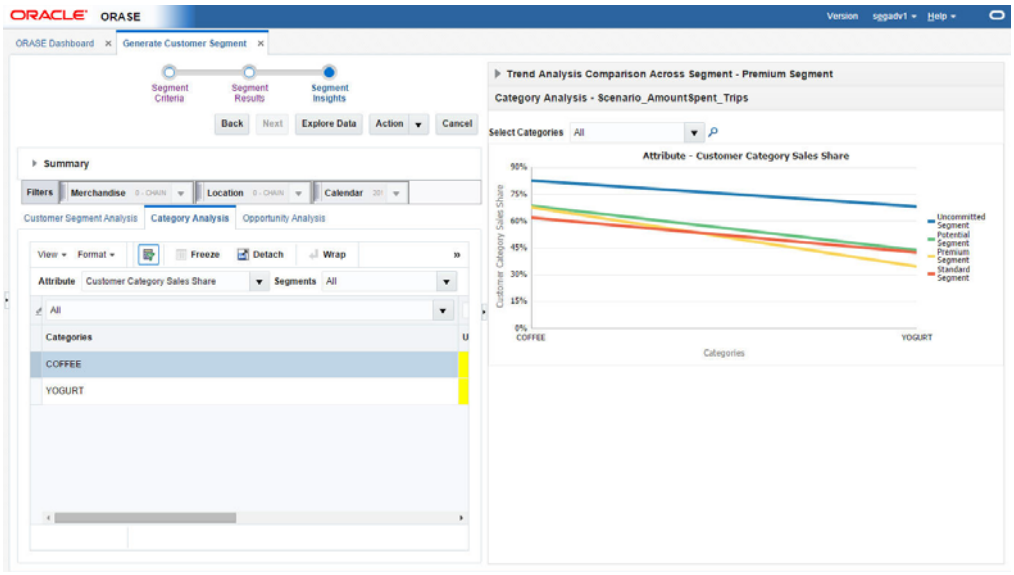
Table 5–29 Properties

Axis	Description
x-axis	Sales Retail $ contribution, calculated using the sales revenue share of each attribute in the segment with respect to the total segment sales revenue.
y-axis	Sales Unit $ contribution, calculated using the sales unit share of each attribute in the segment with respect to the total segment sales units
z-axis	Total Sales Retail $ or number of customers of each attribute in the selected segment, indicating, via the bubble, the magnitude of the sales contribution.

Category Analysis BI

You can perform category/segment analysis by comparing categories using a selected attribute (such as category sales share, category promotion sales share, average amount spent per month, or number of trips). This can indicate the segments and categories that customers are shopping frequently.

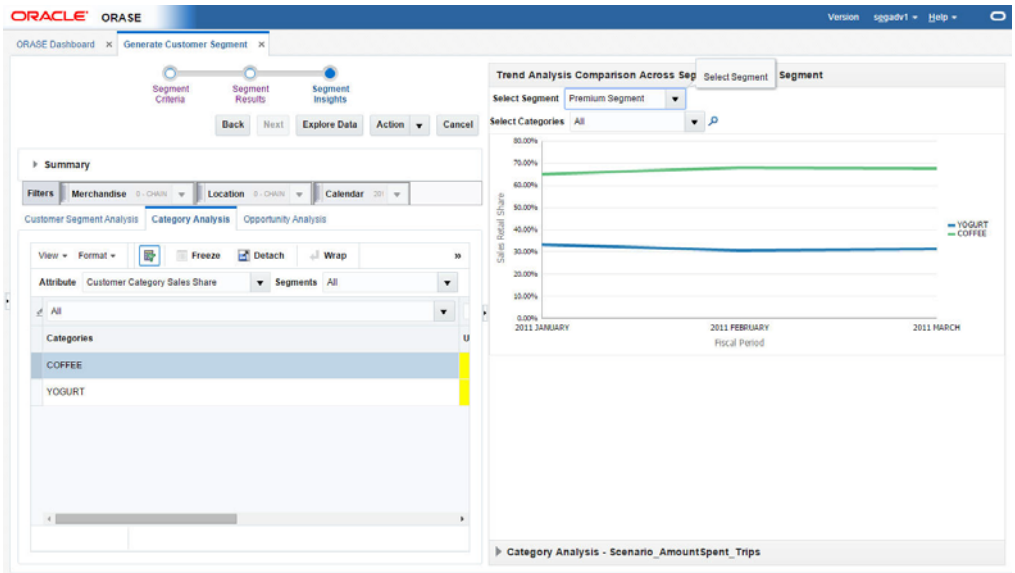
Figure 5–39 Category Analysis



Trend Category Analysis BI

This BI compares sales for selected categories in a segment. You can compare sales by <Configurable Period (Period or Quarter)> to view trends for each category. You can limit the number of categories by selecting a subset of the categories in the drop-down list. You can use this information to plan certain categories for certain seasons by comparing sales with previous periods or by modifying a category plan that is underperforming for a segment in a certain season. The time period displayed is driven by the selection you make in the Segment Criteria. If you select a source time period to include two years worth of data then you will also see the data for the previous year.

Figure 5–40 Trend Category Analysis



6

Assortment and Space Optimization

Oracle Retail Assortment and Space Optimization Cloud Service is used to determine the optimal selection and arrangement of products within stores by optimizing the product assortment and product placement on a virtual planogram. It uses information about available space, product dimensions, expected demand, replenishment schedules, merchandising rules, fixture space, and category goals in order to create a virtual planogram that optimizes total performance.

The products and product assortments that are selected for the virtual planograms are the ones that ASO Cloud Service recommends for the finalized assortment. The recommended product level data is used inside CMPO, while the virtual planograms created in ASO Cloud Service are also then available to be used in space planning.

It supports the following fixture types:

- Shelves
- Pegboards
- Freezer chests
- Shelf/pegboard combinations

Optimizations are run at either the Store level or at the Space Cluster level. Space Clusters are ad hoc groups of stores used for optimization. Each space cluster includes stores that are in the same assortment, have the same product list, and have the same current planogram length (or same POG length, height, and depth, if selected). Space clusters are typically a level between assortment cluster and store. ASO Cloud Service creates these groups by splitting stores from an assortment cluster into smaller groups of stores that have the same product list and the same current planogram length.

At a high level, ASO Cloud Service starts with an assortment that is ready for optimization. The assortment is mapped to one or more planograms, and one or more optimization runs occur. Approved runs are then available for assortment analysis and can then be finalized and exported.

Overview

This section provides information to help you understand the functionality of ASO.

Administering ASO

For information about the administration of ASO Cloud Service, see *Oracle Retail Science Cloud Services Administration Guide*.

Users and Roles

The following four user login roles are available:

- Micro Space Optimization Analyst – main business user responsible for day-to-day micro-space optimization activities
- Category Manager – product-assortment-centric user who is interested in viewing ASO results and in the translation of data between CMPO, Retail Analytics, and ASO
- Administrator – responsible for general system setup and configuration tasks related to the business
- Analytical Super User – responsible for analytical configuration, testing, and model diagnosis

Table 6–1 User Roles and Permissions

Privilege	Data Access	Micro ASO Analyst	Category Manager	Administrator	Analytical Super User
Create new ASO run	All runs	X			X
Modify existing ASO run	Runs they created	X			X
Modify the name and description of an existing ASO run.	Runs they created	X			X
View existing ASO run	All runs	X	X	X	X
Submit or re-submit ASO run	Runs they created	X			X
	Runs they created with technical failures	X			X
Approve ASO runs	Runs they created	X			
Delete saved ASO runs	Runs they created	X			X
	All runs			X	
View list of saved ASO runs	All runs	X	X	X	X
View mapping of products to POG sets	All assortments	X	X	X	X
Modify mapping of products to POG sets	All assortments		X		X
Toggle assortment status between Ready for ASO and POG Mapping Needs Review	All assortments		X		X
View rules for mapping products in POG sets	All assortments	X	X	X	X

Table 6–1 (Cont.) User Roles and Permissions

Privilege	Data Access	Micro ASO Analyst	Category Manager	Administrator	Analytical Super User
Modify rules for mapping products to POG sets	All assortments		X		X
Analyze space-optimized assortments	All assortments	X	X	X	X
Finalize and export space-optimized assortments	All assortments		X		
Roll back finalized and exported space-optimized assortments	All assortments			X	
View list of space-optimized assortments	All assortments	X	X	X	X
Delete unprocessed assortment	All assortments			X	

Table 6–2 User Constraints

Role		Privilege		Object	Status
Users	Cannot	Modify existing ASO run	When	Run is	Running, approved, or finalized
Users	Cannot	Submit or resubmit ASO runs	When	Run is	Running, approved, or finalized
Users	Cannot	Approve ASO runs	Unless	Run is	Complete
Users	Cannot	Delete saved ASO runs	When	Run is	Running, approved, or finalized
Administrator	Cannot	Delete saved ASO runs	When	Run is	Finalized
Users	Cannot	Modify mapping of products in POG sets	Unless	Assortment is	With status of POG Mapping Needs Review
Users	Cannot	Toggle assortment status between Ready for ASO and POG Mapping Needs Review	When	Assortment is	With status of POG Mapping Needs Review and also has mappings with errors (both conditions must be true)
Users	Cannot	Modify rules for mapping products to POG sets	When	Assortment is	Finalized or exported

Table 6–3 Privilege Descriptions

Privilege	Description
Create new ASO run	Set up and execute optimizations. All steps on the main train can be carried out, including creating a new run via the copy run or the save as functionality.
Modify existing ASO run	Open an existing run and modify settings. The user can visit and change settings on any step in the main train. This includes performing VPOG overrides for the optimal VPOGs that are generated by the run.
Modify the name and description of an existing ASO run	Open an existing run and modify the name and description of that run.
View existing ASO run	Open an existing run and visit any step in the main train. The user can view data and search/aggregate/filter but cannot change anything. This includes viewing the optimal VPOGs that are generated by the run and comparing them to historical VPOGs.
Submit or re-submit ASO run	Submit a new or modified run or restart a failed run. This includes validating a run.
Approve ASO runs	Approve run results.
Delete saved runs	Delete run setup information and results from the database.
View list of saved runs	View a list of summary details about each saved run.
View mapping of products to POG sets	View the mapping of products to POG sets by assortment.
Modify mapping of products to POG sets	Modify the mapping and related data of products to POG sets by assortment. This includes altering the mappings by using the re-map functionality, which refreshes an assortment's mappings based on the pre-configured rules.
Toggle assortment status between Ready for ASO and POG mapping needs review	Toggle an assortment's status between Ready for ASO and POG mapping needs review.
View rules for mapping products to POG sets	View the rules that determine how products are mapped to POG sets.
Modify rules for mapping products to POG sets	Edit the rules that determine how products are mapped to POG sets.
Analyze space optimized assortments	Aggregate approved runs to the assortment level prior to approval for export to CMPO.
Finalize and export space optimized assortments	Finalize assortments and export the aggregated assortment data to export to CMPO.
Rollback finalized and exported space optimized assortments	Roll back a finalized assortment or an exported assortment.
View list of space optimized assortments	View a list of summary details about each of the space-optimized assortments.
Delete unprocessed assortment	Delete assortments that have not moved into ASO processing. Only assortments that have the following statuses can be deleted: Received, Data Errors, Ready for Mapping, Discarded, and POG Mapping Needs Review.

Overview of the User Interface

The ASO Cloud Service user interface, at a high level, consists of

- Task pane. Located on the left side of the application. Use this to navigate between the four main tabs.

- The four tabs that you use to do the application work: Optimization List, Assortment List, Execute Space Optimization, and Assortment Mapping.
- Contextual information. Located on the right side of the application. Use the various tabs to set up runs and to view graphical and tabular data that can help you as you interpret the optimization results. In addition, the tools can be used during the run setup to provide additional information about data involved in the run setup.

Icons

The following icons are used in user interface. Certain icons have slightly different definitions, depending on the context.

Table 6–4 Icons

Icon	Icon Description
	Add. Add row.
	Compare to history.
	Copy a run. Add a planogram. On the Assortment Mappings screen, it is used to create a like planogram.
	Create. Create a run.
	Detach.
	Delete. Delete a run. Delete planogram.
	Edit. Change status.
	Embedded help.
	Export to Excel. Export all.
	Freeze.
	Go to top.
	Go up.
	Move all items to other list.
	Move selected items to other list.
	Progress indicator. The indicator changes mode when the application is processing data.
	Query by example.
	Re-map.

Table 6–4 (Cont.) Icons

Icon	Icon Description
	Refresh.
	Remove all items from list.
	Remove selected items from list.
	Required.
	Select date.
	Show as top.
	View configuration. View results.
	Wrap.

Buttons

Buttons are used to perform certain actions and for navigation.

Table 6–5 Buttons

Button	Button Description
Action	Provides access to the following actions: Submit, Re-Submit, Validate, Save As, and Approve.
Add All Stores	Within the Add Stores dialog box, used to add all stores instead of a sub-group of stores.
Apply	In Optimization Objectives, used with the Edit Objectives drop-down list.
Back	Used to return to the previous train stop.
Cancel	Used to close a dialog box without making a selection.
Choose Stores	In Optimization Setup, when creating a new run, provides access to a dialog box for adding stores.
Clear Selection	In Assortment Analysis, clears the Selected Runs list.
Details	Radio button that toggles the display between a detailed view and a summary view.
Edit Constraints	In Optimization Objectives, provides access to the Edit Constraints dialog box, which you use to edit settings.
Export	In Assortment Analysis, makes the selected finalized run available for use by other applications.
Finalize	In Assortment Analysis, finalizes the selected run.
Next	Moves to the next train stop.
No Template	Clears the template section.
Save	Saves the existing settings.
Summary	Radio button that toggles the display between a detailed view and a summary view.

Table 6–5 (Cont.) Buttons

Button	Button Description
Table View	Radio button that toggles the display between a table view of the data and a tree view of the data.
Tree View	Radio button that toggles the display between a table view of the data and a tree view of the data.
View Results	In Assortment Analysis, displays results for runs in the Selected Runs list.

Pull-Down Menus

The user interface has three pull-down menus that provide access to a variety of functionality.

The **Format** pull-down menu and the **View** pull-down menu can be used to adjust how the display is organized. For example, you can resize the columns or detach a table from the interface.

The **Actions** pull-down menu provide functionality that you use to manipulate the application data. Some of the following functionality is also provided by the icons described in Table 6–4.

Table 6–6 Actions Pull-Down Menu Functionality

Action Name	Description
Create Run	In Optimization List, takes user to Optimization Setup train stop.
Copy Run	In Optimization List, used to create a duplicate of an existing run.
Delete Run	In Optimization List, used to delete a run.
Export to Excel	Used to export to Excel.
Rollback	In Assortment List, used to revert the status of a run.
Re-Map	In Assortment Mapping, re-maps the planogram to product. This action must be executed after adding POGs, mapping data or any other information that has been delivered in order to correct exceptions or errors found by the assortment to POG mapping process.
Add Planogram	In Assortment Mapping, provides access to the Select Planogram Nodes dialog box.
Delete Planogram	In Assortment Mapping, used to delete the selected planogram.
Add Row	Add a row to the display.
Mark for Optimization	In Fixture and Product Data, used to select a product for optimization (indicated with a check mark).
View Configuration	In Fixture and Product Data, displays an empty planogram for the selected product.
Compare to History	In Fixture and Product Data, and in Results and Analysis, displays historical planogram that is populated.
Add	In Objectives and Constraints - Product Groups, used to add data.
Edit	In Objectives and Constraints - Product Groups, used to edit data.

Table 6–6 (Cont.) Actions Pull-Down Menu Functionality

Action Name	Description
Delete	In Objectives and Constraints - Product Groups, used to delete data.
View Results	In Assortment Analysis, processes the selected run. Once the processing is complete, the results are displayed in the table.
Export All	In Results and Analysis, used to export all the results.
Create	In Fixture and Product Data, used to create.
Change Status	In Assortment Mapping, used to update the status.
Save As	Makes a copy of the run that is currently open.

Histograms

The stages have associated histograms available on the right-hand side of the display that can help you analyze the data presented in that stage. You can adjust the way the histogram presents the data by changing the number of bins that are used in the display. Once you change the number of bins, click the **Refresh** button to update the display.

Charts

Certain stages have associated charts available on the right-hand side of the display that list data in a tabular format.

Process Train

The process train displays the stages of Execute Space Optimization. The current stage is highlighted. You can also use the **Back** button and the **Next** button to move through the train.

Embedded Help

Embedded help, which you access by clicking the **Question Mark** icon, provides additional information about the details required by certain fields.

Process Indicator

At the top of the user interface, in the right-hand corner, is a process indicator that you can use to monitor the status of a user action such as clicking **Next** to go to the next stage.

Search

In certain cases, you can customize your search, using advanced search capabilities to specify the search criteria.

Workflow

ASO Cloud Service is used to optimize a category manager's assortment plan by creating optimal planograms. The category manager creates a preliminary assortment and wants to determine how well that list of products can fit in stores, given the available space, product sizes, and merchandising goals, constraints, and rules. The application creates virtual planograms that organize products onto fixtures in a way that best achieves the optimization objectives. Once the manager finalizes the

optimization results, the product level data can be exported for use in planning applications.

Optimization

An optimization can be carried out at one of two location levels: Store or Space Cluster.

In a Store level optimization:

- each store is optimized separately, based on the store's individual data
- one planogram is produced for each store
- store-specific results are produced

In a Space Cluster level optimization:

- every space cluster is optimized separately, based on aggregate store data
- one planogram is produced for each space cluster
- space-cluster-specific results are produced

Space Cluster

A space cluster is an ad hoc group of stores used for optimization. Each space cluster includes stores that are in the same assortment, have the same product list, and have the same current planogram length. Space clusters are typically a level between assortment cluster and store. A fixture configuration can optionally be added to the clustering rules. ASO Cloud Service creates these groups by splitting stores from an assortment cluster into smaller groups of stores that have the same product list, the same current planogram length, and same fixture configuration if activated.

Run Templates

You can pick an existing run that was created for the same POG Set as a template. The system uses the settings for that run and applies them to a new run (which has a new group of products, locations, and data). This allows you to leverage the initial settings from the template run and use them across time, locations, planogram sets, and users.

When you select an existing run as a template, this functionality allows the visual guidelines to be copied over to the new run with a different assortment. If the assortment has not changed, you do not need to copy a run from template; just use Save As to provide exact copy. The idea behind this functionality is that the destination run has a different assortment, hopefully not completely different, compared to the source run, so visual guidelines that have been created earlier can be re-used when the assortments overlap. This copy functionality copies over the best possible match from the template to the destination run, based on the quality of matching on different criteria. Destination and source space clusters are scored by taking a weighted average of the following criteria (in the same order): Hard match on Fixture Type (for example, Pegboard Visual Guidelines are not copied over to Shelf), Primary Attributes, Secondary Attributes, POG dimensions, and Store Count. For the destination space cluster, the source space cluster with the highest score is selected. If you do not like the automatic match provided, you can edit the visual guidelines for the destination space cluster.

Validation

Within each stage, you can perform a validation. This provides you with feedback concerning the settings you have configured. Validation can identify conflicts between

settings. You see details in the validation report that can help you to resolve any issues. Note that the validation is the same, regardless of which stage you invoke it from.

The conflicts identified by validation include warnings and errors. Warnings occur when constraints can impact or limit the optimization process. Errors prevent the optimization process from producing a solution.

Planogram Mappings

ASO Cloud Service supports various planogram mappings, including the simple case in which a single product category is mapped to a single planogram as well as the case in which a product category's assortment is mapped to multiple planograms in many locations in a store.

Smart Start Process

If detailed shelving data is not available, a planogram cannot show the complete details about the shelving. In such cases, ASO Cloud Service creates fixture details using a combination of default values and user selection in order to assign shelves to partial shelf fixture planograms. This process occurs automatically when optimization locations are generated. These fixture configurations are available in the Fixture and Product Data stage. The user can manually override such planograms within the ASO Cloud Service user interface. See Planogram Setup for additional information.

Fixture Types

ASO Cloud Service works with three types of fixtures: shelves, pegboards, and freezer chests. Freezer chests are not the containers in the frozen food section; rather, they are chests that shoppers reach down into in order to select a product. Planograms can support a mixture of shelves and pegboards.

Virtual Planogram

A virtual planogram (VPOG) is a planogram that is created by the application. It contains colored rectangles representing products, product groups, or images of products (when available), with supporting hover text for each product on the virtual planogram. A planogram is the layout of the product on shelves or pegboards with an underlying fixture assumption (for example, the number of shelves or the overall size).

A user can open multiple virtual planograms simultaneously in order to view images for different clusters or stores within the same run. However, only the first virtual planogram that a user opens can be edited, subject to the run status and the permission rules, until the open virtual planograms are closed. The application only supports the opening of multiple simultaneous virtual planograms if the user is using the same browser session, browser tab, and login session for the same optimization run. In other words, the user must open the multiple virtual planograms for a single run in a single run screen. All other combinations of browser sessions, browser tabs, and multiple logins to open multiple simultaneous virtual planogram windows are not supported. Although some combinations may work in practice, consistent UI behavior is not guaranteed and data integrity issues may occur.

Planogram Set

Planogram set is an ASO concept that refers to a planogram node (which is the subcategory name) and concatenates it with the planogram season, with a hyphen between.

Assortment Finalization

The Assortment Finalization process handles the assortment summarization of the results at the assortment set level as well as the transformation of placeholder product names and IDs into final products. This process is initiated by the Assortment Planning and Optimization (APO) module of CMPO. ASO loads and integrates the Assortment Finalization file with database objects. This triggers the following:

- ASO generates aggregated results for all the exported assortments that have the same assortment set ID.
- The only assortments that are aggregated are the ones that have been exported to APO.
- All other assortments that have the same assortment set ID but that have a status other than exported are disabled so they can no longer be used in optimization or analysis.

Assortment List

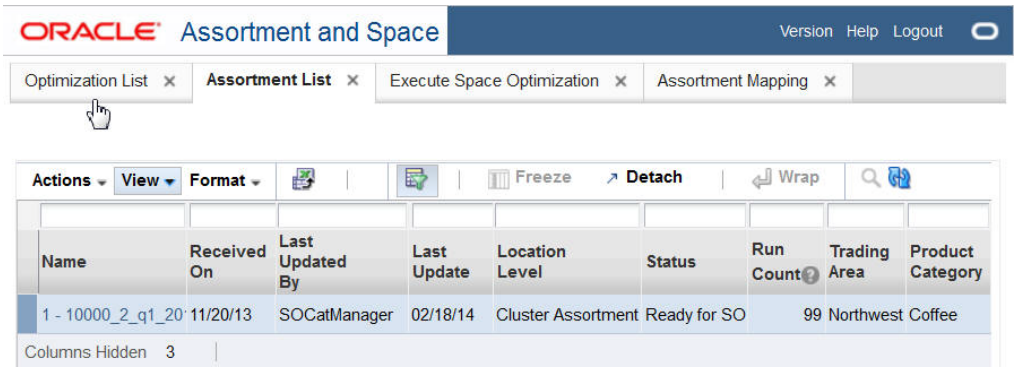
The Assortment List tab has one table that presents a list of all available assortments. You can use this information to view descriptive information and the status of each of the listed assortments. Use the **Refresh** icon to update the display.

You can also access the Assortment Analysis functionality from the Assortment List tab. You use the Assortment Analysis screen to determine how much of an assortment has or has not been optimized, to analyze the results of the optimizations, and to finalize and export your work.

Assortment List

The Assortment List is shown in Figure 6–1.

Figure 6–1 Assortment List



It contains the following columns.

Table 6–7 Assortment List

Field	Description
Name	The display name for the assortment.
Goal	The goal of the assortment.
Role	The role of the assortment.
Tactic	The tactic of the assortment.

Table 6–7 (Cont.) Assortment List

Field	Description
Received On	The date the assortment was received.
Last Updated By	The user ID of the person who most recently updated the assortment.
Last Update	The date when the assortment was most recently updated.
Location Level	The location level for the assortment (Cluster Assortment, Store Assortment, or Space Cluster). This is the level at which the assortment data is delivered. Products are delivered once for all stores within a cluster (Cluster Assortment) or once for every store (Store Assortment)
Status	The status of the assortment in terms of space optimization. Values include Received, Data Errors, Ready for Mapping. POG Mapping Needs Review, User Re-Mapping, Ready for SO, Finalized, Exported, and Assortment Set Finalized. See Table 6–8 for details about each status. Select any assortment with a status of Finalized or Ready for SO in order to display that assortment in Assortment Analysis.
Run Count	The total number of existing optimization runs (regardless of status) that have occurred, based on a planogram set that is included in the assortment.
Trading Area	The trading area for the assortment.
Product Category	The product category for the assortment.

Table 6–8 Assortment Status Descriptions

Status	Description
Received	The assortment has been received by ASO and has been loaded within ASO tables.
Data Errors	The assortment has been validated by ASO against other datasets in the application, and errors have been found. Assortments that have failed global validation checks will have a change of status to Data Errors and will not be included in any further processing until the data is fixed and a new global validation run verifies the data integrity.
Ready for Mapping	The assortment has passed global validation so its status has changed to Ready for Mapping. As a result, the assortment can now be included in the Assortment to POG mapping process.
POG Mapping Needs Review	The mapping process has been executed and some exceptions have been found that require user review. Depending on the system configuration, a review might be required, regardless of mapping process results.
User Re-Mapping	Internal state is not visible in UI. This is an internal indicator for a new mapping request from the user.
Ready for SO	The assortment has been mapped. It is exception free and available for use within the optimization process.
Finalized	An individual assortment reaches this status once the assortment's results have been reviewed and accepted by the user.
Exported	The final status of the individual assortment, indicating that the data is available for export.

Table 6–8 (Cont.) Assortment Status Descriptions

Status	Description
Assortment Set Finalized	This is a signal from APO indicating that the processing for the whole assortment has ended and that the combined results from the related assortments in SET should be created and made available for export.

About Rollbacks

The **Rollback** functionality is available from the **Action** drop-down menu.

If the current status of the assortment is Exported, the status rolls back to Finalized. After the rollback from Exported, the assortment is not available to external application (as it was when it had a status of Exported).

You cannot roll back any run. If an assortment is in Finalized status, you will not be able to create any new runs based on POG sets within that assortment so you must roll the assortment back to Ready for SO status, create more optimization runs, and finalize again if necessary.

An assortment cannot be part of an assortment set that has been finalized. In such cases, the rollback menu option is disabled.

When you roll back an assortment to Ready for SO, you cannot export it but you can finalize eligible runs that are associated with the assortment that were not finalized before. In addition, once you have rolled the assortment back to Ready for SO, you can create new runs that are based on planogram sets associated with the assortment, something you cannot do when the assortment is in status of Finalized or Exported.

Assortment Analysis

Assortment Analysis is available for an assortment if one or more of the runs associated with that assortment have been approved. Such runs have a status of Approved or Finalized. CMPO users and SO analysts have permission to conduct this analysis. The SO Super User can access this screen as read-only.

The Assortment Analysis screen can be used to:

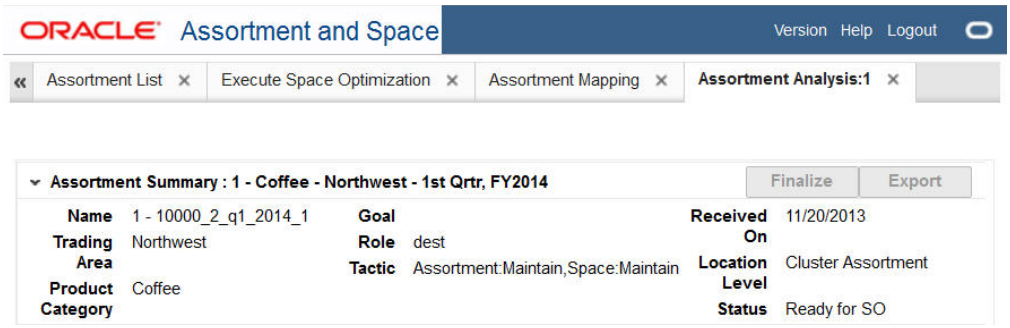
- determine how much of a specific assortment has not been optimized
- examine the aggregated results of the approved runs for a specific assortment
- finalize the optimized results
- export the results

To access Assortment Analysis, select an assortment from the Assortment List that has a status of Ready for SO or Finalized. You see the Assortment Analysis screen. This screen consists of three distinct areas.

Assortment Summary

This section of the Assortment Analysis screen, shown in Figure 6–2, displays a subset of the information contained in the Assortment List. The information describes the assortment you selected.

Figure 6–2 Assortment Summary

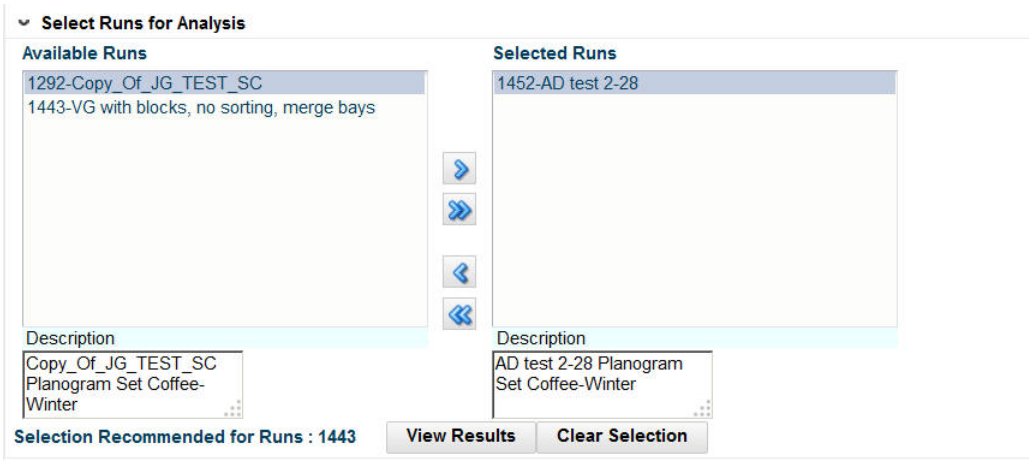


Select Runs for Analysis

This section, shown in Figure 6–3, displays the runs for the assortment you selected. The Available Runs lists all the runs associated with the assortment whose status is Approved or Finalized. Only approved or finalized runs appear here. The Selected Runs lists the runs whose results you are interested in viewing. Overlaps may exist across selected runs. In that case, results from the most recently approved or finalized run take precedence over other runs with the same sku/store combination.

The Selected Runs list is also used to control which runs become finalized. If you select more than a run that covers the same product/location, you will only see results from the run with the most recent results data.

Figure 6–3 Select Runs for Analysis



To move all runs from the Available Runs list to the Selected Runs list, click the double arrow.

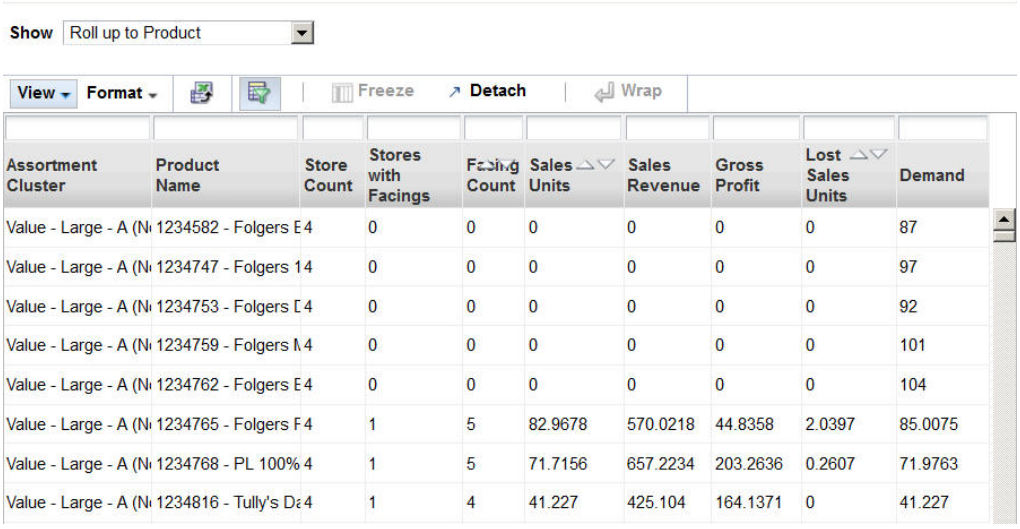
To move a select group of runs from the Available Runs list to the Selected Runs list, highlight the run or runs you want in the Available Runs list and click the single arrow.

To moves runs from the Selected Runs list to the Available Runs list, you should reverse this process.

Run Results

To display the results of a run in the Run Results table, shown in Figure 6–4, highlight the run, and click the **View Results** button. (Note that some columns are hidden so that the figure fits on the page.)

Figure 6–4 Run Results



By default, the results are displayed for all data. To display the data at different levels of aggregation, use the **Show** drop-down list. You can choose from the following options:

- All data
- Roll up to Assortment Cluster (a single line is displayed for the entire cluster)
- Roll up to Product. Product: Results are displayed at the Assortment Cluster/Product level. You see one line for each product within each cluster.
- Roll up to Planogram Set. Results are displayed at the Assortment Cluster/Product/Planogram Set level. You see one line for each planogram set that a product belongs to within an assortment cluster.

To clear the display, click the **Clear Selection** button. This clears the results from the results table and moves all of the runs in the Selected Runs list back into the Available Runs list (resets the user selection).

Results Table

The Results Table displays the following information at the aggregation level you selected. Note that each aggregation level displays only a subset of the information listed in Table 6–9.

Table 6–9 Results Table

Column Name	Description	Displayed at Aggregation Level
Assortment Cluster	The name that identifies the cluster.	All, Cluster, Product, Planogram Set
Product Code	The code assigned to the product.	All, Product, Planogram Set
Product Name	The name that identifies the product.	All, Product, Planogram Set
Planogram Set	The name that identifies the planogram set. The name is a concatenation of the subcategory name and the season name.	All, Planogram Set
Store Code	The code assigned to the store.	All

Table 6–9 (Cont.) Results Table

Column Name	Description	Displayed at Aggregation Level
Store Name	The name that identifies the store.	All
Run ID	The ID that identifies the specific run. This is useful for assortments with more than one run.	All
Store Count	Tells the user how many stores the data from that row represents.	Product, Planogram Set
Stores with Facings	The number of stores whose facing count is greater than 0 (see Store Count description).	Product, Planogram Set
Facing Count	The facing count.	All, Cluster, Product, Planogram Set
Sales Units	The number of the standard sales category for a product.	All, Cluster, Product, Planogram Set
Sales Revenue	The amount of revenue for the product.	All, Cluster, Product, Planogram Set
Gross Profit	The difference between sales revenue and cost.	All, Cluster, Product, Planogram Set
Lost Sales Units	The lost sales units.	All, Cluster, Product, Planogram Set
Demand	The demand.	All, Cluster, Product, Planogram Set

Finalizing an Assortment

Which runs are finalized is based on which runs are used to generate the results data (the set of unique Run IDs in the View Results table). In order to determine the runs used to generate the data, the most current run for each product/store is determined. For example, when you view the results for All, you may see records that have different Run IDs.

The Selection Recommended for Runs should include all runs used to generate the results. A run must be in this list and in the Selected Runs list in order to become finalized when the assortment is finalized.

Once you are satisfied with the results and finalize an assortment, it is available for export to CMPO and Space Planning.

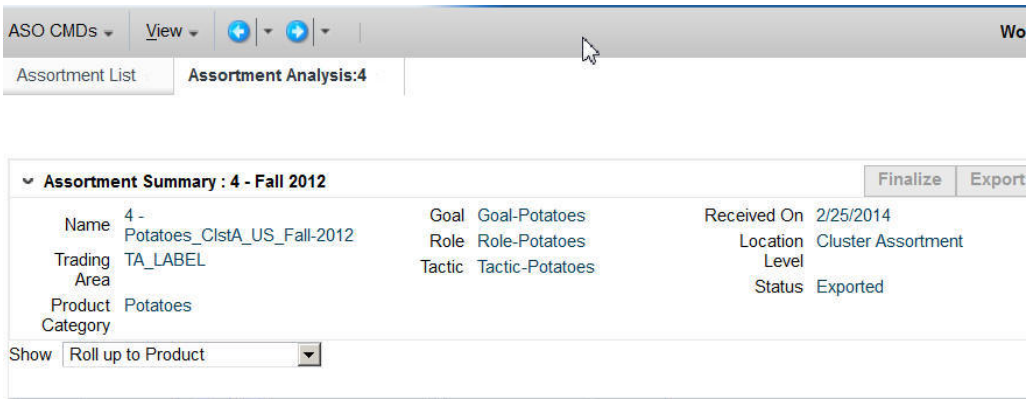
Once an assortment is finalized, the status as displayed in the Assortment List changes to Finalized.

To finalize an assortment, open it in the Assortment Analysis screen, select the runs you want to finalize, and click the **Finalize** button.

Exporting the Run Results

Once an assortment is finalized, the Export button is enabled and the Finalize button is disabled. Click the **Export** button in order to export the results. Once the assortment is exported, the results of the export are displayed in the Assortment Analysis table, as shown in Figure 6–5, and you can no longer select the runs for analysis. At this point, the Assortment Analysis display consists only of the Summary and the Results. The Select Runs for Analysis section is not displayed.

Figure 6–5 Assortment Analysis After Export



Once an assortment is exported, the status of the run as displayed in the Assortment List changes to Exported. The status is also shown on the Assortment Analysis screen, in the Summary section.

Once the assortment is finalized and exported, it is no longer available in the Select Planogram Set in Execute Space Optimization.

An assortment cannot be part of an assortment set that has been finalized. In such cases, the export button is disabled.

The runs that are finalized may represent a subset of the group of approved runs.

Assortment Analysis Metrics

The metrics displayed in the Assortment Summary list and in the Run Results list can be different. This can occur when there is more than one run, as the run results roll-up to Space Cluster or Store (optimization location) and the assortment analysis rolls up to Assortment Cluster.

Specifically, the metrics Units, Sales Revenue, Gross Profit and Lost Sales Units obtained from run results are summarized across all optimization location for all stores. These values match the assortment analysis rollup to assortment cluster when that run is the only selected. The metrics Facing and Demand actually match but are presented at different levels within the two screens. The run results bottom panel - included products shows those values as an average per store, while the assortment analysis screen shows those values just as the other metrics: across all the stores within the cluster.

Optimization List

The Optimization List is the dashboard for the ASO runs. In this tab, you can see a list of all existing runs, along with details that describe each run. The list includes runs created by other users, which you can open in read-only mode. You can create a run, copy a run, open a run, or delete a run.

Optimization List Interface

The optimization list, shown in Figure 6–6, includes the columns described in Table 6–10.

Figure 6–6 Optimization List

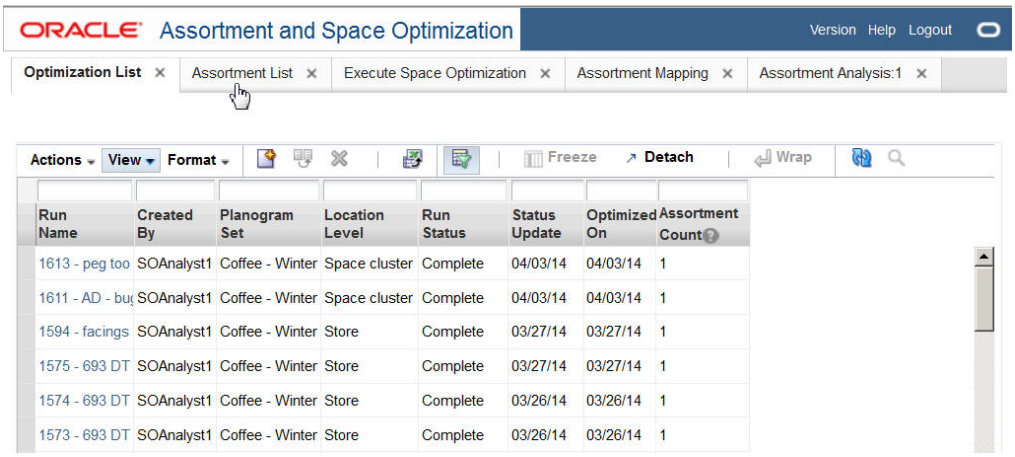


Table 6–10 Optimization List

Field	Description
Run Name	The Run Name is built by concatenating the internal Run ID and the user-provided name.
Created By	The login ID of the user who created the run.
Planogram Set	A description of the POG set. The name is a concatenation of the subcategory name and the season name.
Location Level	The location level selected for optimization (Store or Space Cluster).
Optimized On	The date when the run was optimized.
Status Update	The date when the status was last updated.
Run Status	The current status of the run. Values include Setup in Progress, Submitted, Running, Complete, Data Error, Technical Failure, Approved, and Finalized. See Table 6–11for details about each status.
Assortment Count	The number of assortments to which the run applies (this is based on how many assortments include the planogram set that is set for the run)
Description	The user-provided description for the run.

Table 6–11 Run Status Description

Status	Description
Setup in Progress	User is setting up the run, which has not yet been submitted for optimization.
Submitted	The user has submitted the run for optimization.
Running	The run optimization process is currently running.
Complete	The optimization process is complete, and the run results are ready for review.
Data Error	Execution was suspended/aborted due to data errors.
Technical Failure	An unexpected condition (technical) has caused the optimization run to abort execution.

Table 6–11 (Cont.) Run Status Description

Status	Description
Approved	The run results have been reviewed and approved for assortment analysis.
Finalized	The run results have been included in the assortment finalization and are ready for export to other applications.

Create a Run

To create a new run, click the **Create Run** icon. You are taken to the Optimization Setup stage in the Execute Space Optimization tab.

In order to create a run, you must have either Micro-space opt analyst permission or analytical super user permission.

Copy a Run

To copy an existing run, highlight that run in the displayed list and click the **Copy Run** icon. Alternatively, you can also open the run you want to copy and then choose Save As from Action menu. You are taken to the Optimization Setup stage in the Execute Space Optimization tab. The fields are populated with copies of all inputs and data. Any existing results are not included. The list of runs is updated with a run that has a new ID; all other fields match the existing run. Only the micro space opt analyst and analytical super user can create new runs.

Open a Run

To edit an existing run, highlight that run in the displayed list and click the run ID. Any user can open a run; however, only users with micro space opt analyst or analytical super user permissions can open a run in edit mode. Edit mode is only allowed for runs that a user created and only when the run does not have a status of Running, Approved, or Exported. All other runs are opened in read-only mode.

Delete a Run

Runs can be deleted depending on the status of the run and the permissions a specific user has.

A user with Administrator permissions can delete any run except for one that has a status of Setup in Progress.

Users with micro space opt analyst or analytical super user can delete any runs they have created that do not have a status of Running, Approved, or Exported. They cannot delete runs created by another user.

More than one run can be selected for deletion at the same time.

Execute Space Optimization

The Execute Space Optimization tab provides a series of four stages that you progress through in order to set up, run, and analyze the results of the optimization run:

- Optimization Setup. Used to pick a planogram set, the locations, and the optimization level.
- Fixture and Product Data. Used to view or change fixtures and product merchandise settings.

- Objective and Constraints. Used to view or choose optimizations, objectives, and constraints.
- Results and Analysis. Used to view results and override, approve, or revisit prior steps in order to make changes.

Optimization Setup

The Optimization Setup stage is used to pick a planogram set, the locations, and the optimization level. It is organized into three major sections: the Summary section, the Setup section, and the Review Optimization Locations section.

During this stage, if a planogram does not have detailed shelf information, ASO Cloud Service, using the smart start process, automatically assigns shelves to these partial shelf fixture planograms (PSFP) using configured parameters. These planograms are flagged in the Partial Fixtures column in the Planogram Fixtures tab (see Fixture and Product Data).

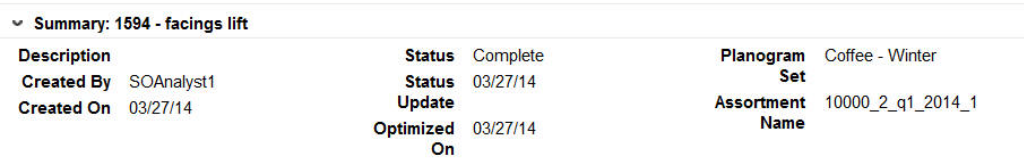
This tab is read-only if the run has been submitted to the system, if the user does not have permission to edit runs, or if the user has permission to edit runs but not create runs.

After you have made your selections, you can go to the next stage by clicking the **Next** button, submit the run for optimization, or validate the data. Use the **Action** menu to access Submit and Validate. If you submit the run for optimization at this point, the default rules and settings are used.

Summary

The Summary section, shown in Figure 6–7, identifies the run you are configuring. This section is populated once you provide the required information in the Setup section. It contains the fields listed in Table 6–12.

Figure 6–7 Summary



The Summary section is displayed at the top of each tab in the Execute Space Optimization stage.

Table 6–12 Summary Section Fields

Field	Description
Description	The name for the run.
Created By	The user ID of the person who has created the run.
Created On	The date when the run was created.
Status	The status of the run. Values include Setup in Progress, Running, Complete, Approved, and Finalized.
Status Update	The date when the status for the run was last updated.
Optimized On	The date when the run was last submitted for optimization.

Table 6–12 (Cont.) Summary Section Fields

Field	Description
Planogram Set	The name of the planogram set, which is selected from the Select Planogram Set dialog box. A planogram set name is a concatenation of the subcategory name and the season name.
Assortment Name	A comma-separated list of all assortments associated with the run.

Setup

The Setup section, shown in Figure 6–8, is used to define and configure the run and provides access to dialog boxes that you use to select the Planogram sets and the optional Template to use for the run. In addition, you identify the Location Level for the run.

Figure 6–8 Setup



Click the **Save** button to save changes to the Name and Description fields. You also use this button to save changes to the planogram set, location level, and template run fields.

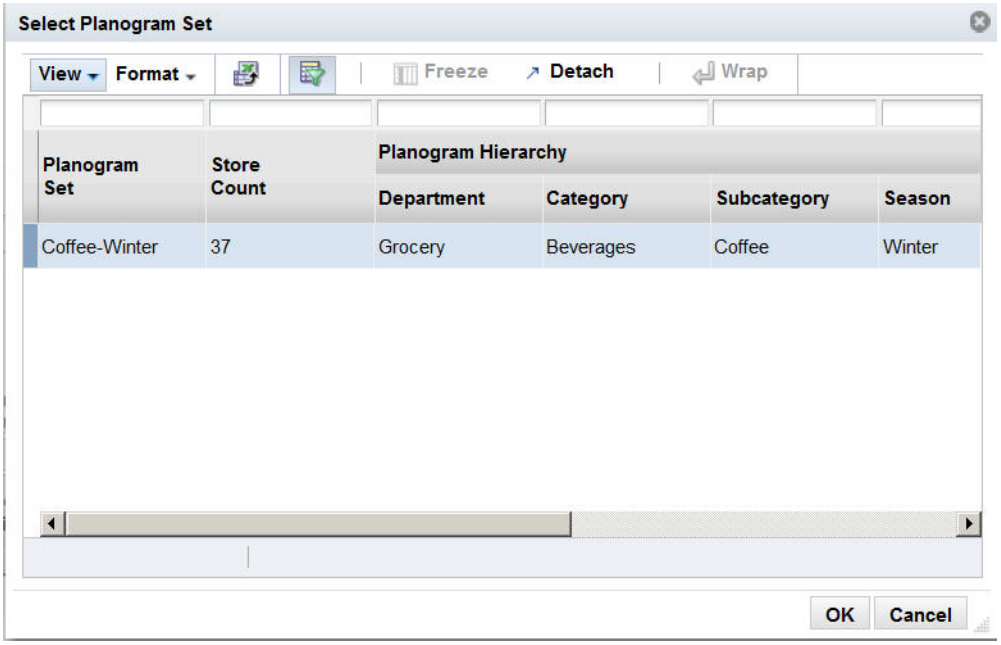
Table 6–13 Setup

Field	Description
Name	The name identifying the run. The name must be 80 characters or less and use alphanumeric characters only. It is case sensitive.
Description	A brief description of the run. The value must be 200 characters or less and use alphanumeric characters only.
Planogram Set	The field displays the name of the planogram set you select for the run. The name is a concatenation of the subcategory name and the season name. Click the **Pencil** icon to access the Planogram Set dialog box, from which you make a selection.
Template Run	The field displays the name of the template you optionally select to use when creating a new run. Click the **Pencil** icon to access the Template Run dialog box, from which you make a selection.
Location Level	Choose the Location Level for the run from the drop-down list. Values are Space Cluster and Store.

Use Additional POG Attributes Use this check box to activate or deactivate the use of the POG configuration in the clustering process.

Planogram Set The Select Planogram Set dialog box, shown in Figure 6–9, which you access by clicking the **Pencil** icon adjacent to Planogram Set, lists all the planograms that are available to you for the run. You can select only one.

Figure 6–9 Select Planogram Set



Highlight the appropriate row and click **OK**. If you change this selection after either selecting stores or building optimization location tables, then you must re-select the stores and the optimization locations for those stores.

Table 6–14 Select Planogram Set

Field	Description
Planogram Set	The display name for the planogram set. The name is a concatenation of the subcategory name and the season name.
Store Count	The number of stores associated with the planogram set.
Planogram Hierarchy Department	The planogram set belongs within a hierarchy that consists of Department, Category, Subcategory, and Season.
Planogram Hierarchy Category	The planogram set belongs within a hierarchy that consists of Department, Category, Subcategory, and Season.
Planogram Hierarchy Subcategory	The planogram set belongs within a hierarchy that consists of Department, Category, Subcategory, and Season.
Planogram Hierarchy Season	The planogram set belongs within a hierarchy that consists of Department, Category, Subcategory, and Season.
Assortment Details Product Category	A comma-separated list of the assortment IDs that the planogram set applies to.
Assortment Details ID	The external ID of the planogram set.
Assortment Details Name	A comma-separated list of the assortment external IDs that the planogram set applies to.

Template Run The Select Template Run dialog box, which you access by clicking the **Pencil** icon adjacent to Template Run, lists all the previous runs that are available to you to use as a template in creating a new run. You can select only one. Highlight the appropriate row and click **Save**. The new run you initialized uses values and settings that were used in the template run and apply to the new run. This is optional. You can use the No Template button to clear the template selection.

Table 6–15 Select Template Run

Field	Description
ID	The ID that identifies the run.
Name	The name of the run.
Description	A brief description of the run.
Planogram Set	The planogram set associated with the run.
Created By	The user ID of the person who created the run.
Run Status	The status of the run. Values include Setup in Progress, Running, Complete, Approved, and Finalized.
Status Update	The date when the status for the run was last updated.

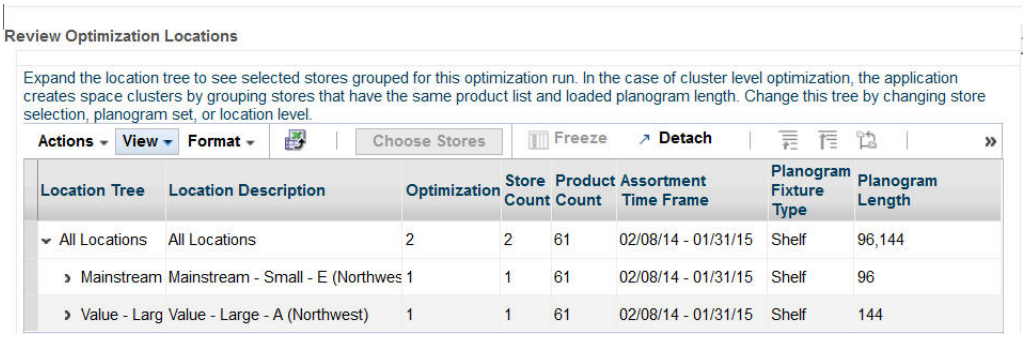
Location Level The location level identifies the level at which the optimization occurs. You select either Space Cluster or Store from the drop-down list. If you change the value of this field after building the optimization locations, then the optimization locations are re-built using the store you selected and at the new location level.

Use Additional POG Attributes Use this check box to enable space cluster generation based on the product list and POG length (base criteria) plus the additional POG attributes height, depth, and fixture type. If you select this check box and save the changes, the optimization location tables are refreshed based on the new setting.

Review Optimization Locations

The Review Optimization Locations table, shown in Figure 6–10, displays the optimization locations for the optimization.

Figure 6–10 Review Optimization Locations



You can use this table to edit the list of existing locations or add new locations. The locations are listed in a tree structure that shows the hierarchy the optimization locations are part of. If the location level is set to Space Cluster, then the lowest level of the tree is space cluster. If the location level is set to Store, then the lowest level of the tree is store.

Table 6–16 Review Optimization Locations

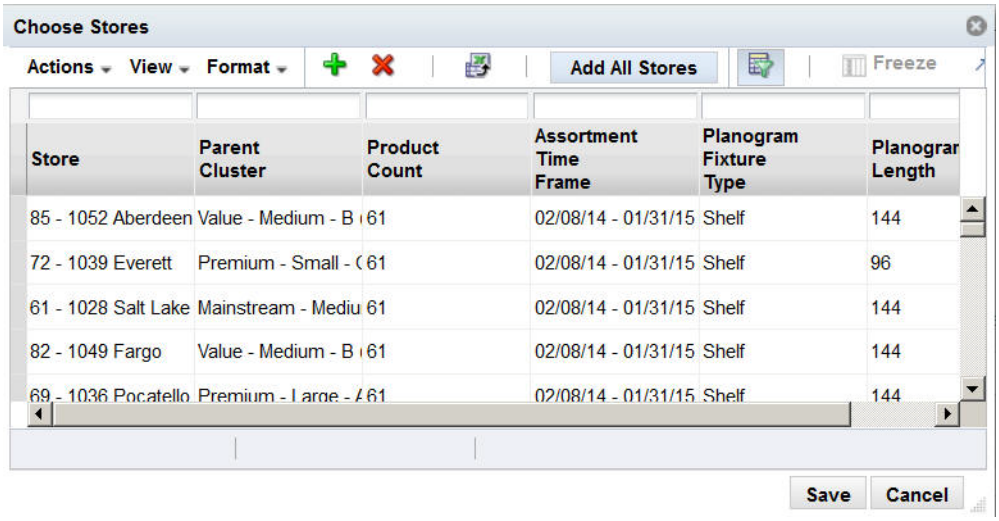
Field	Description
Location Tree	Displays the node name in the optimization location hierarchy tree. The lowest level is either space cluster of store.
Location Description	A brief description of the optimization location hierarchy node.

Table 6–16 (Cont.) Review Optimization Locations

Field	Description
Optimization Locations	The number of optimization locations under the optimization location hierarchy node. The lowest level is always a value of 1.
Store Count	The number of stores in the optimization location hierarchy node. If the location level is store, then it will be a value of 1 for the lowest level.
Product Count	The number of products for the optimization location, based on the associated assortment.
Assortment Time Frame	The start date and end date for which the optimization location's associated assortment is valid, across products.
Planogram Fixture Type	A comma-separated list of unique planogram fixture types that exist for the optimization location hierarchy node.
Planogram Length	A comma-separated list of unique planogram lengths that exist for the optimization location hierarchy node.

Choose Stores Click the **Choose Stores** button to add locations to the optimization run. The Selected Stores dialog box, shown in Figure 6–11, lists the stores that you have already selected for the optimization locations.

Figure 6–11 Choose Stores



If you have not yet selected any stores, then no stores are listed and the table is blank. This can occur if you have changed the planogram set for the run or in the case of new runs (in which stores have not yet been added).

The Choose Stores dialog box contains the following fields:

Table 6–17 Choose Stores

Field	Description
Store	The identification for the store. It is a concatenation of the store code and the store name.
Parent Cluster	The name of the cluster that the store belongs to.
Product Count	The number of products that the store carries that belong to the assortment associated with the planogram set that you selected.

Table 6–17 (Cont.) Choose Stores

Field	Description
Assortment Time Frame	The start date and end date during which the assortment associated with the planogram set you selected is valid.
Planogram Fixture Type	A comma-separated list of the unique planogram fixture types available for the associated planogram set's assortment at the store.
Planogram Length	A comma-separated list of the unique planogram lengths available for the associated planogram set's assortment at the store.

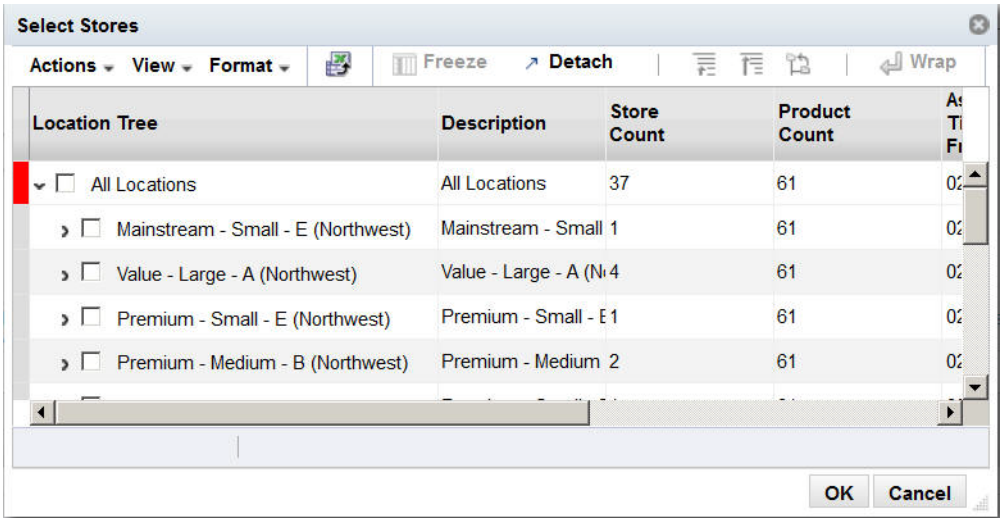
You use the Choose Stores dialog box for adding stores and deleting stores. It is enabled once you provide the required information in the Summary section. To add stores, click the **Add** icon in order to access the Adding Stores functionality. To delete stores, highlight one or more stores from the Choose Stores list and click the **Delete** icon. You can also click the **Add All Stores** button in order to add all the stores for which the planogram set you selected applies.

When you click **OK** in the Choose Stores dialog box, the list of selected stores is saved and the optimization locations are built or re-built. As long as the run is in a state in which the optimization locations can be edited, you can update the list and re-build the optimization locations.

When you click **Cancel** in the Choose Store dialog box, all the changes you have made to the list of selected stores are reverted and no changes are made to the optimization locations.

Adding Stores You add stores through the Select Stores dialog box, which you access by clicking the **Add** icon in the Choose Stores dialog box. The Select Stores dialog box, shown in Figure 6–12, displays a list of stores that can be added to the optimization run, based on the run's Planogram Set value.

Figure 6–12 Select Stores



The Select Stores dialog box contains the following fields:

Table 6–18 Select Stores

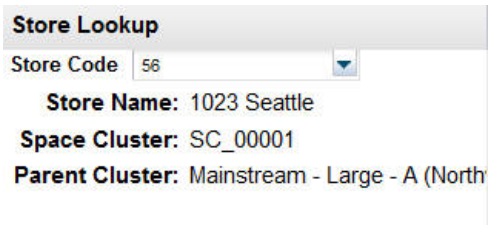
Field	Description
Location Tree	The name of the location hierarchy node. If it is the lowest level, it is space cluster or store.
Description	A brief description of the location hierarchy node.
Store Count	The number of stores in the optimization location hierarchy node. If the location level is store, then it will be a value of 1 for the lowest level.
Product Count	The number of products in the location, based on the associated assortment.
Assortment Time Frame	The start date and end date during which the location's associated assortment is valid.
Planogram Fixture Type	A comma-separated list of unique planogram fixture types for the location hierarchy node.
Planogram Length	A comma-separated list of unique planogram lengths for the location hierarchy node.

Contextual Information

The following information is only available when the optimization locations have been defined.

Store Lookup You can use Store Lookup, shown in Figure 6–13, to access information about a store. You must provide the store code, either by entering it into the text field provided or by selecting it from the drop-down list. This list is populated based on the optimization location selected for the run.

Figure 6–13 Store Lookup



Once you select the store, you see the following information:

- Store name
- Optimization location (if location level is Space Cluster)
- Parent cluster

Space Cluster Lookup You can use the Space Cluster Lookup, shown in Figure 6–14, to see the list of stores associated with a specific space cluster. You must provide the name of the space cluster, either by entering it into the text field provided or be selecting it from the drop-down list. The list is populated based on the optimization location selected for the run.

Figure 6–14 Space Cluster Lookup



Once you select the space cluster, you see a list of the component stores.

Charts Two charts are provided in which you can see the variation in store counts by optimization location and product counts by optimization location. The charts are populated with data once the optimization location has been configured. Figure 6–15 provides an example of the store counts chart and Figure 6–16 provides an example of the product counts chart.

Figure 6–15 Store Counts

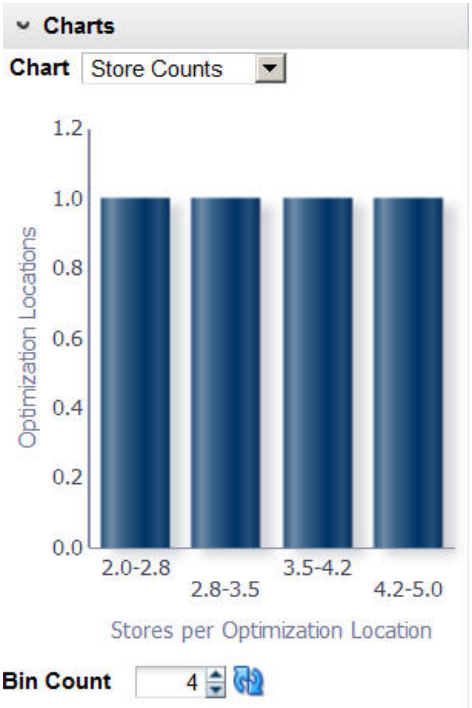
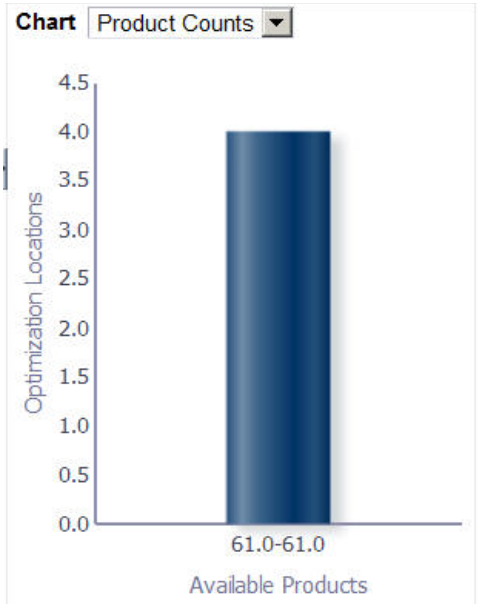


Figure 6–16 Product Counts



Fixture and Product Data

The Fixture and Product Data stage is divided into a Summary section, a Planogram Fixtures tab, and a Product Merchandising Data tab. The Summary section identifies the characteristics of the run. The Planogram Fixtures tab is used to review the available planogram lengths for each optimization location and to update the selected configuration for each available length from a set of configurations for each length. The Product Merchandising Data tab is used to assign merchandising options to products. These assignments can vary by product or optimization location.

This tab can only be edited by the user who created the run and only before the run has a status of Running, Approved, or Finalized.

Summary

The Summary section, shown in Figure 6–7, identifies the run you are configuring.

Table 6–19 Summary Section Fields

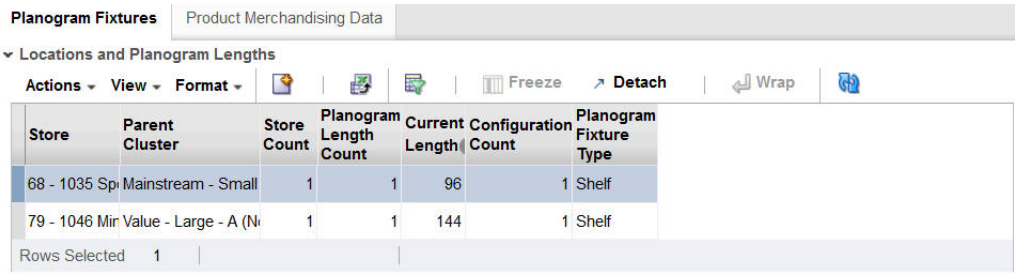
Field	Description
Description	The name for the run.
Created By	The user ID of the person who has created the run.
Created On	The date when the run was created.
Status	The status of the run. Values include Setup in Progress, Running, Complete, Approved, and Finalized.
Status Update	The date when the status for the run was last updated.
Optimized On	The date when the run was last submitted for optimization.
Planogram Set	The name of the planogram set, which is selected from the Select Planogram Set dialog box.
Assortment Name	The name of the assortment associated with the planogram set.

Planogram Fixtures

The Planogram Fixtures tab contains two sections: the Locations and Planogram Lengths table and the Available Configurations table.

Locations and Planogram Lengths This table, shown in Figure 6–17, displays the run's optimization location and provides metrics about the available planogram configurations for each run.

Figure 6–17 Locations and Planogram Lengths



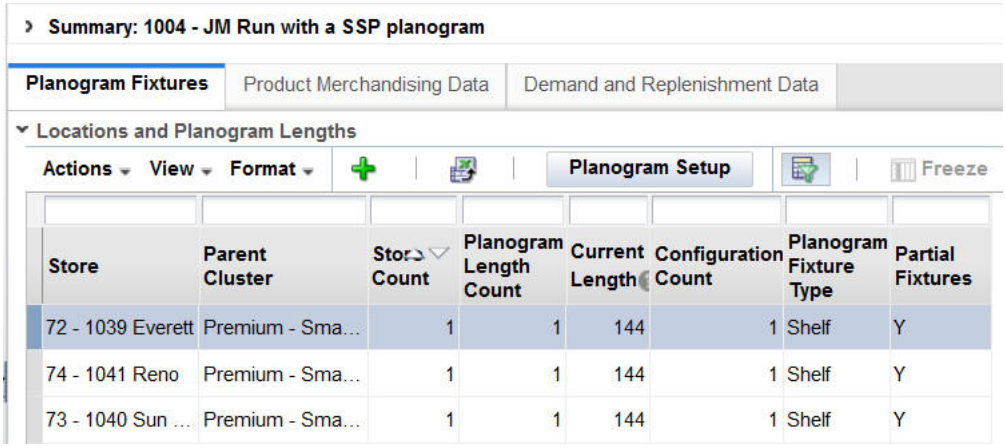
You can create new lengths, which can be either shorter or longer than the current lengths (subject to validation), and refresh the table display.

Table 6–20 Locations and Planogram Lengths

Field	Description
Space Cluster/Store	The name of the space cluster or store.
Parent Cluster	The name of the parent cluster of the optimization location.
Store Count	The number of stores in the optimization location.
Current Length	The length of the configuration that is currently ready for optimization.
Planogram Length Count	The number of unique planogram lengths in the optimization location.
Configuration Count	The number of planogram configurations in the optimization location.
Planogram Fixture Type	A comma-separated list of the unique fixture types present across all configurations that are ready for optimization. This list includes fixture types that have been created by the user.
Partial Fixtures	Y indicates that the planogram does not have detailed shelving data and that the fixture configuration has been generated based on configuration parameters. N indicates that the planogram has complete shelving data.

Planogram Setup

Figure 6–18 Planogram Setup Button



Select **Planogram Setup** to access the Planogram Setup dialog box. This button in enabled after you select at least one optimization location row from the table. The Planogram Setup dialog box contains the parameters used during the smart setup process to generate the Partial Shelf Fixture Planogram (PSFP). You can edit these parameters and rerun this process using this dialog box. These parameters are described in Table 6–21. Note that the allowable minimum and maximum values permitted in this dialog box are configured in the database.

Figure 6–19 Planogram Setup Dialog Box

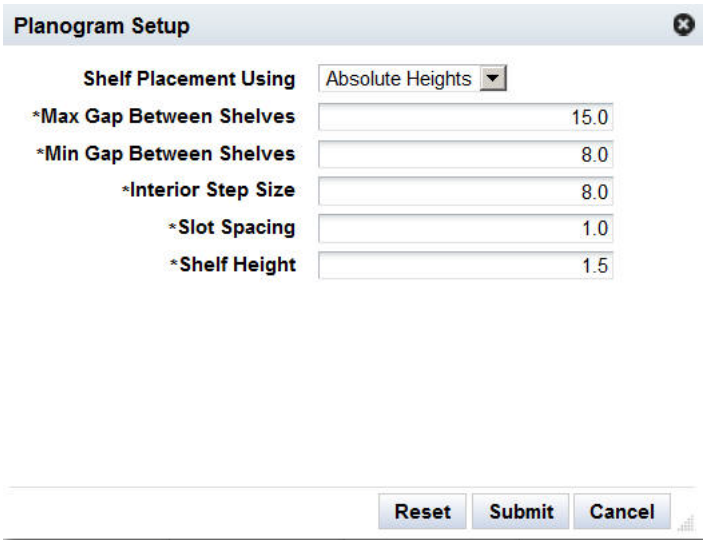


Table 6–21 Planogram Setup Parameters

Field	Description
Shelf Placement Using	This drop-down menu as two options, Height Percentile and Absolute Height. The choice you make between these two options determines the types of values (percentage or numeric) used by some of the parameters and which parameters are displayed.

Table 6–21 (Cont.) Planogram Setup Parameters

Field	Description
Max Gap Between Shelves	When using Height Percentile: Select the Nth percentile from the distribution of product heights in the assortment as the maximum gap between two shelves. When using Absolute Height: Use the provided value as the maximum gap between two shelves
Min Gap Between Shelves	When using Height Percentile: Select the Nth percentile from the distribution of product heights in the assortment as the minimum gap between two shelves. When using Absolute Height: Use the provided value as the minimum gap between two shelves
Interior Step Size	When using Height Percentile: Reduce the height percentile to use for each subsequent shelf placement (e.g. with a max of 100% and step of 20%, the next shelf placement will use the 80th percentile of product heights to determine the gap) When using Absolute Height: Determine the gap between interior shelves using the step size
Absolute Min Gap	The absolute minimum gap between two shelves, regardless of product height percentiles.
Slot Spacing	Distance between slots in the bay. Final shelf placement will round to the nearest slot position.
Shelf Thickness	Thickness of all shelves placed using the Smart Start Process. Shelf thickness is considered when calculating the gaps between shelving.

The following buttons are available in this dialog box:

- **Reset** – Use this button to update the values for all the parameters to the configured default values. You must then click **Submit** to rerun the smart setup process.
- **Submit** – Use this button to rerun the smart start process for the optimization locations you selected, based on the parameter values you entered here. When you click this button, the dialog box closes and the available configurations table is refreshed with the new configurations generated by the smart setup process. The new values that you provided for the parameters are saved at the optimization location level.
- **Cancel** – Use this button to close the dialog box without making any changes.

If some of the locations you select do not have partial fixtures, they are not processed by the smart setup process. You will see a message if this occurs.

If some of the locations you select also have overrides that apply, these overrides are also applied. You will see a warning message if this is the case.

Planogram Setup Shelf Placement

When planogram setup parameters are submitted, the system calculates shelf positions using the selected process (percentile or absolute). The shelf placement rules for each process are summarized in Table 6–22 and Table 6–23 in the order in which they are applied.

Table 6–22 Height Percentile Shelf Placement Rules

Shelf Position	Rule Description
Top Shelf	Place the shelf at (fixture height - (top shelf stacking limit + shelf thickness)), rounded down to the nearest slot position.
Bottom Shelf	The bottom shelf goes in position 0.
Interior Shelves	Starting from the bottom, calculate the gap to the next shelf using product height percentiles. The first interior shelf is placed using the product height at the Max Gap percentile, and each subsequent gap is decreased based on the step size percentile, rounded up to the nearest slot.

Table 6–23 Absolute Height Shelf Placement Rules

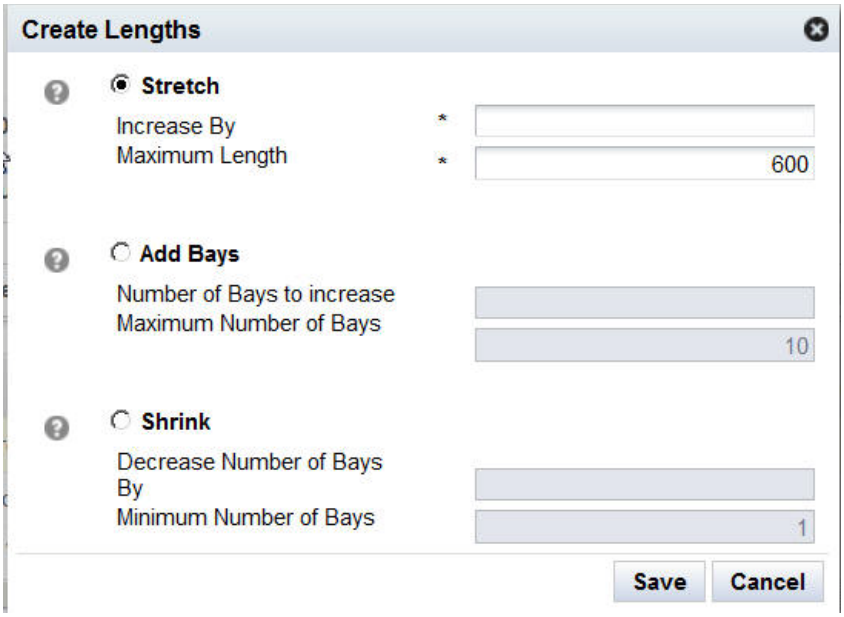
Shelf Position	Rule Description
Top Shelf	Place the shelf at (fixture height - (top shelf stacking limit + shelf thickness)), rounded down to the nearest slot position.
Bottom Shelf	The bottom shelf goes in position 0.
Interior Shelves	Starting from the bottom, calculate the gap to the next shelf using the max gap values and step size. The first interior shelf is placed using the max gap size, and each subsequent gap is determined by the step size, rounded up to the nearest slot.

The same process is applied to all empty bays in the planogram, such that multiple bays of the same height have identical shelving. Other fixture types (such as pegboards) are not considered during this process, and bays that already have fixtures are not modified. After the process has been run, the user can override the results by editing the planogram directly. See Fixture Overrides for details.

Create Lengths

Select **Create** from the Actions menu to access the Create Lengths dialog box, shown in Figure 6–20. You use this to create new planogram configurations for the optimization locations you select. This involves stretching or shrinking a configuration that is marked for optimization. In addition, you can add new bays to the right. If you want to create new lengths for a configuration that has not been marked for optimization, you must mark that configuration for optimization. but it must be a loaded configuration, which can be identified by the fact that the Parent Planogram field is empty or null.

Figure 6–20 Create Lengths



Stretch

To stretch the length, select the **Stretch** radio button. Stretching changes the planogram by adding length to the last bay (including internal objects). It adds as much length as defined within the Increase By field up to the maximum length.

A new configuration is created that is a copy of the base configuration. The planogram is shrunk by removing the bays from right to left. The length, in generic measurement length units, is adjusted based on the bays that are removed. That copy is then modified based on the values you enter in the following two fields:

Table 6–24 Create Lengths: Stretch

Field	Description
Increase By	Use to specify how much to stretch the base configuration in order to create the new configuration. The value must be less than or equal to the value you enter in the Maximum Length field. The unit of measurement is agnostic.
Maximum Length	Use to specify the maximum length possible for the new configuration. The default value that is displayed is configurable during implementation. You can change the default value, but the new value must be less than or equal to the default.

Add Bays

To add bays, select the **Add Bays** radio button. The add bay process creates a copy of the right-most bay, including the fixtures, shelves, pegs, and freezers. The copy is added after the last existing bay on the right. You determine the number of copies added by entering the appropriate value into the dialog box, as described in Table 6–25.

Table 6–25 Create Lengths: Add Bays

Field	Description
Number of Bays to Increase	Indicates the number of bays to add. Must be an integer value.

Table 6–25 (Cont.) Create Lengths: Add Bays

Field	Description
Maximum Number of Bays	Indicates the maximum number of bays the new POG is permitted to have. This adds as many bays as indicated in the Number of Bays parameter, as long as the overall bay count for the POG does not exceed the maximum value provided for this parameter. Must be an integer value.

Note that adding bays has the effect of increasing the number of shelves, which in turn changes the count of shelves displayed in the UI. Other metrics may also change.

Shrink

To shrink the length, select the **Shrink** radio button. Shrinking changes length (decrease) by eliminating starting from the last bay as many bays as defined by the Decrease Number of Bays, making sure that at least the minimum number of bays remain.

A new configuration is created that is a copy of the base configuration. The changes are made to the bays (not to the length). The copy is modified based on the values you enter in the following two fields:

Table 6–26 Create Lengths: Shrink

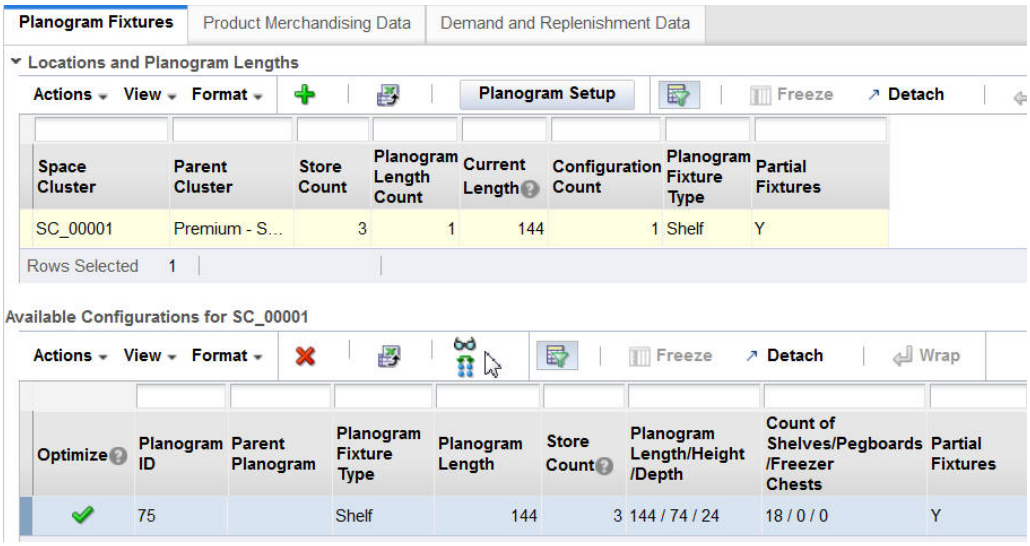
Field	Description
Decrease Number of Bays By	Use to specify, in terms of the number of bays, how much to shrink the base configuration in order to create the new configuration.
Minimum Number of Bays	Use to specify the minimum number of bays for the new configuration. The default value that is displayed is configurable during implementation. You can change the default value, but the new value must be less than or equal to the default.

If the values you specify are the same as those for a configuration (same parent planogram and same length), no new configuration is created. If the values you specify do not fall within the defined minimum and maximum values for shrinking or stretching, the new configuration will be defined by the minimum or maximum value instead.

The new configuration is marked for optimization automatically by the system. Other configurations with the same length are un-marked. If the new configuration is actually the same as the existing one, the existing configuration is marked for configuration.

Fixture Overrides You can make certain modifications to the run-specific POG data that is used in optimization. These modifications include adding, deleting, editing, and moving shelves associated with shelf fixtures as well as restoring prior settings.

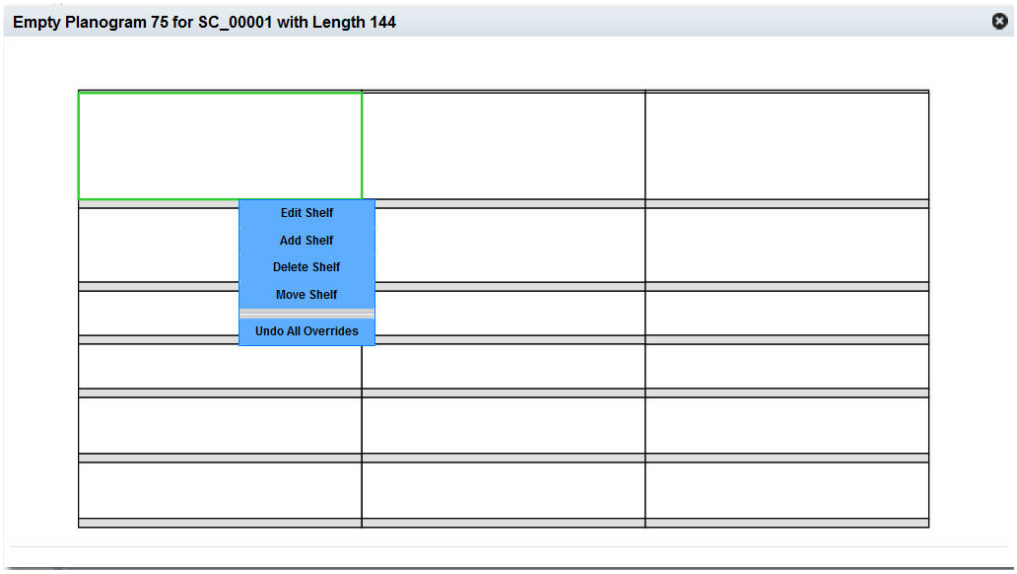
Figure 6–21 Fixture Overrides



Fixture overrides exist at the optimization level. This means that store-level optimization requires store-by-store fixture overrides and cluster level optimization requires cluster-by-cluster fixture overrides.

To make these modifications, highlight the space cluster of interest and click the **Goggle** icon. An empty planogram is displayed. Right-click on a shelf to access the Edit dialog box.

Figure 6–22 Edit Dialog Box for Fixture Overrides



Add

You can add a new shelf to an existing fixture. To do this, select the fixture and then use the right-click menu to access the Add functionality. You see a dialog box the shows defaults. You can change the depth, thickness, and elevation of the new shelf. Note that the new shelf must fit within the existing fixture and it cannot overlap with other shelves.

Delete

You can delete a shelf from an existing fixture. To do this, select the shelf you want to delete and then use the right-click menu to access the Delete functionality. You will be asked to confirm the deletion.

After you delete the shelf, the space where the shelf existed will no longer show a shelf. If a bottom shelf was deleted, then the shelf and the space up to the next shelf (or top of fixture if no shelves) will be replaced by a grey background.

Edit

You can edit an existing shelf. To do this, click on a shelf and use the right-click menu to access the Edit functionality. You see a dialog box where you can change the depth and the thickness of the shelf.

Move

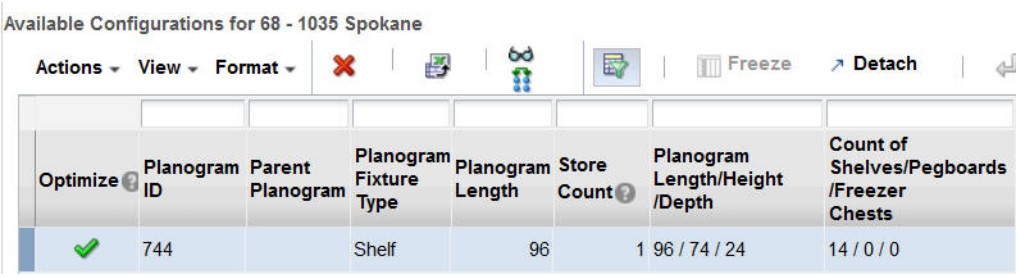
You can move an existing shelf to a new location. To do this, select the shelf you want to move and use the right-click menu to access the Move functionality. You see a dialog box that you can use to enter a new value for elevation.

Restore

You can restore a POG to the most recently loaded or updated version of that historical POG, based on the most current data feed updates made to the historical POG configuration. You cannot restore POGs that were created per macro changes. In such cases, you can delete the run and start over with the most recently loaded or updated version of that historical POG.

Available Configurations This table, shown in Figure 6–23, displays details about the configuration available for the row you select in the Locations and Planogram Lengths table. If you select more than one row, you will see the message "No Data to Display."

Figure 6–23 Available Configurations



The Available Configurations table contains the fields listed in Table 6–27.

Table 6–27 Available Configurations

Field	Description
Optimize	Indicates that the configuration is marked for optimization.
Planogram ID	The ID that identifies the configuration (planogram).
Parent Planogram	The base planogram configuration used to create the new configuration. It is always a loaded configuration. The value is Null for an externally loaded configuration.
Planogram Fixture Type	A comma-separated list of the unique fixture types in the configuration.

Table 6–27 (Cont.) Available Configurations

Field	Description
Planogram Length	The length of the planogram.
Store Count	The number of stores that use the configuration for this run's assortment.
Planogram Length/Height/Depth	The dimensions of the configuration.
Count of Shelves/Pegboards/Freezer Chests	Count for the number of each type of fixture that exists in the configuration.

You can perform the following actions with respect to the Available Configurations table. You can initiate each action from the Actions menu or in some cases by clicking the appropriate icon.

Mark for Optimization

You must mark a configuration to make it available for optimization. Only marked configurations are displayed in the Results and Analysis tab. Note that only one configuration for a given length can be marked for optimization. If one configuration for a given length is currently marked for optimization, and you select another configuration of the same length, the first configuration is un-marked.

Delete

You can delete any user-created lengths (identified by non-null values in the parent Planogram column) in the table. If you delete a configuration that is marked for optimization, all configuration of that configuration's length will be deleted.

View Configuration

You can view a diagram of the selected configuration by clicking the **View Configuration** icon.

Compare to History

Click the **Compare to History** icon to see a diagram of the historical planogram used for the selected configuration.

Contextual Information A chart is provided that displays summary data for the planogram lengths that are being used for the optimization. This information can help you understand what the current macro space allocation is and what has been added.

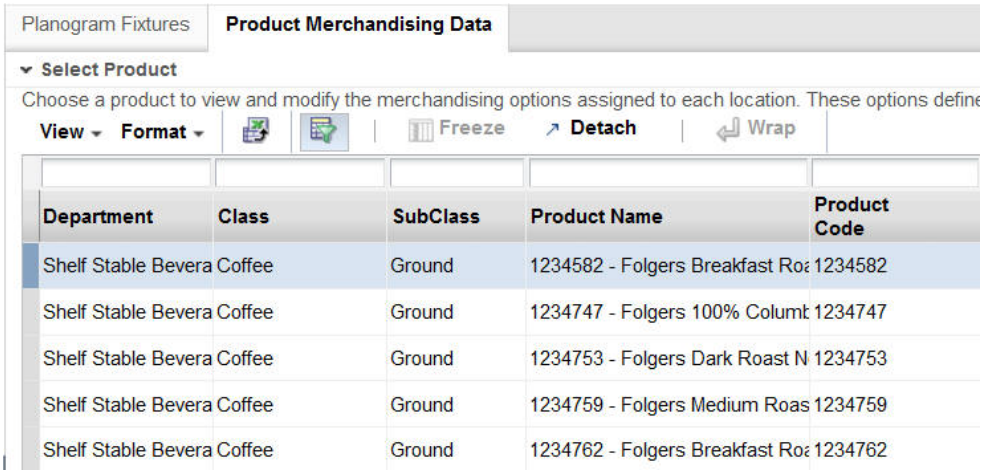
Three series for this chart are available. One series shows the count of current lengths and a second series shows the count of added lengths for a given planogram length. If there are no added lengths, then the chart should only show the current lengths. The count is across all optimization locations. The total count of all current lengths should be the number of optimization locations. You can also see product heights, charted by percentage or numeric values, and representing the entire run.

Product Merchandise Data

The Product Merchandise Data tab contains two main sections: The Select Products table and the Merchandising Options table. This stage addresses how products are placed on fixtures.

Select Product This table, shown in Figure 6–24, lists all available products. Each product is identified by the details of its hierarchy (level 1-3, department, class, and sub-class), as well as product name and product code.

Figure 6–24 Select Product



Merchandising Options This table, shown in Figure 6–25, displays the merchandising options for all optimization locations for the product you select in the Select Product Table table. A merchandising option is the combination of fixture type, display style, main orientation, and capping orientation.

Figure 6–25 Merchandising Options

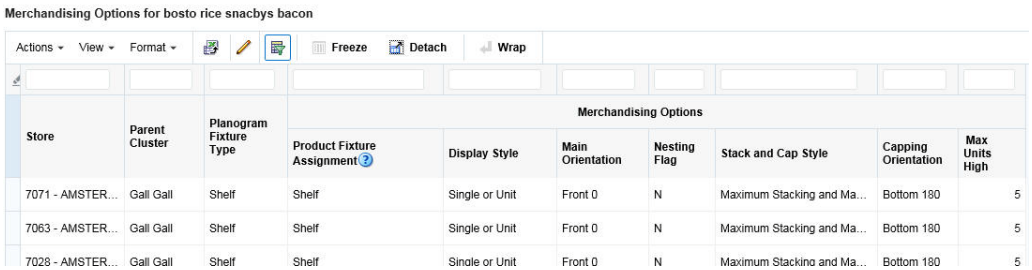


Table 6–28 lists the fields for Merchandising Options.

Table 6–28 Merchandising Options

Field	Description
Space Cluster/Store	Indicates the location level for the product.
Parent Cluster	The name of the parent cluster.
Planogram Fixture Type	A comma-separated list of the unique fixture types in the configuration.
MO: Product Fixture Assignment	The type of fixture.
MO: Display Style	The type of display unit.
MO: Main Orientation	How the product faces the front of the fixture.
MO: Nesting Flag	Indicates that the product will be nested.

Table 6–28 (Cont.) Merchandising Options

Field	Description
MO: Stack and Cap Style	Indicates the stacking and capping style to be used for the product.
MO: Capping Orientation	The orientation for products that are stacked.
MO: Max Units High	Sets the limit for the number of the product units vertical height.

You can edit the merchandising options if you have the appropriate permissions.

1. Select a product from the Select Product table.
2. Select one or more optimization locations from the Merchandising Options table.
3. Click the **Edit** icon to access the Edit Merchandising dialog box.
4. You see the possible options for each of the Merchandising Options that are available for the product you selected. For each option you select the **Select** check box in order to edit that option. The columns you can edit correspond to the columns that are displayed in the table.

Contextual Information The following three sections providing additional information about product merchandising.

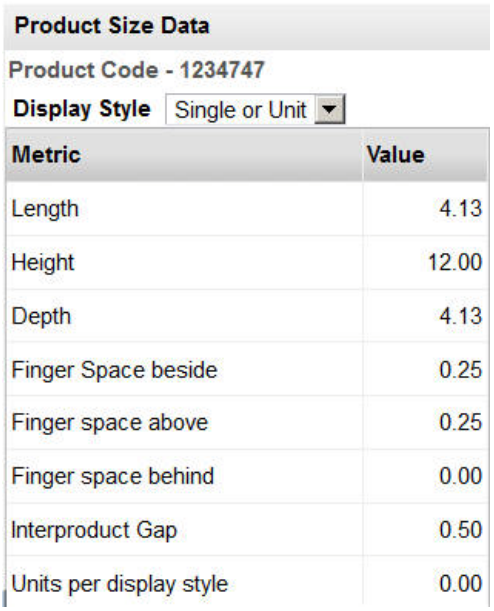
Product Size Data This table, shown in Figure 6–26, displays the sizes of the products for each of the display styles as well as the related spacing data. The data that is displayed depends on the product you select from the Select Product table. The data includes length, height, depth, finger space beside, finger space above, finger space behind, inter-product gap, and units per display style.

Finger space beside refers to the space that is needed by a specific product on the left and right side of each unit.

Finger space above refers to the space that is needed by a specific product between the top of a unit and the bottom of the shelve above it.

Finger space behind refers to the space that is needed by a specific product to the front and to the back of a unit.

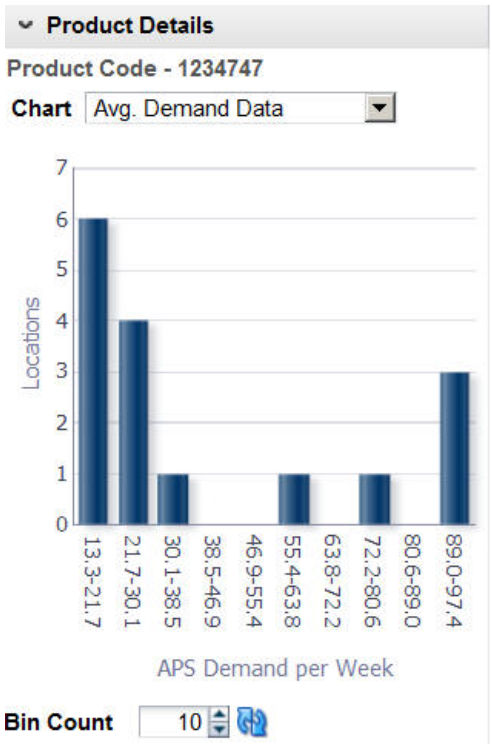
Figure 6–26 Product Size Data



Product Details This table, shown in Figure 6–27, describes the overall demand and replenishment data for the selected product. The data displayed is static and is based on the product you select from the Select Product table.

- Average Demand Data. The size of the histogram bar represents the number of optimization locations.
- Price. The size of the histogram bar represents the number of optimization locations.
- Casepack Size. The size of the histogram bar represents the number of optimization locations.
- Replenishment Frequency. The size of the histogram bar represents the number of optimization locations.

Figure 6–27 Product Details



Location Statistics This table, shown in Figure 6–28, describes the user demand and replenishment data for the selected product and optimization location. You can use this information to help in making decisions at the lowest level. The data includes store count, APS demand, price, case pack size, replenishment frequency, Item Performance Index (IPI) score, and priority. The data displayed is based on the product you select from the Select Product table and the location you select from the Merchandising Options table.

Figure 6–28 Location Statistics

Location Statistics

1234753 / SC_00001

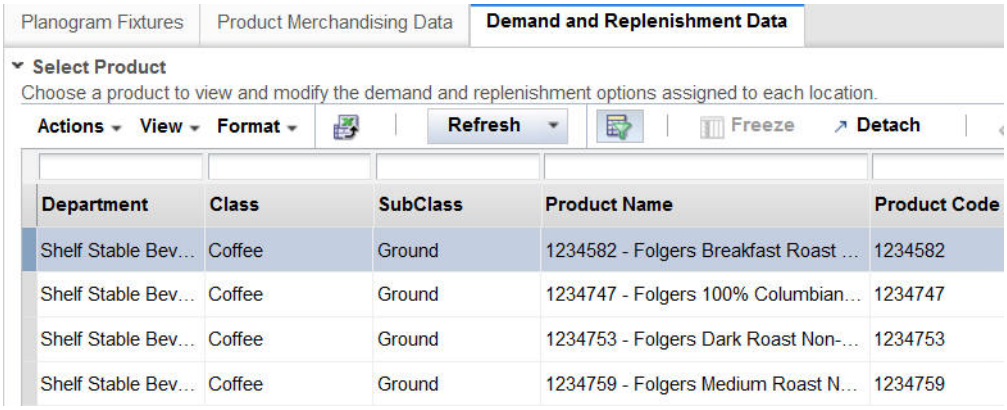
Metric	Value
Store Count	3
APS Demand	35.247
Price	6.8704
Casepack size	6
Replenishment frequency	2
IPI score	1.1774
Priority	Optional

Demand and Replenishment Data

The Demand and Replenishment Data tab lets you refresh the demand and replenishment parameters, make overrides to selected parameters, and understand the replenishment settings and their effects on results. This tab provides a Select Product table that displays all available products and a Demand and Replenishment Options table that displays the replenishment details for the product you select in the list of available products.

Select Product Table The Select Product table lists the available products and includes the following information about each product: Department, Class, Subclass, Product Name, and Product Code. When you select a specific product from this table, the Demand and Replenishment Options table is populated with the replenishment details for the product you select.

Figure 6–29 Select Product Table



Action Menu

The Action menu provides access to two functions: Demand and Replenishment.

Click **Demand** to refresh all the underlying demand data for all the products in the current run. You see updated data derived from data feeds that occurred after the most recent run was created.

Click **Replenishment** to refresh all the underlying replenishment data for all the products in the current run. This action overrides any changes you have made to replenishment data.

Refresh Button

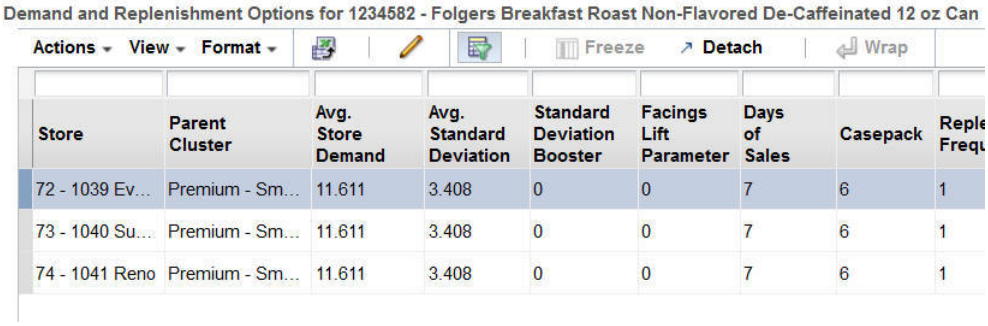
The **Refresh** button provides the same functionality as the Action menu.

Click **Demand** to refresh all the underlying demand data for all the products in the current run. You see updated data derived from data feeds that occurred after the most recent run was created.

Click **Replenishment** to refresh all the underlying replenishment data for all the products in the current run. This action overrides any changes you have made to replenishment data.

Demand and Replenishment Options Table The Demand and Replenishment Options table displays one row for each location that carries the product you select in the Select Product table. Each row contains the aggregated data for all the stores in that optimization location.

Figure 6–30 Demand and Replenishment Options



The following information is displayed:

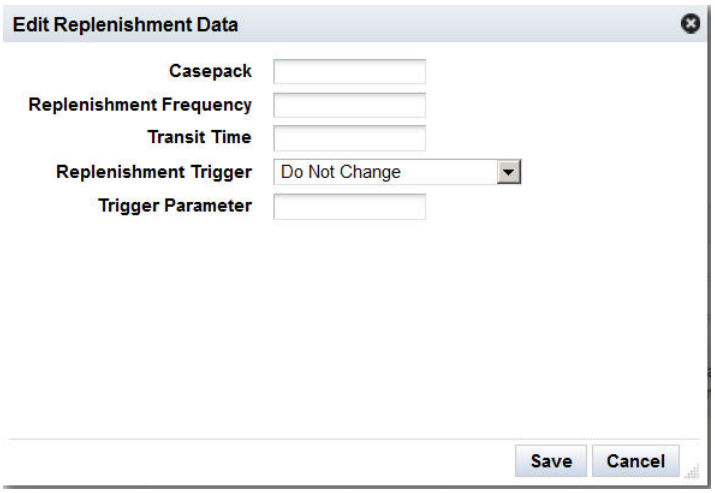
Table 6–29 Demand and Replenishment Options Table

Column Name	Description
Space Cluster	Optimizations are run at the store level or the space cluster level. A space cluster is an ad hoc group of stores with the same assortment cluster, product list, and current planogram length.
Parent Cluster	The name of the cluster that the store belongs to.
Average Store Demand	This is a weekly average.
Average Standard Deviation	A higher value indicates the need for higher safety stock to meet service levels.
Standard Deviation Booster	This is calculated across all stores in the optimization location.
Facings Lift Parameter	The impact of facings on demand.
Days of Sales	The number of days in a week that are available for store sales.
Casepack	This is calculated as a weighted average of store casepacks. It is weighed by total adjusted store demand within an optimization location. The number of units in a case.
Replenishment Frequency	The number of replenishments per shelf per week.
Replenishment Type	DC/Vendor or Backroom.
Transit Time	The number of days an order takes to go from the source location to the shelf.
Replenishment Trigger	Demand-based, target capacity-based, or case pack-based.
Trigger Parameter	Trigger value for replenishment trigger.

Editing the Demand and Replenishment Options Table

To edit the demand and replenishment data in the table, select the row or rows containing the data you want to edit and then click **Edit**. You see the following dialog box.

Figure 6–31 Edit Replenishment Data



You can edit the following fields. If you leave a field blank, then no changes will be made to that value.

Table 6–30 Edit Demand and Replenishment Data

Field Name	Description
Casepack	Value must be greater than zero.
Replenishment Frequency	Value must be greater than zero.
Transit Time	Value must be greater than zero.
Replenishment Trigger	Choose from the following: Do Not Change, Cover Demand, Percent Capacity, Casepack Fit.
Trigger Parameter	Value must be a percentage greater than zero and less than or equal to 100.

Click **Save** when you are finished editing the replenishment data. The data displayed in the Demand and Replenishment Options table is updated. The contextual information is also updated.

Contextual Information The Product Detail chart illustrates the overall demand and replenishment for the selected product. You can select the type of data to display from the drop-down menu: Average Demand Data, Standard Deviation Booster, Facings Lift Parameter, Days of Sales, Casepack Size, Replenishment Frequency, Replenishment Type, Transit Time, Replenishment Trigger, and Trigger Parameter.

Objective and Constraints

The Objectives and Constraints stage is used to modify or set the constraints (business rules) for an optimization run and to display the settings associated with the run. You specify three types of optimization constraints, product level constraints, product group constraints, and planogram constraints. Even though all constraints are ultimately applied at the optimization location level, you generally set or review the data at higher location levels.

In addition, you also pick an objective for the optimization. The objective defines what metric is used as the goal of the optimization.

Once you are satisfied with all the settings in this stage, you can validate a run or submit the run for optimization. The run closes and the optimization list is displayed. When the status for the run is Complete, you can view the results.

Summary

The Summary view, shown in Figure 6–7, provides identifying information for the optimization run.

Table 6–31 Summary Section Fields

Field	Description
Description	The name for the run.
Created By	The user ID of the person who has created the run.
Created On	The date when the run was created.
Status	The status of the run. Values include Setup in Progress, Running, Complete, Approved, and Finalized.
Status Update	The date when the status for the run was last updated.
Optimized On	The date when the run was last submitted for optimization.
Planogram Set	The name of the planogram set, which is selected from the Select Planogram Set dialog box.
Assortment Name	The name of the assortment associated with the planogram set.

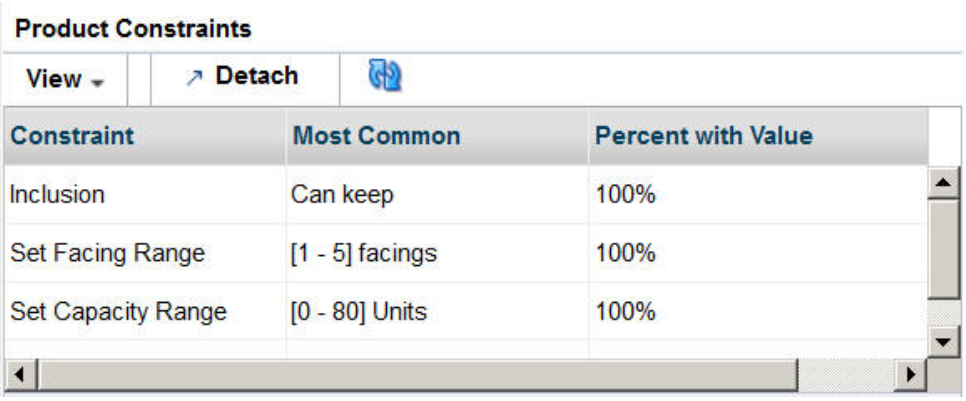
Objectives and Constraints Summary

When you select the **Summary** radio button, you see the following information displayed.

Settings Across All Locations and Products This section contains four panels that display a summary of the settings.

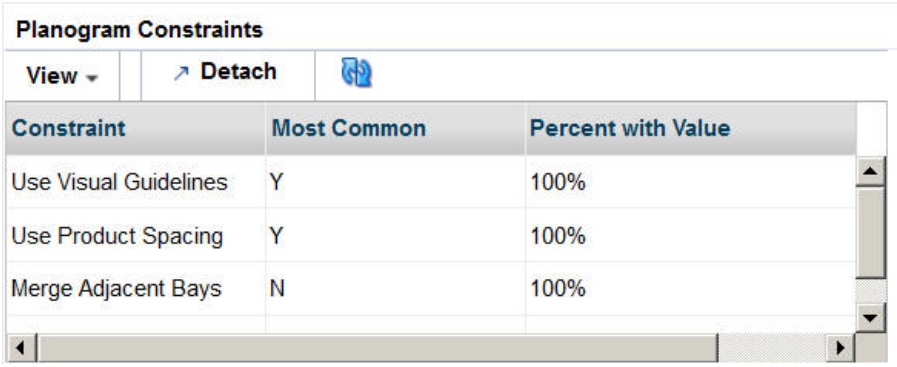
Product Constraints This table, shown in Figure 6–32, lists each product constraint, along with values for the most common value, the percentage with the most common value, and the percentage of optimization locations that have the most common value. The product constraints include Inclusion, Facing Range, Capacity Range, and Elevation Range.

Figure 6–32 Product Constraints



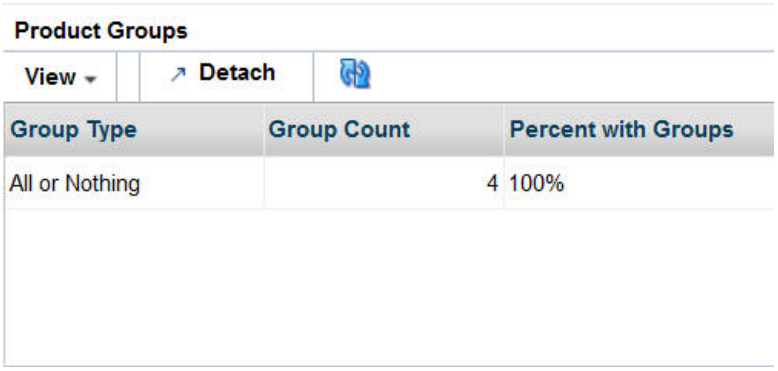
Planogram Constraints This table, shown in Figure 6–33, lists each planogram constraint, along with values for the most common value and the percentage of optimization locations that have the most common value. The planogram constraints include Use Visual Guidelines, Use Product Spacing, Merge Adjacent Bays, and Set Usable Space.

Figure 6–33 Planogram Constraints



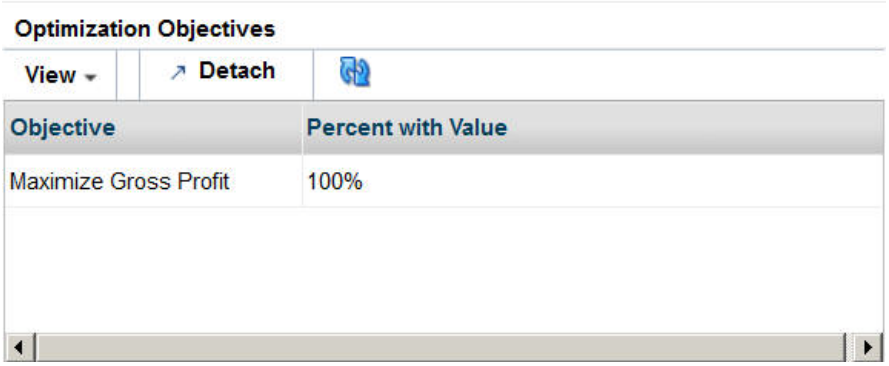
Product Groups This table, shown in Figure 6–34, lists each type of group that is used, along with values for the total number of product group/optimization location combinations for each group and the number and percentage of optimization location that have a group of that type. The product groups include All or Nothing, Match Facings, Same Shelf, and Choose From. This table does not list product groups created at levels above the optimization location, even though those product groups may be inherited by the optimization locations.

Figure 6–34 Product Groups



Optimization Objectives This table, shown in Figure 6–35, lists the objective functions that are being used, along with the number and percentage of optimization locations that use each objective function. The objective functions include Maximize Sales Units, Maximize Sales Values, Maximize Margin Value, Maximize Total Sales/On Hand Value, Maximize Sales Units (Weighted), Maximize Sales Value (Weighted), Maximize Margin Value (Weighted), and Maximize Total Sales Value/On Hand Value (Weighted).

Figure 6–35 Optimization Objectives



Select Locations and Products

When you select the **Details** radio button, this section displays. It consists of two tables: Select Locations and Select Products, shown in Figure 6–36 and Figure 6–37. You must select a row from the location tree and a product row in order to display details about your selections in Objectives and Constraints details section.

The Locations table has three levels. When the first level is selected, the product constraints details across all optimization locations are displayed in the Product Constraints table. When the second level is selected, only the product constraints details for the second level location are displayed. When any optimization location is selected, only the product constraint details for the selected optimization location are displayed. In the same way, the product constrains in the Products table are displayed. So product constraint details are fetched after both the location and the product levels are select, location first and product second.

Figure 6–36 Select Locations

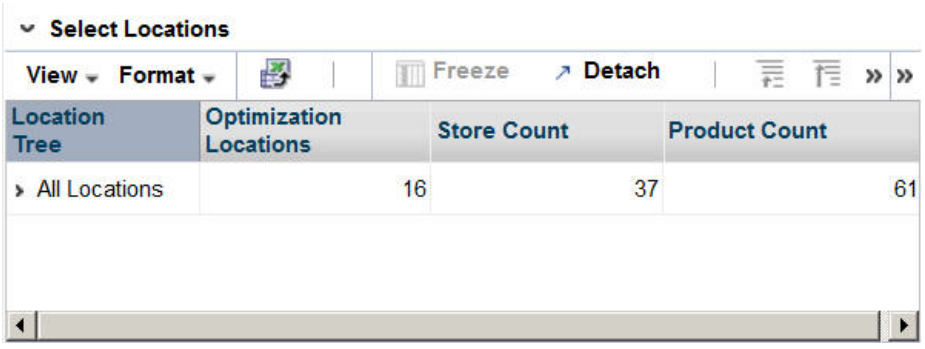
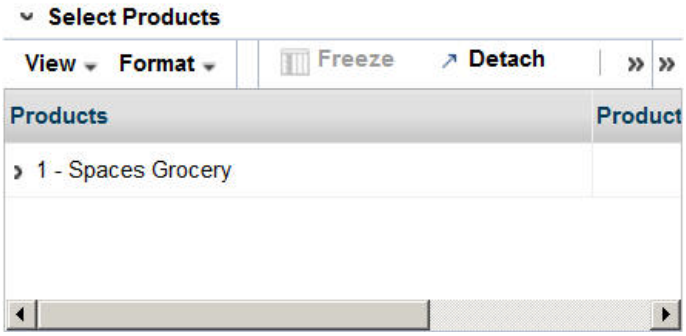


Figure 6–37 Select Products

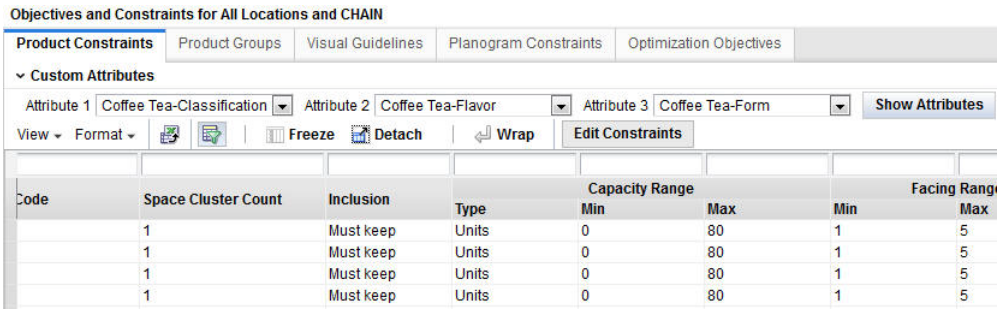


Objectives and Constraints Details

This section is used to specify optimization constraints that define the business rules that determine acceptable virtual planograms. It has five tabs; each tab displays detailed information that you can modify if you created the run.

Product Constraints The Product Constraints tab is shown in Figure 6–38. It includes Custom Attributes, which allow the user to select up to three attributes associated with the selected category of products. Select the attributes from the drop-down menus and click **Show Attributes** to add the attributes to the Product Constraints table.

Figure 6–38 Product Constraints



For each product, it defines the current product constraints, as listed in Table 6–32.

Table 6–32 Product Constraints

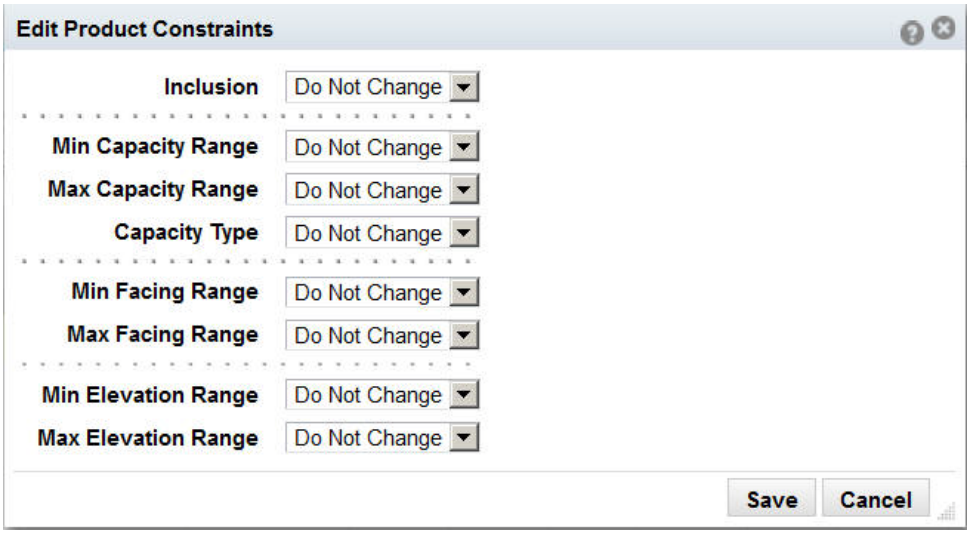
Column	Description
Department	Merchandise/product hierarchy node.
Class	Merchandise/product hierarchy node.
Subclass	Merchandise/product hierarchy node.
Product Name	Product identification.
Product Code	External product code.
Store or Space Cluster Count	Number of stores/space clusters that carry the product.
Inclusion	Supported inclusion values: Must keep (product must not be dropped), Can keep (product can be dropped), and Do not include (product must not be included).
Capacity Range	Type, minimum, and maximum value for capacity range.

Table 6–32 (Cont.) Product Constraints

Column	Description
Facing Range	Minimum and maximum value for facing range.
Elevation Range	Minimum and maximum value for elevation range.

Click the **Edit Constraints** button to access the Edit Product Constraints dialog box, shown in Figure 6–39. After selecting one or more rows, you can edit the constraints if you created the run and if the run has not been submitted.

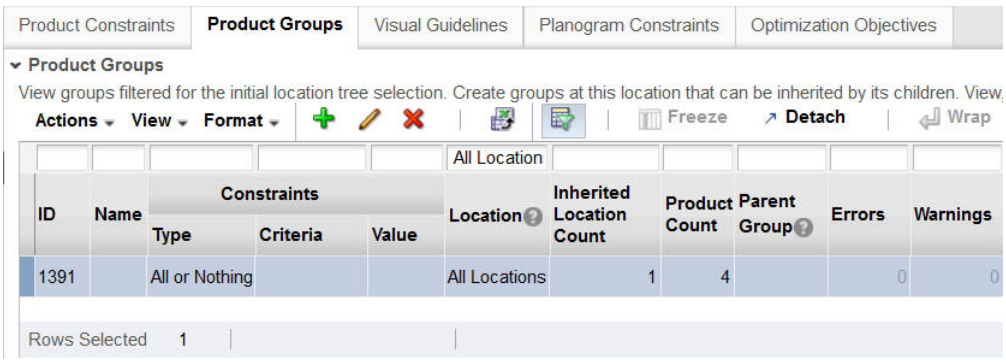
Figure 6–39 Edit Product Constraints



Use the drop-down lists, which contain the valid options for each constraint, to change the settings for any constraints, as appropriate. After you complete your edits, select **Save** from the Action menu. The system validates the changes you make and adjusts other settings as necessary. If you enter minimum and maximum values that are inconsistent, you will see an error message and must make changes before you can successfully save your edits.

Product Groups The Product Groups tab, shown in Figure 6–40, is used to define parent groups, which, when defined at a location node above the optimization location, can be inherited by the optimization locations that are the children of the parent location node.

Figure 6–40 Product Groups



You can add, edit, or delete product groups as well as add products to and delete products from product groups. Inherited groups can also be added or deleted.

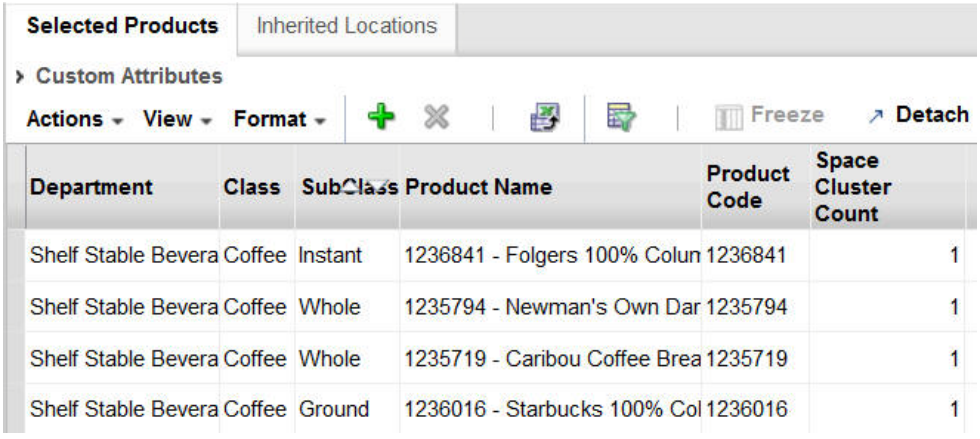
Table 6–33 Product Groups

Column	Description
ID	Unique numerical identification for product group.
Name	Name of product group.
Location	Store or space cluster for product group.
Constraint Type	Values are: All or Nothing, Match Facings, Same Shelf, and Choose From.
Constraint Criteria	If Type is Choose From, values are: At Least, At Most, and Exactly.
Constraint Value	If Type is Choose From, then a numerical value is required here.
Inherited Location Count	If the product group belong to either a Level 1 or Level 2 optimization location, the product group can be copied to or referenced by Level 3 optimization locations. This column shows the count of such optimization locations.
Product Count	The number of products in the product group.
Parent Group	If this product group is created by copying or referencing an existing product group of a higher level (1 or 2) then this column shows the ID of parent product group.
Errors	Number of errors for product group.
Warnings	Number of warnings for product group.

Selected Products

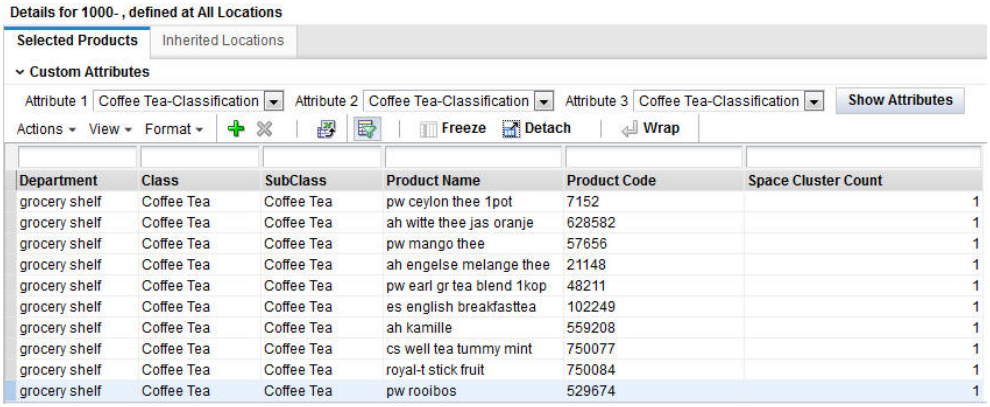
Use this section, shown in Figure 6–41 to add or remove products from a selected product group.

Figure 6–41 Selected Products



You can see and use the product attributes by accessing the Custom Attributes dialog box shown in Figure 6–42 when you are creating product groups. When you make a selection, the Product Groups table is populated.

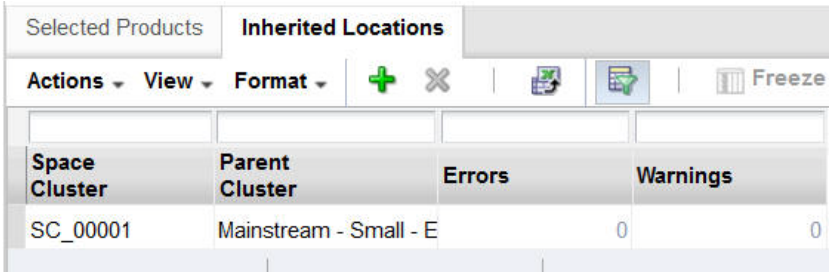
Figure 6–42 Custom Attributes



Inherited Locations

Use this section, shown in Figure 6–43, to add or remove inherited locations.

Figure 6–43 Inherited Locations



Visual Guidelines Visual Guidelines are used to provide attribute-based vertical and horizontal blocking for placing products on a shelf fixture planogram. You are allowed up to four levels for the blocking criteria. The visual guidelines are created at the optimization location level. Each optimization can have only one visual guideline. The sorting attributes are defined at the top level and are applied across all of the levels below the top level. You can add, delete, and edit the visual guidelines.

Blocking is defined as horizontal or vertical. Vertical is the default. You can specify primary and secondary blocking criteria. The secondary blocking is optional. If the primary blocking is vertical then the secondary blocking is horizontal; if the primary blocking is horizontal then the secondary blocking is vertical. Horizontal blocking supports pegboards.

Primary and secondary blocking can have up to two attributes each. Each attribute for a specific blocking strategy is defined in the same way, horizontal or vertical. A sequence number is assigned to each attribute value for each blocking level. The sequence numbers determine the order of the blocks. In ascending order, the blocks are positioned from top to bottom (vertical) and left to right (horizontal). When two are blocks are assigned the same sequence number, the blocks are combined (merged).

Example 1:

Primary blocking: Vertical blocking by brand.

Merge blocks: Brand A, Brand B, and Brand C

Additional primary blocking: Vertical blocking by type.

Secondary blocking: Horizontal blocking by size for all vertical blocks.

Example 2:

Primary blocking: Horizontal blocking by size.

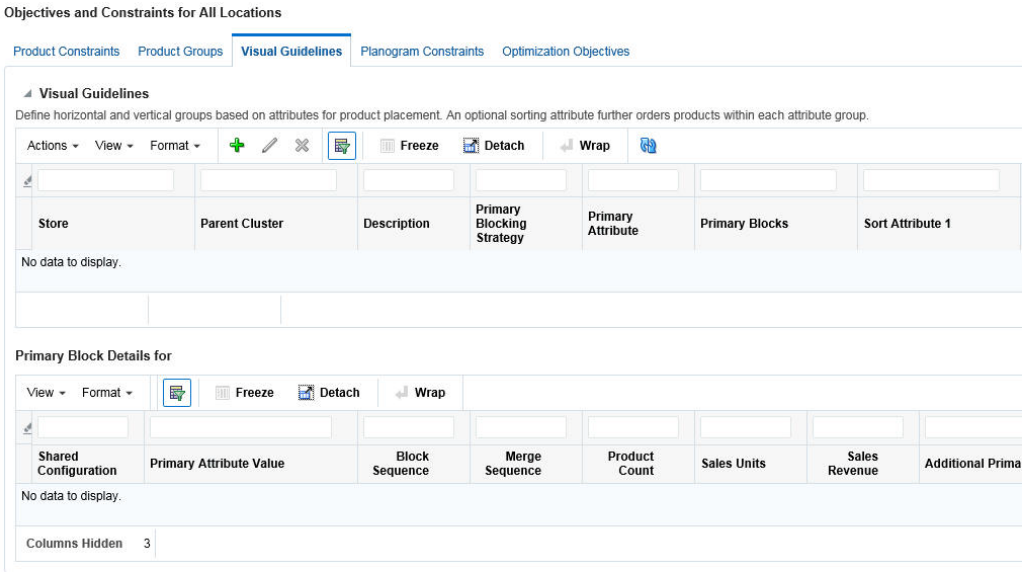
Additional primary blocking: Horizontal blocking by flavor.

Secondary blocking: Vertical blocking by brand for all vertical blocks.

Visual Guidelines Table

The Visual Guidelines table provides information about each visual guidelines row.

Figure 6–44 Visual Guidelines



The columns for this table are described in Table 6–34.

Table 6–34 Visual Guidelines

Column	Description
Space Cluster/Store	The name of the optimization location.
Parent Cluster	The name of the parent cluster for the optimization location.
Description	The name of the visual guideline.
Primary Blocking Strategy	Vertical or horizontal.
Primary Attribute	The product attribute used.
Primary Blocks	The first block used for the configuration
Sort Attribute 1	The attribute you select that is used to determine the product display order.
Sort Attribute 2	The attribute you select that is used to determine the product display order.
Sort Attribute 3	The attribute you select that is used to determine the product display order.

The Primary Block Details table provides additional information.

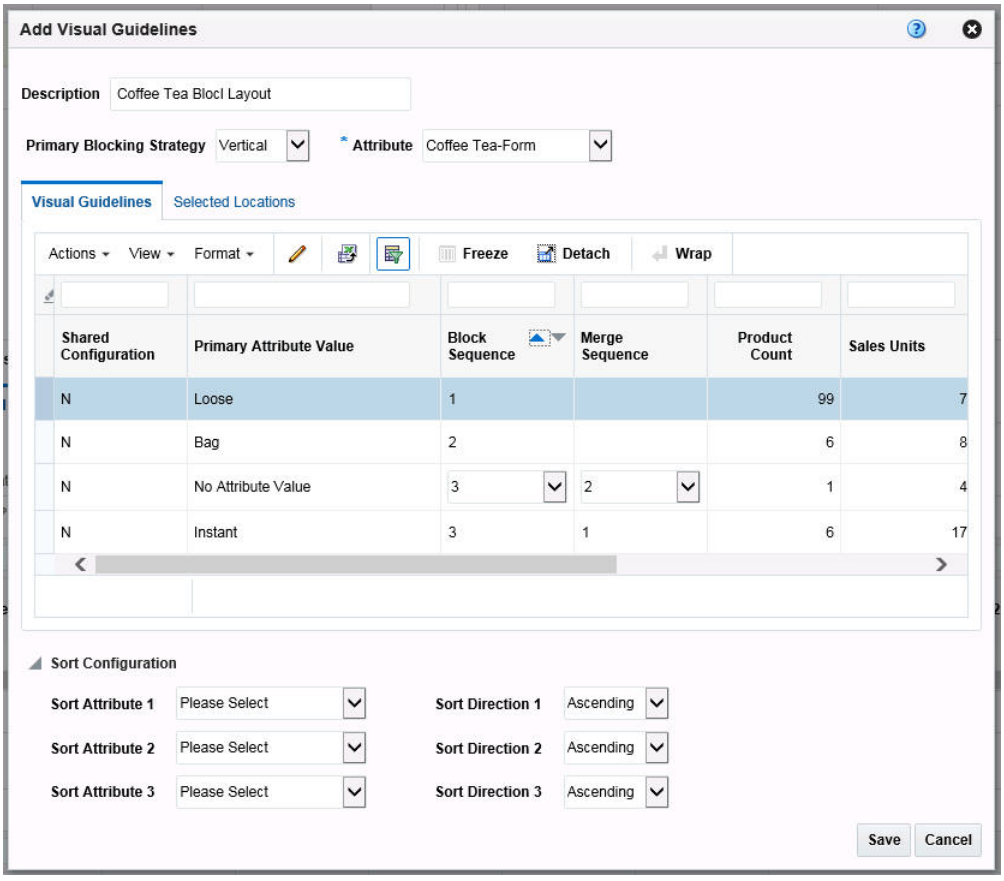
Table 6–35 Primary Block Details

Column	Description
Shared Configuration	Indicate whether or not two attributes share the same configuration.
Primary Attribute Value	Primary attribute.
Block Sequence	Defines the block and the sequence for that block.
Merge Sequence	Blocks with the same merge value are combined.
Product Count	The number of products in the block.
Sales Units	The average sales units for products in the block.
Sales Revenue	The average sales revenue for products in the block.
Additional Primary Attribute	Additional primary attribute.
Secondary Attribute	Secondary attribute.
Additional Secondary Attribute	Additional secondary attribute.

Adding Visual Guidelines

You can add visual guidelines through the Add Visual Guidelines dialog box. You access this by selecting the **Detail** radio button and clicking the **Add** icon.

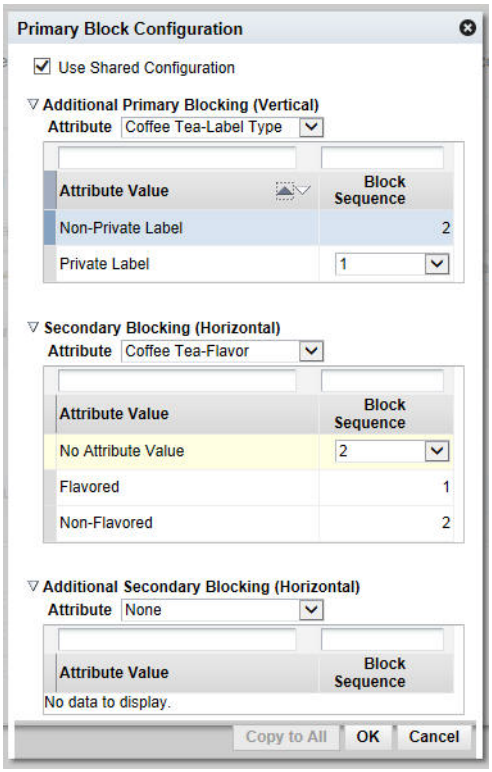
Figure 6–45 Add Visual Guidelines



To add visual guidelines, complete the following steps:

1. Enter the description of the visual guideline in the Description text box.
2. Select the primary blocking strategy, horizontal or vertical, from the Primary Blocking Strategy drop-down list.
3. Select the attribute to use for blocking from the Attribute drop-down list. This populates the table with the attribute values. At this point, the values for all sequences are assigned a default value of 1.
4. Define the attribute sequence number for each attribute by selecting a value from the drop-down list in the Block Sequence column. This defines the block and the sequence for the block. To define merged blocks, assign the same value to each attribute to be merged.
5. Optionally, define the merge sequence value by selecting a value from the drop-down list in the Merge Sequence column. This functionality is only enabled if you have assigned the same sequence number to two of the primary blocks.
6. Define up to three sort attributes and the sort direction, ascending or descending, for each.
7. Choose locations from Selected Locations.
8. Optionally, prior to saving the configuration, you can configure the primary block. To do this, select the attribute value in the table and click the **Pencil** icon. You see the Primary Block Configuration dialog box.

Figure 6–46 Primary Block Configuration



After the primary blocking strategy and attribute are defined, the primary block can be configured. When primary blocking is vertical, then horizontal blocking is

specified for each primary block. When primary blocking is horizontal, then vertical blocking is specified for each primary block.

9. When two attributes have the same sequence number, they share the same configuration. If you change one of the attribute's block sequence numbers to an unassigned number, it still maintains that configuration. If you change the attribute's block sequence number to an assigned number, then it takes the new configuration of the block it joins.

   The Use Shared Configuration check box functions is different ways, depending on the following:

   - When the primary blocking strategy is vertical, the check box is enabled. Check this box if you want all of the primary blocks to share the same secondary and additional blocking. If you do not check this box, then each primary block is configured separately.
   - When the primary blocking strategy is horizontal, the check box is disabled. When two or more attributes share the same sequence number, all the products that have these attributes are combined.

10. To configure Additional Primary Blocking, select the attribute from the drop-down list. The table populates with the attribute values. All sequence numbers default to a value of 1. If you select None for the attribute, then blocking at this level is disabled.

11. Define the attribute sequence for each attribute value by selecting the sequence number from the drop-down list in Block Sequence. If you assign the same sequence number to more than one attribute, then all the products with these attributes are combined.

12. To configure Secondary Blocking, select the attribute from the drop-down list. The table populates with the attribute values. All sequence numbers default to a value of 1. If you select None for the attribute, then blocking at this level is disabled.

13. Define the attribute sequence for each attribute value by selecting the sequence number from the drop-down list in Block Sequence. If you assign the same sequence number to more than one attribute, then all the products with these attributes are combined.

14. You can only configure Additional Secondary Attributes if you have configured Secondary Blocking. To configure Additional Secondary Attributes, select the attribute from the drop-down list. The table populates with the attribute values. All sequence numbers default to a value of 1. If you select None for the attribute, then blocking at this level is disabled.

15. Define the attribute sequence for each attribute value by selecting the sequence number from the drop-down list in Block Sequence. If you assign the same sequence number to more than one attribute, then all the products with these attributes are combined.

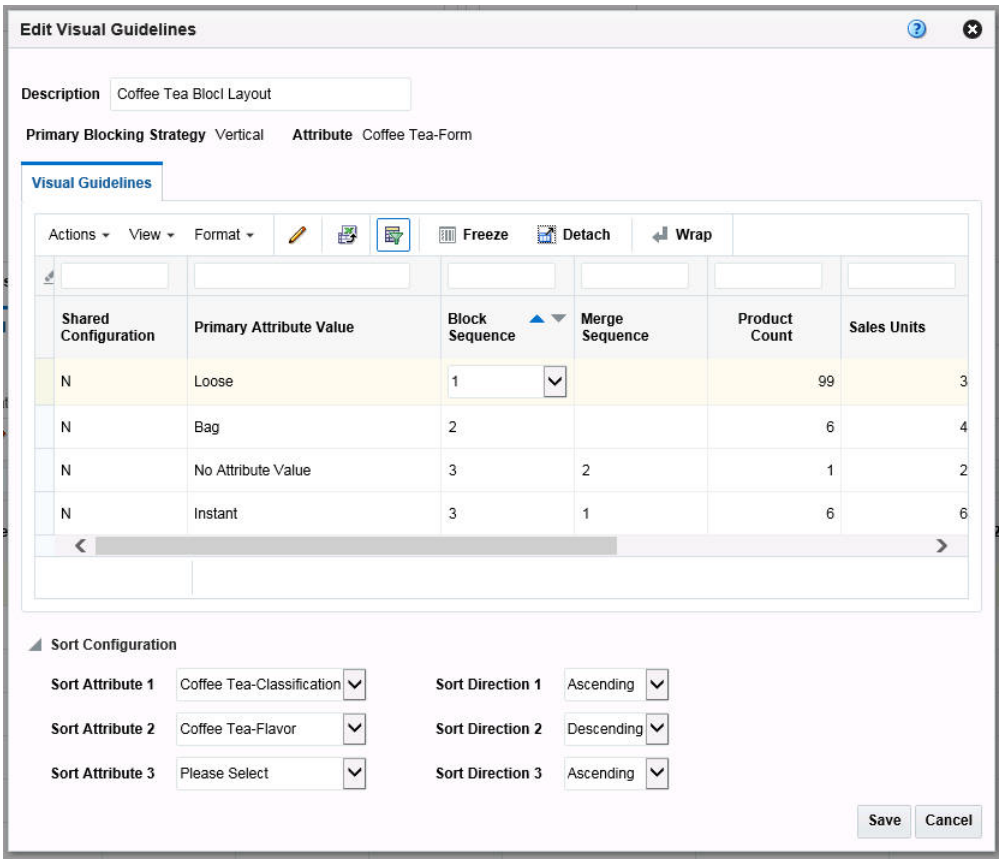
16. Click **OK**. The Add Visual Guidelines table is updated with the selected attributes.

17. Click **Save** to save the configuration. This overwrites any existing visual guidelines configuration for the selected location.

Editing Visual Guidelines

You can modify the characteristics of a visual guideline through the Edit Visual Guidelines dialog box. Select a row to indicate which guideline you want to edit and click the **Pencil** icon. You see the Edit Visual Guidelines dialog box.

Figure 6–47 Edit Visual Guidelines

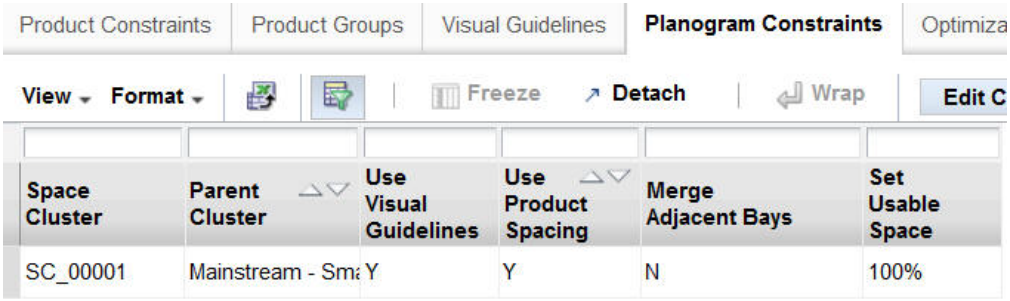


You can modify any values except the Primary Blocking Strategy and the Primary Attributes following same approach you use to create those values. Once you have completed making changes, click **Save** to save your changes and update the Visual Guidelines table. You can click the **Cancel** button to cancel the edits.

Planogram Constraints The Planogram Constraints tab, shown in Figure 6–48, displays a list of optimization locations along with values for four constraints: Use Visual Guidelines, Use Product Spacing, Merge Adjacent Bays, and Set Usable Space.

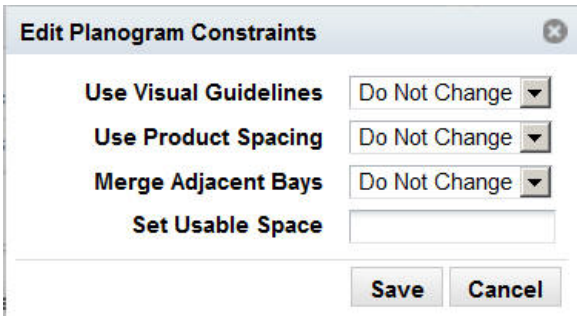
Service levels can be defined at the planogram level using the Planogram Constraints screen.

Figure 6–48 Planogram Constraints



You can edit these constraints by selecting the row or rows whose settings you want to change and then click **Edit Constraints**. You see the dialog box shown in Figure 6–49.

Figure 6–49 Edit Planogram Constraints



Use the drop-down menus to select a new setting. For Set Usable Space, you must enter a value between 1% and 100%. Select Do Not Change if you do not want to change a particular constraint. Click **Save** to save your changes.

Table 6–36 Planogram Constraints

Column	Description
Space Cluster/Store	The name of the optimization location.
Parent Cluster	The name of the parent cluster for the optimization location.
Use Visual Guidelines	Indicates whether or not to use visual guidelines. Values are Yes and No.
Use Product Spacing	Indicates whether or not to use product spacing. Values are Yes and No.
Merge Adjacent Bays	Indicates whether or not to merge adjacent bays. Values are Yes and No. The optimization process provides the option of merging adjacent bays.
Set Usable Space	A value between 1 and 100 that indicates the usable space. available to the optimization process. A values less than 100 forces the optimization process to leave some fixture space unused for aesthetic or other reasons.

Optimization Objectives The Optimization Objectives table, shown in Figure 6–50, includes:

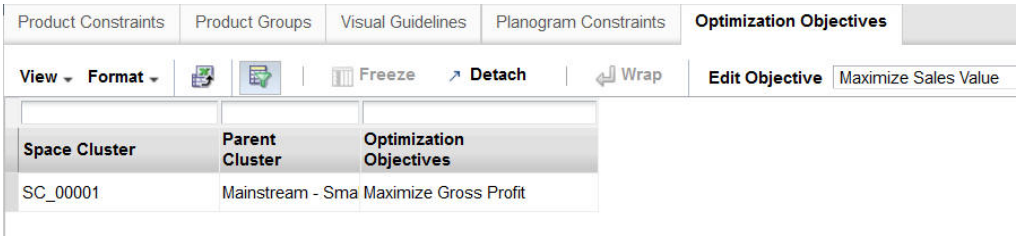
- Maximize Sales Units. This tells the optimization to fill the planogram with items that will result in best possible sales units for selected items/facings.
- Maximize Sales Revenue. For an item, sales revenue is calculated as price times the sales units. This tells the optimization to fill the planogram with items that will result in the best possible sales revenue for selected items/facings.
- Maximize Gross Profit. For an item, gross profit is obtained by multiplying the difference between price and cost, and sales units. This tells the optimization to fill the planogram with items that will result in the best possible gross profit for selected items/facings.
- Maximize Sales Revenue/On Hand Units. For each item, the ratio is calculated and added together. This objective tells the optimization to fill the planogram with items that result in the best possible revenue for the inventory units carried for each item.
- Maximize Gross Profit Return on Investment. This objective tells the optimization to fill the planogram with items that result in the best possible total profit, but that at the same time minimize the total inventory cost at the planogram-level. As one

can imagine, carrying too much inventory will result in higher revenues or profits but at a higher cost of carrying more inventory. This metric lets the user strike a balance between these two metrics. There are two versions supported: the generic version in which the user is unable to provide the store-level inventory. To use the generic version GENERIC_GMROI_FLG is set to 'Y' in RSE_CONFIG.

- Maximize Sales Units (Weighted). This objective is similar to 'Maximize Sales Units' except that here each item's contribution to objective is weighted by the IPI values provided by CMPO.
- Maximize Sales Revenue (Weighted). This objective is similar to 'Maximize Sales Revenue' except that here each item's contribution to objective is weighted by the IPI values provided by CMPO.
- Maximize Gross Profit (Weighted). This objective is similar to 'Maximize Gross Profit' except that here each item's contribution to objective is weighted by the IPI values provided by CMPO.
- Maximize Sales Revenue/On Hand Units (Weighted). This objective is similar to 'Maximize Sales Revenue/On Hand Units' except that here each item's contribution to objective is weighted by the IPI values provided by CMPO.
- Maximize Gross Profit Return on Investment (Weighted). This objective is similar to Maximize Gross Profit Return on Investment except that here each item's contribution to the objective (in the numerator and denominator) is weighted by the IPI values provided by CMPO.

Weighted objectives are weighed using the IPI values provided by CMPO.

Figure 6–50 Optimization Objectives



Select one or more rows in order to modify specific objectives using the **Apply** button.

Table 6–37 Optimization Objectives

Field	Description
Space Cluster/Store	The name of the optimization location.
Parent Cluster	The name of the parent cluster of the optimization location.
Optimization Objectives	The name of the current optimization objective.

Contextual Information

Each of the five tabs within the Objectives and Constraints stage have associated contextual information. The information is displayed based on the selections you make in the location hierarchy table.

The Product Constraints information includes:

- Constraint Values, which displays a list of the constraints and associated metrics
- Constraint Charts, which provide a summary for each constraint

- Product Details

Product Groups information includes:

- Product Group Summary, which displays counts for the group types
- Product Lookup, which shows, for a given product, the associated group ID, group name, and group type

Visual Guidelines information includes:

- Visual Guidelines Summary, which includes an Attributes histogram for each attribute, and Attributes Group histogram
- Attributes Lookup, which provides details about product counts for individual attributes
- Visual Guidelines Detail, which analyzes attribute values and product counts by sequence number

Planogram Constraints information includes:

- Constraint Charts, which provide summaries for each constraint
- Planogram Space, which can help you understand the relationship between the available space and the set of products.

Optimization Objectives information includes:

- Optimization Objectives
- Product Details, which includes details for IPI values, price, and margin percentage

Results and Analysis

After the optimization run is complete, the results are displayed in the Results and Analysis stage. You can see summary performance data and view planograms for the optimization runs. In this stage, you can make a limited number of overrides to the table view or the planogram view. Substantive changes require that you re-configure settings and complete another optimization run.

The results are displayed in two tables. The upper table displays the total optimized performance. The data is summed for all the stores in a space cluster. You can access the virtual planogram by clicking the **View Planogram** icon. The lower table displays detailed results with product-level data. The data you see is the average for store numbers, not totals for space clusters.

You can drop products from the lower table for one or more locations; these dropped products are listed in a table.

Alerts are also provided, along with supporting metrics.

Actions To Take

From the Results and Analysis stage, you can:

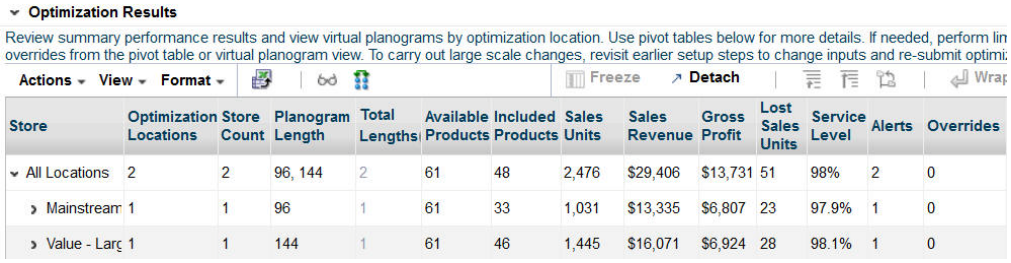
- Revisit earlier stages in order to make changes to various settings and then re-submit the optimization. You can also validate a run in order to look for constraints that produce errors or warnings. See Validation for more details.
- Copy and save the run using the **Save As** functionality in the Action menu. Use the copy to make changes and then compare the results.
- Approve the run and make it available for assortment analysis.

- Leave the run in the optimization list without taking any actions. You cannot delete the run; you must be in Optimization List in order to delete a run.

Optimization Results

The Optimization Results table is shown in Figure 6–51.

Figure 6–51 Optimization Results



You can toggle between a tree view and a table view to examine the results. The tree view displays the results hierarchically. The results are either a sum of lower level values, an average of lower level values, or a distinct count.

Table 6–38 Optimization Results

Field	Description
Space Cluster/Store	The name of the optimization location: Store or Space Cluster.
Optimization Locations	The number of optimization locations.
Store Count	The number of distinct stores for a given level.
Planogram Length	A comma-separated list of selected planogram lengths
Total Lengths	The total count of planogram lengths for which the optimization was run. If more than one total is displayed, you can click the number to access a pop-up window. You can modify selected planogram lengths for the specified optimization location.
Available Products	A count of the available products for a given space cluster level.
Included Products	A count of the included products in an assortment for a given space cluster level.
Sales Units	The value for sales units in the optimization results. This value is aggregated at higher levels by summing the values for the lower levels.
Sales Value	The value for sales. This value is aggregated at higher levels by summing the values for the lower levels.
Margin Value	The value for margin. This value is aggregated at higher levels by summing the values for the lower levels.
Lost Sales Units	The value for lost sales units. This value is aggregated at higher levels by summing the values for the lower levels.
Service Level	The value for the service level. The value is aggregated at higher levels by averaging the values for the lower levels.
Alerts	The number of alerts. This value is aggregated at higher levels by summing the values for the lower levels.
Overrides	The number of user overrides. This value is aggregated at higher levels by summing the values for the lower levels.

Included Products

The Included Products table lists metrics at various levels of aggregation. You can select specific products from the table and click the **Delete** button in order to remove them from the results. You see a list of the products you want to delete and are prompted to confirm the action.

Dropped Products

The Dropped Products table, shown in Figure 6–52, lists all the products that have been dropped per optimization process rules including but not limited to user override action or if the Inclusion Product constraint is set to Do Not Include.

Figure 6–52 Dropped Products

		Demand	Demand Transferred
› Value - Large - A (Northwest)	› 1236841 - 1236841 - Folgers 100% Columbian Non-Flavored Regular - Caffeinated 12 oz Jar	109	0
	› 1234759 - 1234759 - Folgers Medium Roast Non-Flavored Regular - Caffeinated 12 oz Can	101	0

Alerts

The Alerts table, in Figure 6–53, shows the metrics provided for alerts.

Figure 6–53 Alerts

		Used Space Percent	Dropped Products Percent	Dropped Products Count	Service Level Value	Service Level Units
› Mainstream - Small - E (Northwest)	Dropped Percent, Dropped Count	96.7%	45.9%	28	1	1

The Used Space Percent can help you see if there is too much open space.

The Dropped Products Percent and Dropped Products Count can help you see if too many products were eliminated from an assortment.

The Service Level Value and Service Level Units can help you understand stockout levels.

You can configure the following parameters to control the thresholds for the alerts.

Table 6–39 Alerts Parameters

Parameter Name	Description
ALRT_LESS_THAN_PCT_USED_SPACE	An alert will be triggered if the run optimization results use less space than the value specified by this global parameter.

Table 6–39 (Cont.) Alerts Parameters

Parameter Name	Description
ALRT_LESS_THAN_SERVICE_LEVEL_AMT	An alert will be triggered if the run optimization results have a sales service level lower than he value specified by this parameter.
ALRT_LESS_THAN_SERVICE_LEVEL_QTY	An alert will be triggered if the run optimization results have a quantity service level lower than he value specified by this parameter.
ALRT_MORE_THAN_CNT_PRODUCT_DROPPED	An alert will be triggered if the run optimization results dropped more products than he value specified by this parameter.
ALRT_MORE_THAN_PCT_PRODUCT_DROPPED	An alert will be triggered if the run optimization results dropped a percentage of product higher than he value specified by this parameter.

Contextual Information

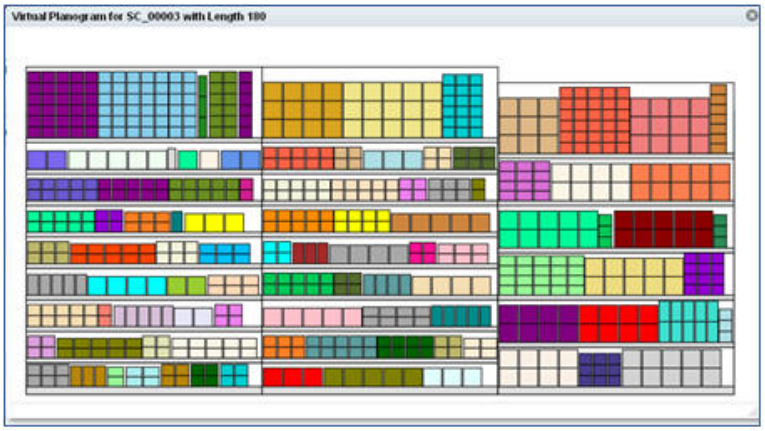
The Results and Analysis stage provides summary information about sales units, sales value, service level, and included products.

- Results Summary
- Inclusion Summary

Planograms

When you click the **View Results** icon, you see a pop-up window displaying a planogram your run. Depending on your permissions, you can edit the planogram.

Figure 6–54 Planogram



Selection View

Left-click a facing to highlight it.

Fixture Details

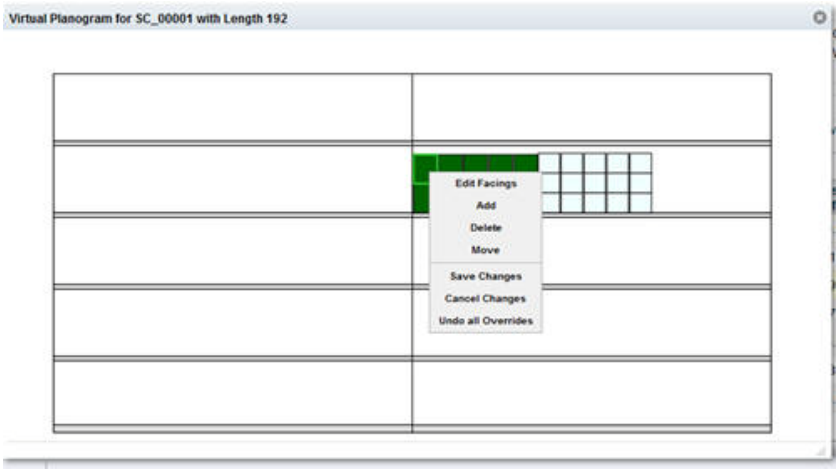
When you mouse over a fixture, the following information is displayed: fixture type, length, depth, and elevation.

Product Details

When you mouse over a product, the following information is displayed: product name, product code, facing count, facing ID, display style, main orientation, dimensions, stacking option, and elevation. When applicable, visual guideline information such as definition, attribute value, and sorting attribute data is shown.

Actions When you right-click a facing to select it, you have access to the following actions: editing facings, add, delete, move, save changes, cancel changes, and undo all changes.

Figure 6–55 Virtual Planogram with Action Menu



Edit Facings When you select this action, you can change the facing quantity.

Delete Use this action to delete a product. Deleted products can be added back using the Add Products action.

Add When you select this action, you see a list of product codes for dropped products from which to choose. The product you select is restored to its original position and with the same number of facings as when it was dropped.

Move If you move a product on a shelf, it moves to the 0 location on the shelf. For a facing move, the product is moved to the first available shelf location after the last product facing.

Save The save action commits changes to the database.

Cancel The cancel action rolls back any changes.

Undo Overrides This action restores the planogram to the original results. If you close the application without saving changes, the next time you log in, you are prompted to save the changes.

Historical Planograms Click the **Compare to History** button to view a historical planogram. Such planograms cannot be modified.

The historical planogram contains colored rectangles representing products. You can change the color mapping in the historical planogram and select images of products (if available). To do this, click the **Color** button. You see the Color Mapping pop-up, from which you can choose one of the following color mappings:

- Product ID (the default)
- Product Image

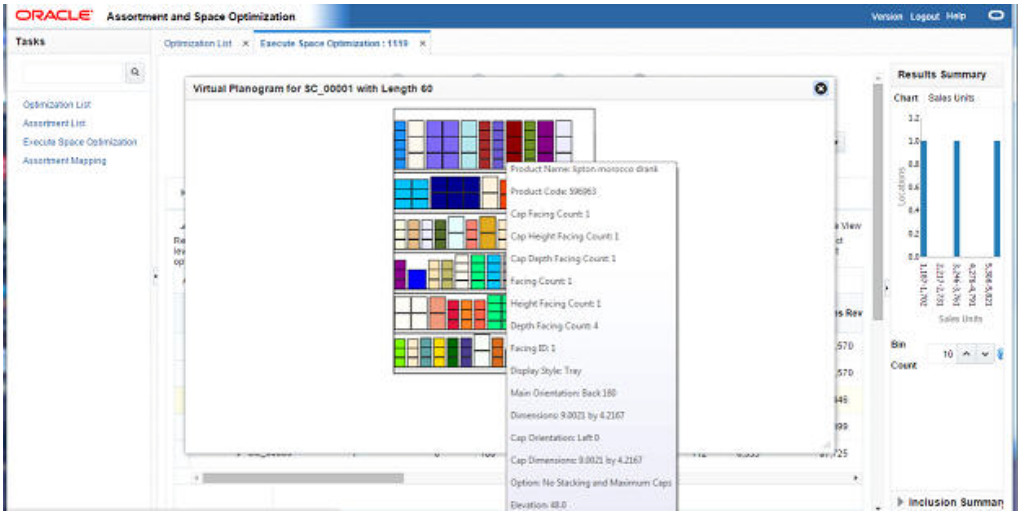
Capping Capping is an alternative way of defining the top layer(s) of a product on a shelf after the bottom layer(s) has been placed. Capping layers may differ from the corresponding bottom layers, in both numbers of rows and columns as well as the individual dimensions.

The VPOG pop-up displays the applicable capping layers. It also provides capping information in the hover pop-up. This only occurs when the style indicates capping, and it is applicable to the whole product display style, independent of the hover location. The following capping information is shown:

- Capping Facing Count. The horizontal facing count for the capping layer.
- Capping Height Facing Count. The depth facing count for the capping layer.
- Capping Depth Facing Count. The depth facing count for the capping layer.
- Capping Dimensions. Product dimensions for the corresponding capping layer.
- Capping Elevation. Elevation of the corresponding capping layer.
- Capping Orientation. Orientation of the corresponding capping layer.

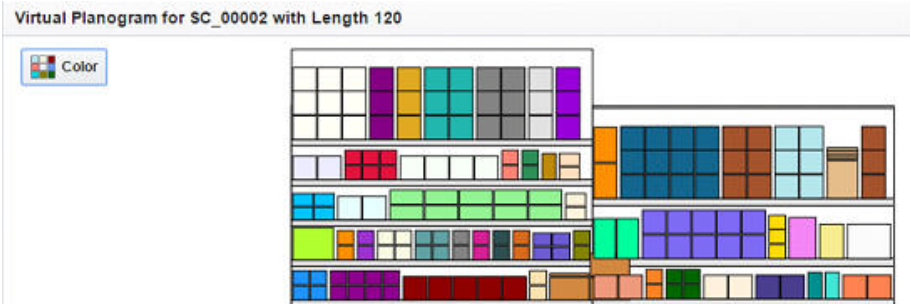
Figure 6–56 provides an example of the VPOG hover pop-up.

Figure 6–56 VPOG Hover Pop-Up



Changing Color Mapping You can change the colors of the displayed boxes or display product images in the virtual planogram. To do this, **click** the **Color** button.

Figure 6–57 Color Button



You see the Color Mapping pop-up, from which you can choose the color mapping you want:

- Product ID (the default)
- Primary Attribute Value
- All Visual Guidelines
- Product Image

Figure 6–58 Color Mapping Pop-Up



Validation

You can validate your data and constraint settings within each stage of Execute Space Optimization by using the **Validate** option in the Action menu. Validation lets you identify problems with potential data or constraints prior to initiating a complete optimization run. The Validate functionality is enabled after you create a run, provide it with a name, and populate it with a POG set and a group of stores or clusters. The Validate functionality is disabled for a run once it has a status of Finalized or Approved.

The validation process provides you with a quick way to determine if an optimization problem has any obvious incorrect data or parameters before you submit a run and request an optimization result. The process is not designed to find all possible errors or warnings, as some infeasible scenarios are only identifiable by the optimization process itself. However, the validation does provide you with the opportunity to find potential issues, and to then navigate back to the Fixture and Product Data or Objective and Constraints stage in order to address the setup issues. This can help you avoid submitting a run for optimization and then discovering that changes are required.

The validation process checks on issues with price, cost, demand and replenishment data. Specifically it looks for invalid values, missing data and generates an error or warning, depending on whether the product is mandatory or optional. Except for replenishment data which can also be modified in the UI, other data issues will require the user to correct and re-load the data interface files.

Validation identifies two types of conflict:

- Errors. Identifies data or constraints that prevent an optimization from completing successfully. You must correct all errors before you execute the run.
- Warnings. Identifies data or constraints that you may want to review and possibly change. Warnings do not necessarily cause the optimization to fail but can have unwanted consequences in terms of the results. Thus, it is always a good practice to review the warnings before clicking Submit.

Figure 6–59 Summary with Error



Figure 6–60 Summary with Warning

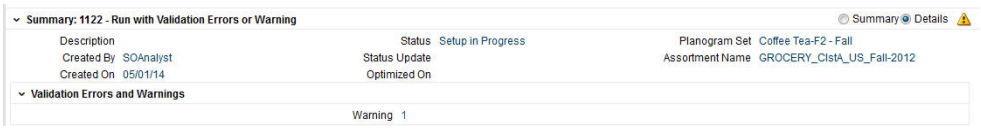
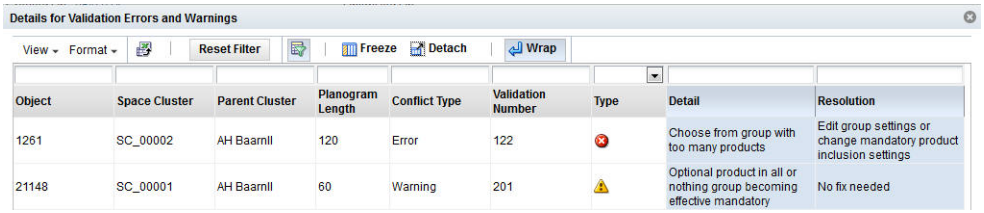


Figure 6–61 Error Details



Details About Errors and Warnings

If the validation you request finds errors or warnings, you will see an error icon in the Summary area of the stage you are in. Some validation checks only apply to shelf fixtures while others apply to both shelves and pegboards. To view the details, expand the Validation Errors and Warning section below the Summary section. Click each link in order to display the Details for Validation Errors and Warnings dialog box, which contains information explaining the error or warning that can help you address the problem.

Warnings

Warnings are informational and indicate that optional items will be handled according to the constraints you set. Here is a list of possible warnings and suggested optional actions to address the warnings.

Table 6–40 Warnings and Suggested Actions

Alert #	Warning Description	Suggested Optional Action
201	Optional product in all or nothing group becoming effective mandatory.	No fix needed.
202	Optional choose from product becoming effective mandatory.	No fix needed.

Table 6–40 (Cont.) Warnings and Suggested Actions

Alert #	Warning Description	Suggested Optional Action
209	Optional product becoming effective drop.	No fix needed.
211	Optional product too wide.	Reduce the minimum number of facings or change the elevation or orientation.
212	Optional product too tall.	Change the elevation or orientation.
213	Optional product lacks shelf.	Change the elevation.
232	No common set of shelves.	Edit the group settings or elevation ranges.
236	No common facing value	Edit the group settings or facing ranges.
1000	Product price is 0 or negative for (all) the store(s)	Optional: Re-load so_assort_proloc_ pricecost_stg.txt file with valid prices otherwise the product will be dropped
1001	Product price is 0 or negative for some stores	Optional: Re-load so_assort_proloc_ pricecost_stg.txt file with valid prices otherwise the product will be incorrectly valued/dropped as affected stores are discarded
1002	Product cost is 0 or negative for (all) the store(s)	Optional: Re-load so_assort_proloc_ pricecost_stg.txt file with valid costs otherwise the product will be dropped
1003	Product cost is 0 or negative for some stores	Optional: Re-load so_assort_proloc_ pricecost_stg.txt file with valid costs otherwise the product will be incorrectly valued as affected stores are discarded
1004	Product price and cost are equal for (all) the store(s)	Optional: Re-load so_assort_proloc_ pricecost_stg.txt file with valid prices or costs otherwise the product will be dropped
1005	Product price and cost are equal for some stores	Optional: Re-load so_assort_proloc_ pricecost_stg.txt file with valid prices or costs otherwise the product will be incorrectly valued/dropped as affected stores are discarded
1006	Product cost is greater than the price for (all) the store(s)	Optional: Re-load so_assort_proloc_ pricecost_stg.txt file with valid prices or costs otherwise the product will be dropped
1007	Product cost is greater than the price for some stores	Optional: Re-load so_assort_proloc_ pricecost_stg.txt file with valid prices or costs otherwise the product will be incorrectly valued/dropped as affected stores are discarded
1008	Product demand is 0 or negative for (all) the store(s)	Optional: Re-load so_assort_proloc_fcst_ stg.txt file with valid demand values otherwise the product will be dropped
1009	Product demand is 0 or negative for some stores	Optional: Re-load so_assort_proloc_fcst_ stg.txt file with valid demand values otherwise the product will be incorrectly valued/dropped as affected stores are discarded

Table 6–40 (Cont.) Warnings and Suggested Actions

Alert #	Warning Description	Suggested Optional Action
1010	Missing price, cost, and demand data for (all) the store(s)	Optional: Re-load the so_assort_proloc_pricecost.txt and so_assort_proloc_fcst.txt files with valid data otherwise the product will be dropped
1011	Missing price, cost, and demand data for some stores	Optional: Re-load the so_assort_proloc_pricecost.txt and so_assort_proloc_fcst.txt files with valid data otherwise the product will be incorrectly valued/dropped as affected stores are discarded
1012	IPI value is 0 or negative for (all) the store(s)	Optional: Re-load so_assort_product_strcltr_stg.txt with valid IPI values otherwise the product will be dropped
1013	IPI value is 0 or negative for some stores	Optional: Re-load so_assort_product_strcltr_stg.txt with valid IPI values otherwise the product will be incorrectly valued/dropped as affected stores are discarded
1014	Casepack value is 0 or negative for (all) the store(s)	Optional: Change casepack data through UI or by re-loading so_prod_loc_repl_param_stg.txt file with valid values otherwise product will be dropped
1015	Casepack value is 0 or negative for some stores	Optional: Change casepack data through UI or by re-loading so_prod_loc_repl_param_stg.txt file with valid values otherwise product will be dropped
1016	Product inventory cost is 0 or negative for (all) the store(s)	Optional: Re-load so_assort_proloc_pricecost_stg file with valid inventory cost otherwise the product may be incorrectly valued/dropped
1017	Product inventory cost is 0 or negative for some stores	Optional: Re-load so_assort_proloc_pricecost_stg file with valid inventory cost otherwise the product may be incorrectly valued/dropped as affected stores are discarded
1018	Product store inventory is 0 or negative for (all) the store(s)	Optional: Re-load so_assort_proloc_pricecost_stg file with valid inventory quantity otherwise the product may be incorrectly valued/dropped
1019	Product store inventory is 0 or negative for some stores	Optional: Re-load so_assort_proloc_pricecost_stg file with valid inventory quantity otherwise the product may be incorrectly valued/dropped as affected stores are discarded
1100	Mandatory product price is 0 or negative for (all) the store(s)	Re-load so_assort_proloc_pricecost_stg.txt file with valid prices otherwise the product will be dropped or make the product optional
1101	Mandatory product price is 0 or negative for some stores	Optional: Re-load so_assort_proloc_pricecost_stg.txt file with valid prices otherwise the product will be incorrectly valued/dropped as affected stores are discarded

Table 6–40 (Cont.) Warnings and Suggested Actions

Alert #	Warning Description	Suggested Optional Action
1102	Mandatory product cost is 0 or negative for (all) the store(s)	Re-load so_assort_proloc_pricecost_ stg.txt file with valid costs otherwise the product will be dropped or make the product optional
1103	Mandatory product cost is 0 or negative for some stores	Optional: Re-load so_assort_proloc_ pricecost_stg.txt file with valid costs otherwise the product will be incorrectly valued as affected stores are discarded
1104	Mandatory product price and cost are equal for (all) the store(s)	Re-load so_assort_proloc_pricecost_ stg.txt file with valid prices or costs otherwise the product will be dropped or make the product optional
1105	Mandatory product price and cost are equal for some stores	Optional: Re-load so_assort_proloc_ pricecost_stg.txt file with valid prices or costs otherwise the product will be incorrectly valued/dropped as affected stores are discarded
1106	Mandatory product cost is greater than the price for (all) the store(s)	Re-load so_assort_proloc_pricecost_ stg.txt file with valid prices or costs otherwise the product will be dropped or make the product optional
1107	Mandatory product cost is greater than the price for some stores	Optional: Re-load so_assort_proloc_ pricecost_stg.txt file with valid prices or costs otherwise the product will be incorrectly valued/dropped as affected stores are discarded
1108	Mandatory product demand is 0 or negative for (all) the store(s)	Re-load so_assort_proloc_fcst_stg.txt file with valid demand values otherwise the product will be dropped or make the product optional
1109	Mandatory product demand is 0 or negative for some stores	Optional: Re-load so_assort_proloc_fcst_ stg.txt file with valid demand values otherwise the product will be incorrectly valued/dropped as affected stores are discarded
1110	Mandatory product missing price, cost, and demand data for (all) the store(s)	Re-load the so_assort_proloc_pricecost.txt and so_assort_proloc_fcst.txt files with valid data otherwise the product will be dropped or make the product optional
1111	Mandatory product missing price, cost, and demand data for some stores	Optional: Re-load the so_assort_proloc_ pricecost.txt and so_assort_proloc_fcst.txt files with valid data otherwise the product will be incorrectly valued/dropped as affected stores are discarded
1112	Mandatory product's IPI value is 0 or negative for (all) the store(s)	Re-load so_assort_product_strcltr_stg.txt with valid IPI values otherwise the product will be dropped or make the product optional
1113	Mandatory product's IPI value is 0 or negative for some stores	Optional: Re-load so_assort_product_ strcltr_stg.txt with valid IPI values otherwise the product will be incorrectly valued/dropped as affected stores are discarded

Table 6–40 (Cont.) Warnings and Suggested Actions

Alert #	Warning Description	Suggested Optional Action
1114	Mandatory product's casepack value is 0 or negative for (all) the store(s)	Change casepack data through UI or by re-loading so_prod_loc_repl_param_ stg.txt file with valid values otherwise product will be dropped or make the product optional
1115	Mandatory product's casepack value is 0 or negative for some stores	Optional: Change casepack data through UI or by re-loading so_prod_loc_repl_ param_stg.txt file with valid values otherwise product will be incorrectly valued/dropped as affected stores are discarded
1116	Mandatory product inventory cost is 0 or negative for (all) the store(s)	Optional: Re-load so_assort_proloc_ pricecost_stg file with valid inventory cost otherwise the product may be incorrectly valued
1117	Mandatory product inventory cost is 0 or negative for some stores	Optional: Re-load so_assort_proloc_ pricecost_stg file with valid inventory cost otherwise the product will be incorrectly valued as affected stores are discarded
1118	Mandatory product store inventory is 0 or negative for (all) the store(s)	Optional: Re-load so_assort_proloc_ pricecost_stg file with valid inventory quantity otherwise the product may be incorrectly valued
1119	Mandatory product store inventory is 0 or negative for some stores	Optional: Re-load so_assort_proloc_ pricecost_stg file with valid inventory quantity otherwise the product will be incorrectly valued as affected stores are discarded

Correcting Errors

Once you have made corrections to the errors, you must run Validate again in order to make sure that all the errors have been corrected. Here is a list of possible errors and recommended corrections.

Table 6–41 Errors and Recommended Solutions

Alert #	Error Description	Recommended Solution
101	Mandatory product too wide.	Reduce minimum number of facings, change product orientation, or make the product optional.
102	Mandatory product too tall.	Change the product orientation or make the product optional.
103	Mandatory product lacks shelf.	Change the elevation setting or make the product optional.
106	Mandatory products combined width too wide for shelf.	Reduce the minimum facings or make some of the products optional.
107	Mandatory products combined area too large for pegboard.	Reduce the minimum facings or make some of the products optional.
121	Choose from group lacking products.	Edit the group settings or add products.
122	Choose from group with too many products.	Edit the group settings or change the mandatory product inclusion settings.

Table 6–41 (Cont.) Errors and Recommended Solutions

Alert #	Error Description	Recommended Solution
127	Same shelf group lacks shelf.	Edit group settings or change the mandatory product elevation settings.
128	Same shelf group is too wide.	Edit the group settings or change the mandatory product elevation or inclusion settings.
129	Same shelf group is too tall.	Edit the group settings or change the mandatory product elevation or inclusion settings.
132	Match facing group without match.	Edit the group settings or change the mandatory product inclusion settings.

Once all errors are corrected, you can submit the optimization run.

Assortment Mapping

The Assortment Mapping tab is used to view and edit the results of the automated mapping between assortments and planograms that the system performs.

Assortment Mapping Interface

The Assortment Mapping interface is divided into two sections. The Select Assortment table lists all assortments with a status of either Ready for SO or POG Mapping Needs Review. When you first access this screen, it is pre-filtered to show only POG Mapping Needs Review assortments, with the QBE field for the Status column already filled in. To see the Ready for SO assortments, you must clear this field and click **Enter** to trigger the update.

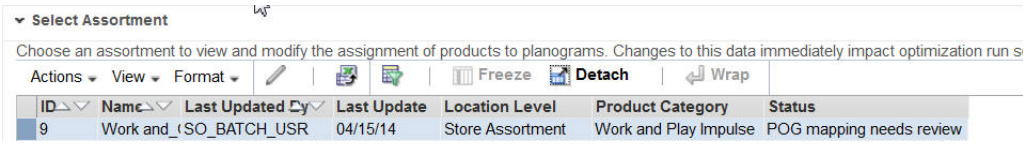
The Assortment to Planogram Mapping table displays the planogram mappings for each combination of location and product in a given assortment. You must have the appropriate permissions to edit the assortments.

Use the **Save** Button to save and validate any edits you make. This updates the mappings and refreshes the mapping table display. You must save your changes for one assortment before editing another assortment.

Select Assortment

The Select Assortment table, shown in Figure 6–62, displays a list of the assortments that require mapping. These assortments have a status of either Ready for SO or POG Mapping Needs Review. Use this list to select an assortment to edit.

Figure 6–62 Select Assortment



When you select an assortment to edit and click the **Edit** button, you see the Edit Status dialog box. You can toggle the status of the assortment between Ready for SO and POG Mapping Needs Review by using the radio buttons. These radio buttons are disabled if errors exist in the assortment mapping table.

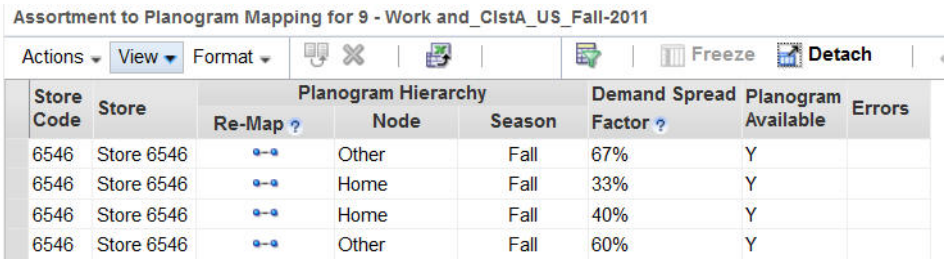
Table 6–42 Select Assortment

Field	Description
ID	The ID of the assortment.
Name	The name of the assortment.
Last Updated By	The user ID of the person who last updated the assortment.
Last Update	The date when the assortment was last updated.
Location Level	The location level for the assortment (Store Cluster or Cluster Assortment).
Product Category	The product category for the assortment.
Status	The status of the assortment.

Assortment to Planogram Mapping

Once you select an assortment from the Select Assortment table, you see the components of that assortment (Figure 6–63). You can add a planogram mapping, delete a mapping, change the planogram node for the mapping (re-mapping), or edit the demand spread factor for the mapping.

Figure 6–63 Assortment to Planogram Mapping



Adding Planogram

Use the **Add Planogram** button to add a planogram mapping row that is similar to the mapping row you select. This action copies all the data from the original row with the exception of Node, Season, Demand Spread Factor, Store Count, Stores with Planograms, and Errors. You are presented with a dialog box in which you specify the planogram node. If you cancel out of the dialog box without selecting a node, the row you are adding is removed.

Deleting

Use the Delete button to delete a mapping. If you try to delete the only mapping for a specific combination of location and product, you will see an error message.

Re-Mapping

Use the Re-map button to change the planogram node (displayed in the Node column) for the mapping. You see the Select Planogram Node dialog box. When you change this value, three other changes occur. The value for the planogram season (displayed in the Season column) is changed to the first available season for the node. The value for the demand spread factor is changed to 100 percent. The values for Store Count and Stores with Planograms are both changed to zero.

You cannot select a node that results in a duplicate mapping for a specific combination of location and product when the duplicate mapping is defined as a mapping that has the same location, product, and planogram node (regardless of planogram season).

Editing the Demand Spread Factor

You can edit the demand spread factor within the Assortment to Planogram Mapping table. You must enter a number between 1% and 100%. For a set of mappings for the same combination of location and product, the total of the demand spread factors for the mappings cannot be greater than 100 percent. If the number exceeds this, you will see an error message after you click **Save** (which validates your edits).

Tables for Stores and Clusters The Assortment to Planogram Mapping tables for Stores and for Clusters are slightly different. The following two tables describe each table.

Table 6–43 Assortment to Planogram Mapping for Stores

Field	Description
Store Code	The store code.
Store	The name of the store.
Assortment Cluster	The name of the assortment cluster.
Product Code	The product code.
Product Name	The name of the product.
Planogram Hierarchy - Re-map	The button you use to select a new planogram node for the mapping.
Planogram Hierarchy - Node	The planogram node, which is the sub-category level of the assortment product hierarchy that the mapping applies to.
Planogram Hierarchy - Season	The season that the mapping applies to.
Demand Spread Factor	The amount of demand that should be allocated to the mapping (within a specific combination of location, product, node, and season.
Planogram Available	Y = planogram is available for a specific combination of location, product, planogram node, and planogram season. N = not available.
Errors	The errors that occur during mapping.

Table 6–44 Assortment to Planogram Mapping for Clusters

Field	Description
Assortment Cluster	The name of the assortment cluster.
Product Name	The name of the product.
Demand Spread Factor	The amount of demand that should be allocated to the mapping (within a specific combination of location, product, node, and season.
Planogram Hierarchy - Re-map	The button you use to select a new planogram node for the mapping.
Planogram Hierarchy - Node	The planogram node, which is the sub-category level of the assortment product hierarchy that the mapping applies to.

Table 6–44 (Cont.) Assortment to Planogram Mapping for Clusters

Field	Description
Planogram Hierarchy - Season	The season the mapping applies to.
Store Count	The number of stores in the assortment cluster. This value is zero for mappings that are edited.
Stores With Planograms	The number of stores in the assortment cluster that have planograms for the selected planogram set, based on the planogram data that has been loaded. This value is zero for mappings that are edited.
Errors	The errors that occur during mapping.
Product Code	The product code.

Mapping Errors

Three errors are identified by the mapping process.

- Total demand spread factor (DSF) out of range. Action: The value must not exceed a 100 percent the DSF within the mapping. Results must be adjusted to stay between 1% and 100%.
- Unmapped product. Action: There is no mapping information for a product. Add the mapping information and manually map the product or remove the product from the assortment feed.
- Unmapped store. Action: There is no POG available for the store. Add store information to the POG feed, remove the store from the assortment, or update the mapping information to map the assortment against a different POG Set with information for the store.

Note: After providing new data (with the exception of the assortment files), such as POGs, mapping data, and so on, that specifically provides a solution for the exceptions and errors found by the mapping process, the user must execute the Re-map action from the pull-down menu. This will trigger a new mapping process run for the assortment against the latest data available. The new data picked by the remapping process can clear the exceptions and errors found by the last mapping process and allow the assortment to move forward for optimization.

7

Attribute Extraction

This chapter describes the Attribute Extraction Cloud Service module.

Introduction

Attribute Extraction (AE) is an enterprise-specific solution that uses machine learning to extract product attributes from free-form product description strings.

The application's embedded science and automation helps you to extract the attributes (such as brand, color, flavor, and so on) of each product in a particular category and to normalize the attribute values by correcting short forms, mis-spellings, and other inconsistencies. The product attributes can be used by Demand Transference, Customer Decision Trees, Advanced Clustering, and other retail applications that require product attributes in a structured format.

The AE Cloud Service module consists of the following tabs: Overview, Edit Labels, Annotation, Errors, Normalization, and Results. You use the Overview tab to select one of the previously added product categories or to add a new category. You use the Edit Labels to define category-specific attributes that you want to extract. In the Annotation and Errors tabs, you follow an iterative process to extract attributes and correct any mislabeled attributes. In the Normalization tab, you can use the embedded List of Values (LOV) or create your own LOV to standardize the attribute values. You use the Results tab to review and export the table of attributes.

Overview Tab

The Overview tab displays general information about the product categories previously added by you and other users. You can view this information to determine which user has created each category, when it was last updated, and whether or not the attribute extraction process is complete.

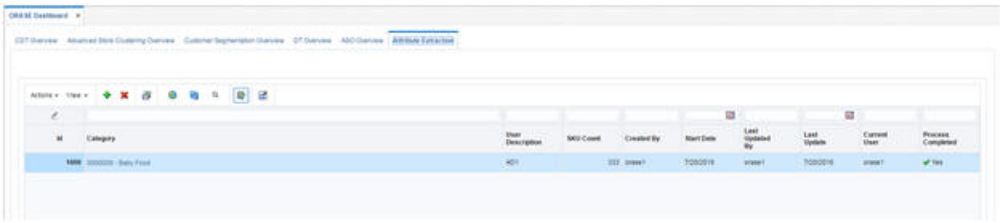
To use the attribute extraction process, you can either select one of the existing categories or you can add a new category. Note that only one copy of each category can exist at any time.

To add a new category, click the **Add** icon. In the pop-up screen, select a category from the list and enter a description. Then click **OK** to return to the Overview tab.

To select a category and start the attribute extraction process, simply click on its name in the overview table.

The AE overview tab is shown in Figure 7–1.

Figure 7–1 Attribute Extraction Overview Tab

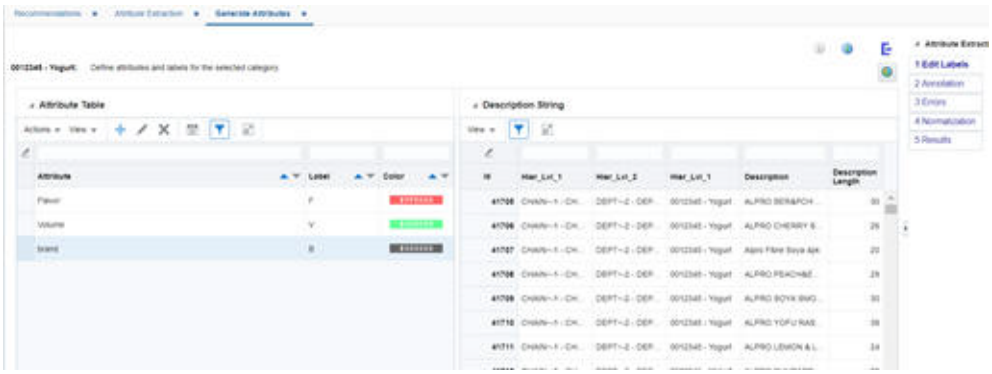


Edit Labels Tab

The Edit Labels tab is used to define category-specific attributes and their labels for the selected category. Defining attributes and labels is the first step in the attribute extraction process. You can return to this tab at any time to add, remove, or modify the attributes.

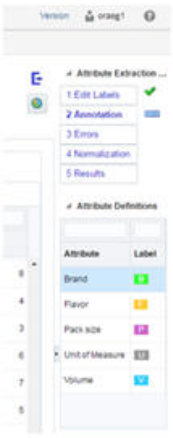
On the right hand side of the tab, you can view the description strings of the products in the selected category as well as three levels of hierarchy for the products. The hierarchy that is shown in the Attribute Extraction tables is based on the product hierarchy data that was loaded for the AE application. (See the Implementation Guide for further information about the hierarchy.) To add an attribute, click the **Add** icon. In the pop-up screen, enter the attribute name and attribute label and select a color. The attribute label must be a single character (the characters "o" or "O" are not permitted) and must be unique for each attribute.

Figure 7–2 Edit Labels Tab



The attributes, labels, and colors that you define in the Edit Labels tab are displayed on the right hand side of all other screens for your reference.

Figure 7–3 Attribute Definitions



Annotation Tab

The Annotation tab is used to assign labels to a single token (referred to as a unigraph) or combination of two adjacent tokens (referred to as a digraph). A token is defined as an alphabetic or numeric sequence of characters in a description string. For example, the string "ABC vanilla yogurt 500gr" is a description that consists of the following five tokens (or unigraphs): "ABC", "vanilla", "yogurt", "500"and "gr". At the same time, it consists of the following digraphs: "ABC vanilla", "vanilla yogurt", "yogurt 500" and "500 gr".

The Annotation tab provides two screens to annotate unigraphs and digraphs.

Unigraphs

In the Unigraph screen in Annotation tab, you can assign attribute labels to tokens, run the machine learning algorithm to find new attribute labels, and review and approve machine-recommended labels.

Assigning User Labels and Reviewing Machine Labels

On the top left of Unigraph screen, you see the Annotation section that contains a table listing all tokens along with the following fields:

Table 7–1 Annotation Fields

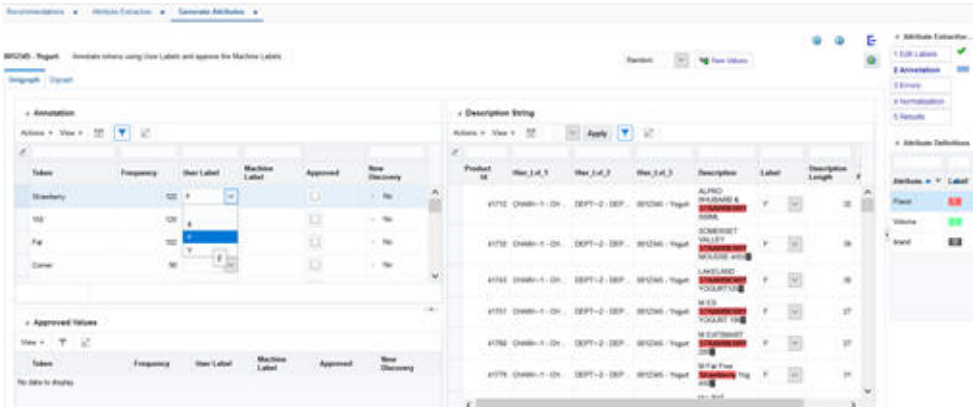
Field	Description
Frequency	The number of times the token appears across all product description strings
User Label	The label you have assigned to the token
Machine Label	The label recommended by the machine learning algorithm
Approved	A check box used to approve and apply the machine label to the token
New Discovery	A Yes/No flag that indicates whether or not the machine-recommended label is a new discovery from the most recent run

In the Annotation section, shown in Figure 7–4, you can perform one of the following tasks for each token:

Table 7–2 Annotation Section Tasks

Task	Description
Change/assign the user label	Select a value from the drop-down menu to change the label or assign a label to a token. (Values in the drop-down menu are the attribute labels that you defined in the Edit Labels tab.) The user label is applied to all instances of the token across all product description strings.
Approve a machine label	Check the Approve check box to assign the machine-recommended value to a token. Once you check the box, the token with the approved machine label is moved from the top tables to the bottom table (the Approved Values section).

Figure 7–4 Annotation Tab with Label Tokens in Unigraph Screen



The Approved Values section on the bottom left displays the tokens that you have approved so far. You can uncheck the Approve check box in this section to remove the approved machine label and move the token back to the top table.

The Description String section on the right side of the Unigraph screen displays all description strings that contain a token selected on the left, along with the length of the strings and position of the selected token in the string, as well as the product hierarchy. The Label column is populated based on the user label or approved machine label for the token. All labeled tokens are colored based on the colors you assigned in the Edit Labels tab. The token selected on the left (in the Annotation section) is shown using a bold font.

You can perform one of the following tasks in the Description Strings section:

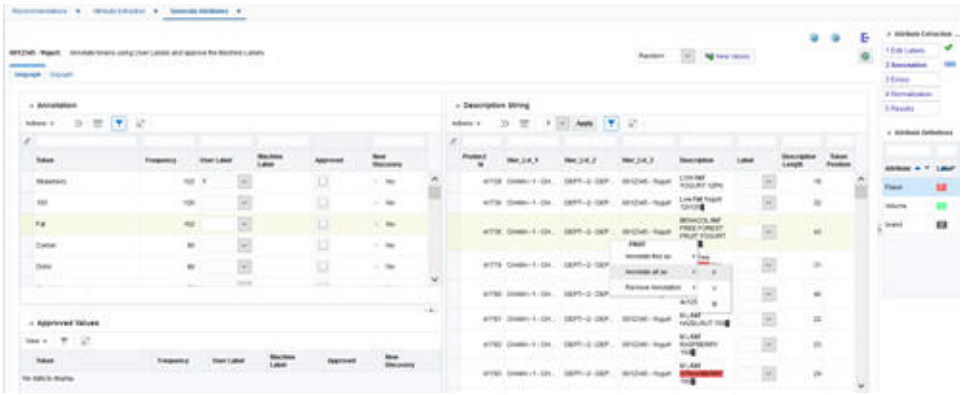
Table 7–3 Description Strings Section Tasks

Task	Description
Change or remove the label for one instance of the selected token	Use the inline drop-down menu in the label column to change or remove the label of the token in a particular description string. Alternatively, right click on any token and use the first option in the menu (Annotate this as) to assign a label to a particular instance. Use the third option in the menu (Remove Annotation -> Remove this label) to remove the label for a particular instance

Table 7–3 (Cont.) Description Strings Section Tasks

Task	Description
Change or remove the label for multiple instances of any of the tokens	Select multiple rows using the Ctrl or Shift keys. Then, change the label using the drop-down list at the top of the table and click **Apply**. To remove the label for multiple rows, change the label to NULL by selecting the first option in the drop-down list at the top of the table. Alternatively, to change or remove the label for ALL instances of a token, right click on any token and use the second option in the menu (Annotate all as) to assign a label to all instances, or use the third option in the menu (Remove Annotation -> Remove all labels) to remove the label for all instances. Example: if you wish to change the label of a token for all products that belong to a certain department, you can first filter the hierarchy columns based on department. Then, select all of the rows using the Shift key and change the label for all selected rows using the drop-down list at the top of table and click **Apply**.

Figure 7–5 Description Strings Tasks



Finding New Values

Once you assigned labels to some of the tokens in Unigraph screen, you may use the New Values button to run the machine learning algorithm that finds new attributes and recommends labels based on the labels that you have assigned and approved so far. Alternatively, you may go to the Digraph screen and label some digraphs before running the algorithm. Note that assigning and approving unigraph and digraph labels and running the algorithm to find new values is an iterative process; so you can repeat these two steps as frequently as you want until you are satisfied with the quality of the results and number of attributes extracted. If the machine learning algorithm does not find enough new values, it is an indication that it requires more information to discover the patterns and be able to recommend new values, and you might want to label some more unigraphs and digraphs before re-running the algorithm.

To run the machine learning algorithm, first select a mode from the drop-down menu next to the New Values button on the top right. You can select from the following three modes:

Table 7–4 Machine Learning Algorithm Modes

Mode	Description
Random	Randomly partitions the data into a training set and a test set to be used by machine learning algorithm.
By attribute	If you select this mode, you will be prompted to select one attribute. The data is then partitioned in such way that all description strings that have one or more tokens labeled as the selected attribute are used as the training set. The remaining descriptions strings are used as test set.
By annotation	The description strings that are 80% labeled are used as training set. The remaining descriptions strings are used as test set.

The By Attribute mode is most effective in the early stages, where very few attribute values have been identified.

Note that after you click the New Values button, it may take few seconds for the machine learning algorithm to run and find new attributes. Then a pop-up message shows how many new values were found. The Machine Label and New Discovery columns in the Annotation section on the top left will be populated to show the new labels and indicate which tokens were labeled by the machine in the most recent run.

Digraphs

The Digraph is used for labeling two adjacent tokens (as opposed to a single token) with the same attribute. The Digraph screen consists of an Annotation section on the top left for labeling digraphs, an Approved Values section on the bottom left, and a Description String section on the right for labeling individual instances of a digraph.

Assigning User Labels to Digraphs

On the top left of Digraph screen, you see the Annotation section that contains a table listing all combinations of adjacent tokens (labeled as "token1" and "token2") where exactly one of the two tokens has been labeled and the other token has no label assigned or approved. The token that already has a label is shown using a bold font.

Other fields in this section are as follows:

Table 7–5 Diagraph Fields

Field	Description
Frequency	The number of times that the two tokens are adjacent to one another across all product description strings.
Label	The label you have assigned to or approved for one of the tokens. (The token that is shown in bold font already has this label.)
Approved	A check box used to approve and apply the label of labeled token to the unlabeled token.

If you think the two tokens are of the same attribute type indicated by the label, check the Approved check box to apply the label to both tokens. Once you check the box, the approved row is moved from the top table to the bottom table (the Approved Values section).

The Approved Values section on the bottom left displays the adjacent tokens that you have approved so far. You can uncheck the approve check box in this section to remove the approved row and return it to the top table.

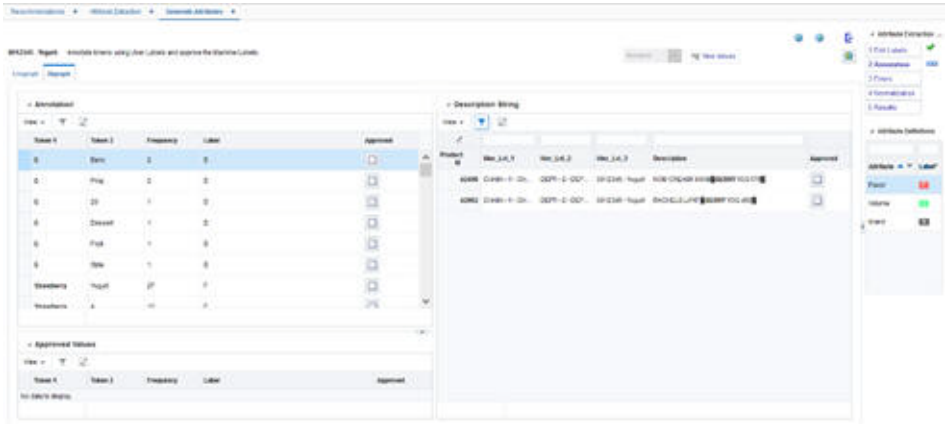
The Description String section on the right side of the Digraph screen displays all description strings that contain the adjacent tokens selected on the left, as well as the product hierarchy. All labeled tokens are colored based on the colors you assigned in the Edit Labels tab. The two tokens selected on the left (in the Annotation section) are shown using a bold font.

You can do one of the following in the Description Strings section:

Table 7–6 Description Strings Section Tasks

Task	Description
Remove the label of the originally unlabeled token for one instance	If you believe the combination of two tokens must not be labeled the same in one or few description strings, un-check the Approved check box to remove the label for the token that was not originally labeled. Note that the token that was originally labeled (i.e., the token displayed in bold font in the Annotation section on the left) remains as labeled and is not affected.
Change or remove the label for one or all instances of any of the tokens	As with the Unigraph screen, you can use the right click menu on any token and use the Annotate this as option to assign a label to a particular instance or use the Annotate all as option to assign a label to all instances of the token. To remove the label for a particular instance or all instances of the token, use Remove annotation -> Remove this label and Remove annotation -> Remove all labels.

Figure 7–6 Annotation Tab Diagraph Screen



Errors Tab

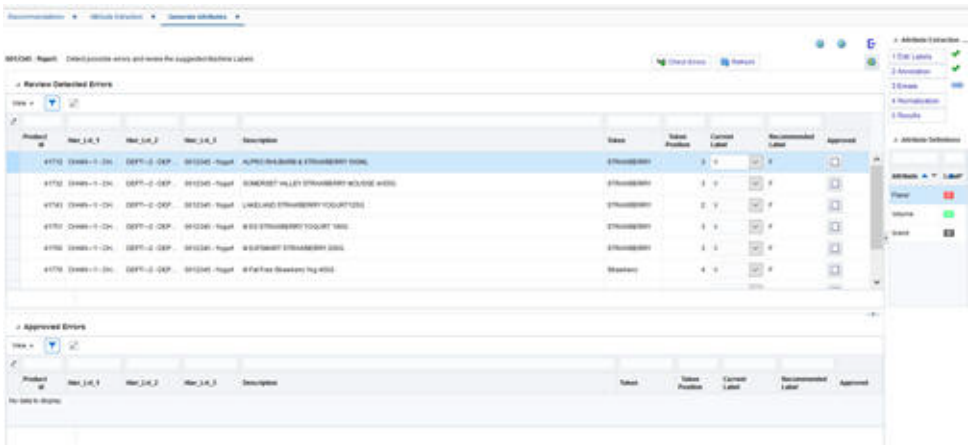
The Errors tab is used to review the instances where the current attribute type assigned to a token is different from the attribute type recommended by the machine. You can review these potential user errors and fix them if needed. For each instance, you can take one of the following actions:

Table 7–7 Error Tab Actions

Action	Description
Approve the Recommended Label	Check the Approved check box to approve the Recommended Label. The approved instance will be moved from the top table to the bottom table. You can use the check box in the bottom table to undo the approval.
Change the Current Label	If neither the Current Label nor the Recommended Label is correct, use the drop-down list to change the Current Label. Do not check the Approved check box in this case. If the Current Label is correct, no action is needed.

It is recommended that you review the errors after you complete the annotation step and extract the attributes. However, you can navigate to the Error tab at any time to review the errors and change or approve the labels for some or all instances.

Figure 7–7 Errors Tab



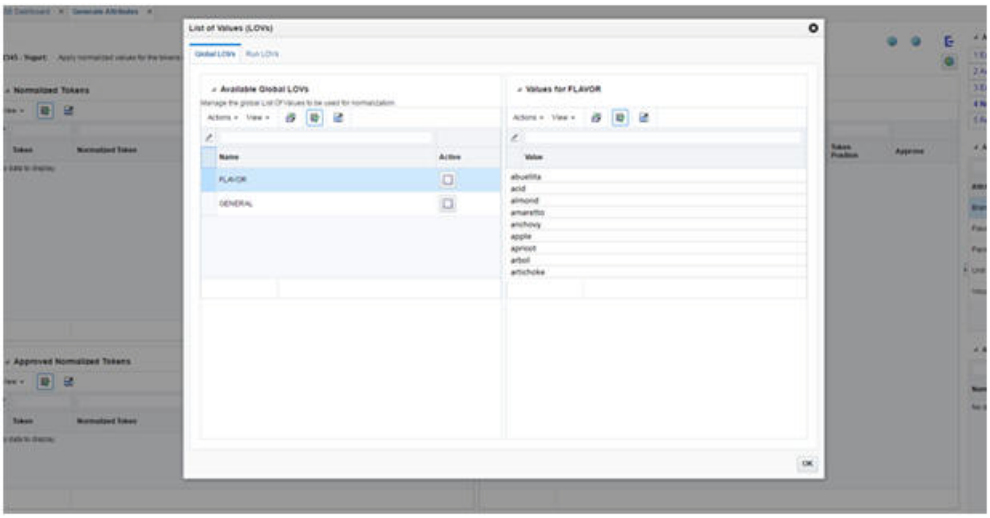
Normalization Tab

After you extract the attributes, you can use the Normalization tab to correct mis-spellings, sort forms, and other inconsistencies in the attribute values to makes sure that each attribute value is displayed in a consistent form across all product descriptions. For example, strawberry is a flavor that may appear in many different forms such as "sberry", "strawb", and ""strberry" across description strings. The purpose of using the Normalization screen is to convert all these different forms to the correct form (i.e., strawberry).

To run the spell-correcting algorithm, you can use one of the pre-defined List of Values (Global LOVs) or you can create your own LOV (Run LOVs) specific to one attribute. To use either of the LOVS, click the List of Values button on the top right.

To select a Global LOV, check the Active check box next to it. There are two Global LOVs: flavor (la ist of all different flavors) and general (a comprehensive list of English words). You can click the row for each LOV to see all the values in the list. If you select the flavor LOV, it will only be used to correct the values of the flavor attribute. If you select the general LOV, it will be used for correcting the values of all attribute types.

Figure 7–8 Normalization Tab with Global List of Values



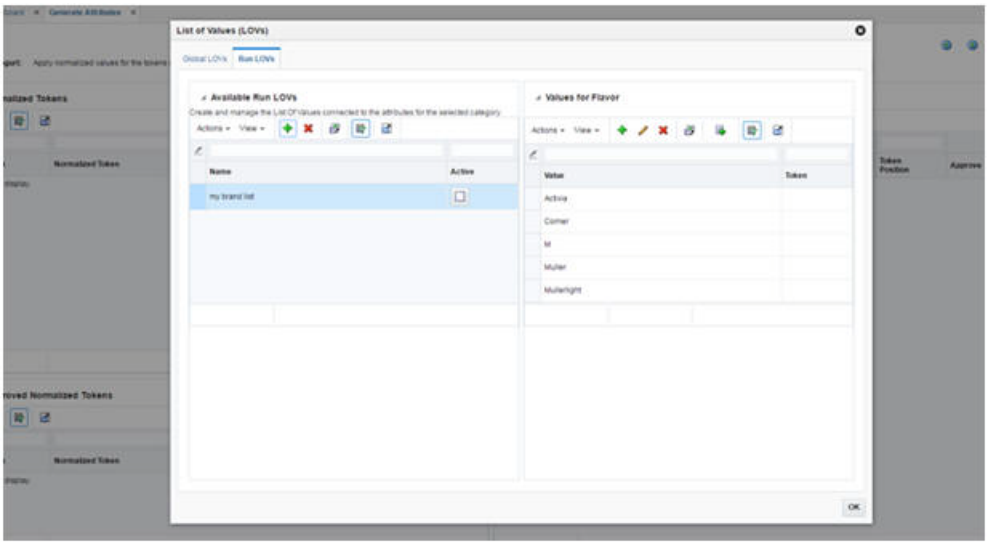
To create a new Run LOV or select an existing Run LOV, navigate to the RUN LOVs tab in the pop-up screen. To create a new list, click the Add button on the left table (Available Run LOVs). Then select the attribute that you want to create the list for and assign a name to the LOV. You can also pre-populate the list with tokens that are labeled as the selected attribute and appear more than a certain number of times across all description strings. The reason for this option is that usually the tokens with high frequency have the correct form and spelling. For example there may be few instances of "sberry" and "strawb", but most certainly there are many instances of "strawberry". So the high frequency tokens are likely to have the correct form and can be used by the spell correcting algorithm.

After you create a Run LOV, you can edit or remove the values in the list or add new values. To edit or remove values, select a row in the right table and use the Edit or Delete button. To add a new value, click the Add button on the top of table on the right.

In addition to adding a new value to the list (to be used by the spelling correction algorithm), you can define a value/token pair so that all instances of the token across all description strings are replaced by the defined value. This is a useful option when the data has lots of abbreviations and short forms that may be difficult to correct using a spelling correction algorithm. For example if the name "Hello Kitty" is a brand that appears as "HK" in many of the description strings, then you can define a pair as "Hello Kitty | HK" to have all instances of "HK" replace by "Hello Kitty".(Note that the correct value and the token must be separated by a pipe delimiter (|)).

Once you are done adding the Run LOV and editing the values, make sure to check the Active check box on the left table to select the list.

Figure 7–9 Normalization Tab with Run List of Values



After you select a Global or Run LOV, click the OK button to close the pop-up screen and return to the Normalization tab. You can see the active LOVs on the bottom right of the screen.

To run the spell correcting algorithm, click the Normalize button. It may take few seconds for the algorithm to run. When the run is complete, the recommended corrections are displayed in the Normalized Tokens table on the top left. This table has the following fields:

Table 7–8 Normalized Tokens

Field	Description
Token	The token identified by the algorithm as misspelled.
Normalized token	The recommended correct value for the token or the value that you defined for replacement (i.e., the value/token pair).
Frequency	The number of times the token appears across all product description strings.
Approved	A check box used to approve and apply the recommended correct value.

You can perform one of the following in Normalization tab:

Table 7–9 Normalization Tab Tasks

Task	Description
Edit the normalized token	If you do not agree with the recommended correction and want to edit the normalized token, you can edit the text in the Normalized Token column before approving it.
Approve/reject the normalized token for all instances	Check the Approved check box to replace all instances of the token with a normalized token. The approved rows will be moved from top table to the bottom table (Approved Normalized Tokens). To undo the approval, uncheck the check box in the bottom table.

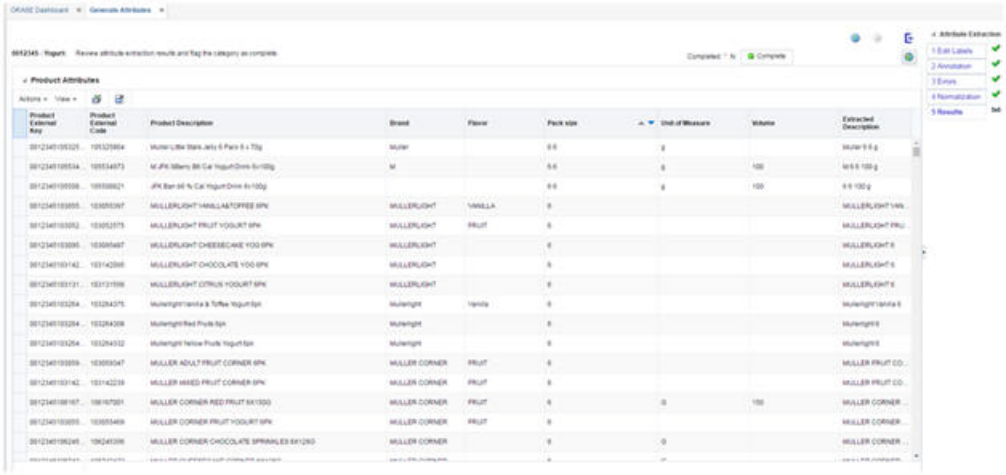
Table 7–9 (Cont.) Normalization Tab Tasks

Task	Description
Approve/reject the normalized token for one instance	When you click on a row in the Normalized Token table, all descriptions that contain the selected token will be displayed in the description strings table on the right. You can approve or reject individual instances of the normalization by checking or un-checking the Approved check box in the right table.
Clear all recommended corrections	To clear all recommended normalized tokens, click the Reset button on top of the Normalized Tokens table on the top left.
Refresh	It is recommended that you click the Refresh button on the top right after you make changes to a LOV or select/deselect a LOV.

Results Tab

The Results tab is used to view the table of attributes. You can export the results into a spreadsheet. You can click the Complete button on the top right to change the status of the category in the overview tab to a value of "complete". This indicates that the attribute extraction process for this category is complete. Other users who log in to the application may decide to work on incomplete categories.

Figure 7–10 Results Tab



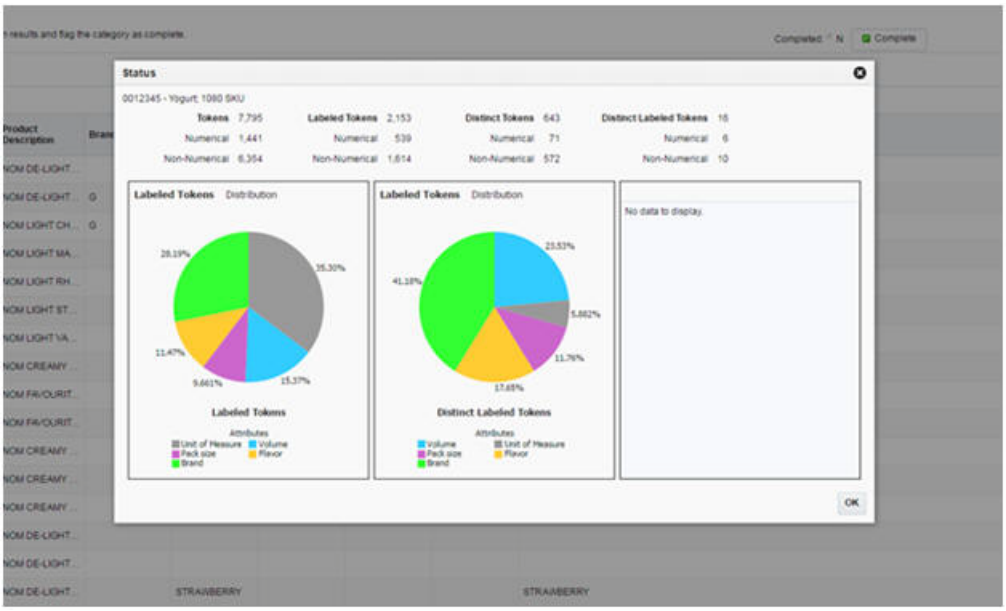
Status

The Status screen is accessible from every tab. To open, click the pie chart button on the top right.

In the Status screen, you can see the progress of the attribute extraction process. It displays different statistics, including total number of tokens, number of distinct tokens, and number of labeled tokens.

You can also see the breakdown of labeled tokens by attribute type. To see the list of all distinct values for a certain attribute, click the respective section of the attribute in any of the pie charts.

Figure 7–11 Status Screen



8

Affinity Analysis

This chapter describes the use of the Affinity Analysis (AA) Cloud Service module.

Introduction

Market basket analysis involves the use of data mining techniques to search for sales patterns between products within a given group of transactions. The output of that analysis provides a rule that defines the association found between products at the subclass or class level of the merchandise hierarchy.

A rule consists of one to three antecedents (IF attributes) and a single consequent (THEN attribute). For example:

IF (milk) and (juice), THEN (cereal)

In other words, if a customer purchases an item from the subclasses milk and juice, the customer will also purchase an item from subclass cereal. After a rule is defined, a user can use the AA interface to understand how strong the affinity is, using rule confidence and support. The probability that a customer will buy milk, juice, and cereal is known as the *support percentage*, while the conditional probability that they will buy cereal when they buy milk and juice is known as *confidence*. Rules with a very high support value occur frequently in your transaction history, while rules with a high confidence value represent a strong affinity between products.

After users have identified selling patterns, they can begin to take action based on those patterns, as well as the needs and goals of their product category. Suppose that a merchant is tasked with bringing in more margin dollars to the cereal category. Using the affinity rule in the preceding example, the merchant might work with the dairy category on a milk promotion to increase sales of milk. This in turn increases the sales of cereal, without sacrificing margin dollars on a cereal promotion. Note that this can require cross-category planning in some cases, as product affinities can sometimes occur between seemingly unrelated products (such as pet food and beer).

Another component of market basket analysis relates to the product assortments being sold in stores. Using product affinities and sales history, AA provides assortment recommendations that improve the revenue or margin of a category by suggesting product additions or removals. Products may be recommended for removal if they are found to be too similar to other products in the assortment (and thus cannibalize the sales of those products). Conversely, products that are not similar to any other items in the assortment may be candidates for inclusion, as they will not divert sales from the existing assortment.

Market basket rules are used to improve the assortment recommendation by showing the potential lift (or halo effect) on your overall sales due to any known affinities on recommended item additions. For example, if AA is analyzing an assortment for

Coffee, and a particular item is part of a market basket rule that drive sales for Milk, then that item has a greater potential value for the lift it brings to the Milk category. AA may then recommend that item over other items in the category, because including it will bring in additional revenue to other assortments without changing those assortments directly.

User Interface

The AA user interface consists of several screens that help the user analyze and take action on the results of the market basket analysis and assortment recommendation processes. The following list summarizes the main purpose of each screen:

- Assortment Recommender - Review the product assortment recommendations made by AA, including the expected halo effects from items in the optimized assortment.
- Top N Class Affinities - Review the market basket affinity rules identified by AA between different product classes, including insights around the frequency and profitability of the top rules.
- Top N Subclass Affinities - Review the market basket affinity rules identified by AA between different product subclasses, including insights around the frequency and profitability of the top rules.
- Top N Promotion Affinities - Review the market basket affinity rules identified by AA between different product subclasses under the effects of a promotion, which can be used to identify the effects promoting a category has on other non-promoted products.
- Top N Customer Segment Affinities - Review the market basket affinity rules identified by AA between different product subclasses under the effects of a promotion and separated by customer segment, which provides insights into the top promoted product affinities for a targeted group of customers.

The way that you interact with the AA user interface depends on your business role and which insights you want to take action on. For example, a category planner looking to make changes to their assortment to increase margin dollars might start from the Assortment Rec screen. The planner might select the category and one of the top-selling locations, and then review the added or dropped items AA has recommended. From there, the planner can choose to move forward with the recommendations and make the necessary adjustments to the assortment plan or dive deeper into the Product Affinities screens to better understand why certain recommendations are being made.

Regardless of which screen you are viewing, a set of global prompts determines the data displayed throughout the application.

On first accessing the application, you make selections from the prompts and click on the **Submit** button, which will load the remaining application screens based on your selections. Modifying and submitting the prompt values after that time will refresh the screens with new data.

Table 8–1 AA Global Prompts

Prompt Name	Description
Top 10 By	Select the primary metric used to rank the market basket results in all places where affinity rules are displayed.
Fiscal Week From	Select the starting week from which results must be displayed. The values displayed are the week of the Fiscal Calendar.

Table 8–1 (Cont.) AA Global Prompts

Prompt Name	Description
To	Select the ending week that results must be displayed for.
Department	Select the department from the merchandise hierarchy that results must be displayed for. All of the classes or subclasses on the THEN side of an affinity rule will be from the selected Department.
Location	Select the location from the location hierarchy that results should be displayed for.

The Top 10 By prompt allows you to select from the following metrics, which are used to determine which affinity rules are displayed and how to sort the results that are shown.

Table 8–2 Top 10 By Options

Metric Name	Description
Net Sales	Chooses the top N rules in each screen, based on the net sales retail amount (e.g., Avg Sales from If Items)
Net Quantity	Chooses the top N rules in each screen, based on the net sales quantity (e.g., Avg Sales Quantity from If Items)
Net Profit	Chooses the top N rules in each screen, based on the net sales profit (e.g., Avg Profit from If Items)

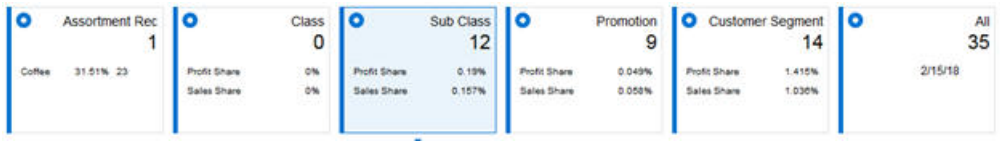
Although this prompt indicates "Top 10", you have the ability to select more than 10 results by changing the Top selection in the Market Basket Filters area of the screen, as discussed later in this chapter.

The Fiscal Week From and To prompts allow you to select a range of dates for which the results are returned. Market basket affinity rules may change over time as new sales transactions are processed, so selecting the weeks to review ensures that only the rules that are applicable for that period of time are displayed.

The Department prompt displays a list of departments from the merchandise hierarchy. Selecting a department will limit the data in all screens such that the consequent (or THEN components) of an affinity rule will be from that department.

After all selections are made and the Submit button is clicked, the interface will refresh and display results. Each section of the application is organized under an informational tile that describes the purpose of the screen, along with summarization metrics.

Figure 8–1 Info Tiles



Each tile has the following general properties.

Table 8–3 Info Tile Components

Component	Description
Title	A descriptive name for the data contained below the selected tile.
Primary Metric	A key measure that highlights important data, such as the number of affinity rules available for review for a given set of prompt selections.
Secondary Metrics	Additional measures that describe the underlying data available when clicking on the info tile.

Clicking on any info tile's title or primary metric will display the associated data in the lower section of the screen. The data shown in each tile in the application is described in Table 8–4.

Table 8–4 Info Tile Metrics

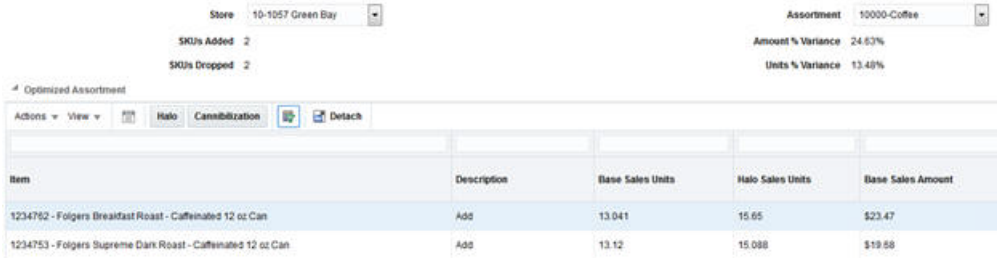
Tile Name	Primary Metric Description	Secondary Metric Description
Assortment Rec	The number of assortment recommendations available for the selected global prompt values.	For the assortments in the department, the percentage the sales amount changed and the amount the sales changed.
Class	The number of class level market basket affinity rules available for the selected global prompt values.	The share of sales profit and sales amount that is represented by all discovered association rules out of the total sales profit and sales amount processed.
Sub Class	The number of subclass level market basket affinity rules available for the selected global prompt values.	The share of sales profit and sales amount that is represented by all discovered association rules out of the total sales profit and sales amount processed.
Promotion	The number of promoted subclass level market basket affinity rules available for the selected global prompt values.	The share of sales profit and sales amount that is represented by all discovered association rules out of the total sales profit and sales amount processed.
Customer Segment	The number of customer segment promoted subclass level market basket affinity rules available for the selected global prompt values.	The share of sales profit and sales amount that is represented by all discovered association rules out of the total sales profit and sales amount processed.
All	The total number of market basket affinity rules available for the selected global prompt values.	The last date that market basket insights were executed on.

Assortment Recommender

The assortment recommender screen provides insights on ways to improve the mix of products in a store assortment, using a combination of data from Demand

Transference and Affinity Analysis. This section describes the data shown in the screen, as well as how to use the results to take action within your business. For more information on Demand Transference, refer to Demand Transference.

Figure 8–2 Assortment Recommender



The first step in using the assortment recommender screen is to select a store and assortment to review. The stores listed in the drop-down menu are limited to those stores that have completed an assortment recommendation calculation for the selected department in the selected period. The assortment drop-down menu is limited to those assortments with a completed assortment recommendation calculation for the date range, department, and location previously selected, and align with the level of the product hierarchy that assortment-planning operations occur at. For example, if we process Coffee on Week 1, Yogurt on Week 2, and Milk on Week 3, and we pick a date range of Weeks 1-3 in the global prompts, we will show results for those three categories.

After selecting a store and an assortment, the screen displays data showing the system-optimized product list and the original product list, along with several summary metrics.

Table 8–5 Assortment Optimization Summary Metrics

Field Name	Description
SKUs Added	The number of SKUs added to the assortment by the optimization process.
SKUs Dropped	The number of SKUs removed from the assortment by the optimization process.
Amount % Variance	The change in sales retail amount for the assortment after SKUs are added and dropped by the optimization, based on average weekly sales.
Units % Variance	The change in sales units for the assortment after SKUs are added and dropped by the optimization, based on average weekly sales.

Table 8–6 lists the summary metrics that the optimized assortment table displays insights into expected product performance if the recommendations are applied to your assortment.

Table 8–6 Optimized Assortment Table

Field Name	Description
Item	The item number and description of the SKU in the assortment.
Description	The recommended change for a SKU in the assortment. This may either be to keep an existing SKU or add a SKU that was not previously in the assortment.

Table 8–6 (Cont.) Optimized Assortment Table

Field Name	Description
Base Sales Units	The average weekly sales units of the SKU at the selected location.
Halo Sales Units	The average weekly sales units of the SKU at the selected location, after adjusting for demand as a result of Halo sales due to market basket affinity rules involving the SKU.
Base Sales Amount	The average weekly sales amount of the SKU at the selected location.
Halo Sales Amount	The average weekly sales amount of the SKU at the selected location, after adjusting for demand as a result of Halo sales due to market basket affinity rules involving the SKU.
Base Sales Profit Amount	The average weekly sales profit of the SKU at the selected location.
Halo Sales Profit Amount	The average weekly sales profit of the SKU at the selected location, after adjusting for demand as a result of Halo sales due to market basket affinity rules involving the SKU.
Incremental Demand	The number of sales units of the SKU that do not transfer to any other SKU in the assortment if it were to be deleted from the assortment.
Substitutable Demand	The number of sales units of the SKU that can transfer to other SKUs in the assortment if it were to be deleted from the assortment.

The primary purpose of the recommendations is to highlight the optimal mix of products, based on a pre-defined optimization target (which can be sales units, amount, or profit). The optimization process takes into account the product similarities and assortment elasticity (as generated by Demand Transference) to determine which products are candidates for addition or removal from the current assortment. It also uses the market basket affinity rules to identify additional sales generated by the SKUs in the assortment in other product categories. For example, a SKU that belongs to Subclass A, which has a strong affinity with sales of Subclass B, is effectively generating a certain amount of additional sales outside of the assortment itself just by being included.

This combination of inputs into the optimization process may result in SKUs being added or dropped from the assortment. A SKU that has been dropped from the assortment may have been found to be too similar to other existing SKUs (thus its removal will transfer the majority of the demand for that SKU to other similar items). A dropped SKU may also have had a low contribution to sales in that assortment, either due to a lack of affinity rules generating additional halo effects, or because other SKUs available for addition provided better results and were swapped in. The system may recommend product additions when it has identified SKUs that will improve the optimization target, such as sales profit. It may also recommend adding a SKU when there are no similar SKUs currently in the assortment, so adding the new SKU will not cannibalize much demand from other items.

If you want to analyze the recommendations further, the system provides pop-up windows displaying Halo and Cannibalization effects relating to the selected item from the Optimized Assortment table. Select a row in the table, and then click the Halo or Cannibalization buttons on the task bar to review the data.

Figure 8–3 Halo Effects Pop-Up

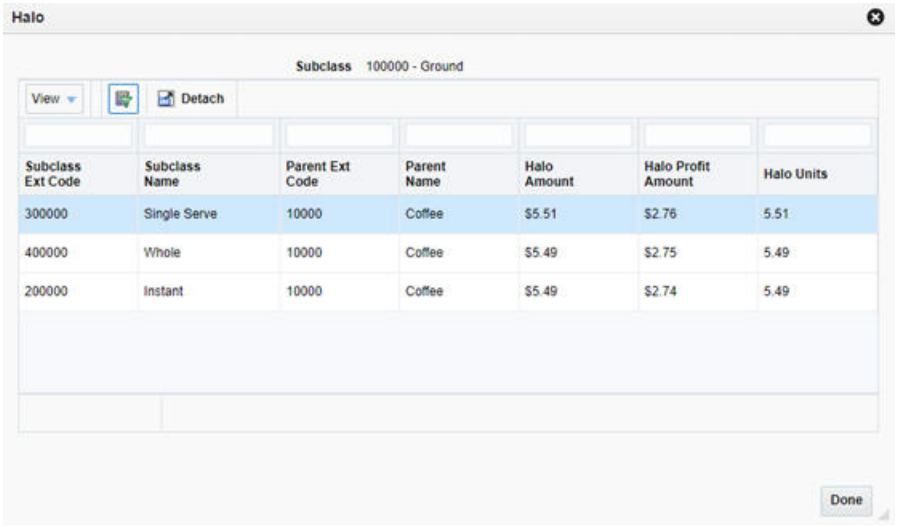


Table 8–7 Halo Effects

Field Name	Description
Subclass Ext Code	The numeric identifier for a subclass with a halo effect from the SKU.
Subclass Name	The descriptive name for a subclass with a halo effect from the SKU.
Parent Ext Code	The numeric identifier of the parent class for a subclass.
Parent Name	The descriptive name for the parent class for a subclass.
Halo Amount	The calculated halo effect on the subclass's sales amount.
Halo Profit Amount	The calculated halo effect on the subclass's sales profit.
Halo Units	The calculated halo effect on the subclass's sales units.

The Halo pop-up window provides insights into the secondary effect of a SKU on other products. Products with strong market basket affinities may generate large halo effects and thus be good candidates for your assortment.

When evaluating the additions and removals on your assortment, it can also be helpful to understand cannibalization effects. These effects are based on the product similarities and assortment elasticities obtained from Demand Transference. Products that are very similar to each other tend to split up the customer demand for this kind of product, since the customer is not likely to buy two similar products. Conversely, products that are not very similar will not have as much overlapping demand. The Cannibalization pop-up window provides some of these insights.

Figure 8–4 Cannibalization Effect Pop-Up

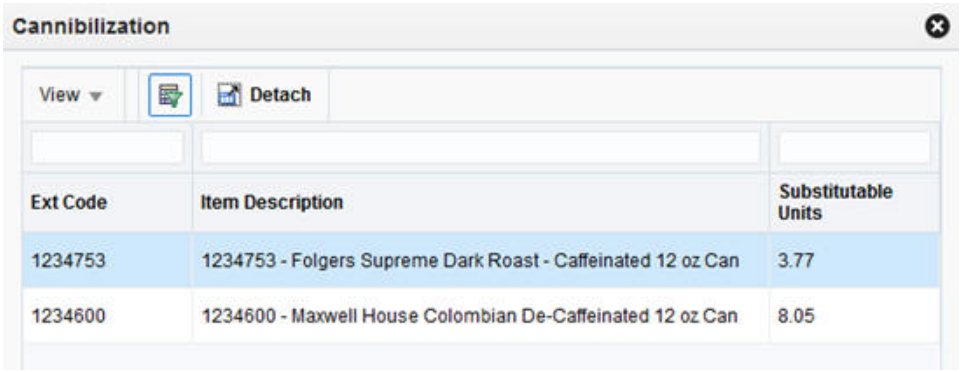


Table 8–8 Cannibalization Effects

Field Name	Description
Ext Code	The numeric identifier for a SKU that is cannibalized by the selected assortment SKU.
Item Description	The descriptive value for the SKU.
Substitutable Units	The calculated number of sales units that can be cannibalized from other SKUs when the assortment SKU is added, or the number of units transferred to these SKUs when the assortment SKU is dropped.

At the bottom of the screen, you can also see the original assortment, along with any SKUs from the assortment that have been dropped in the recommendation.

Figure 8–5 Original Assortment

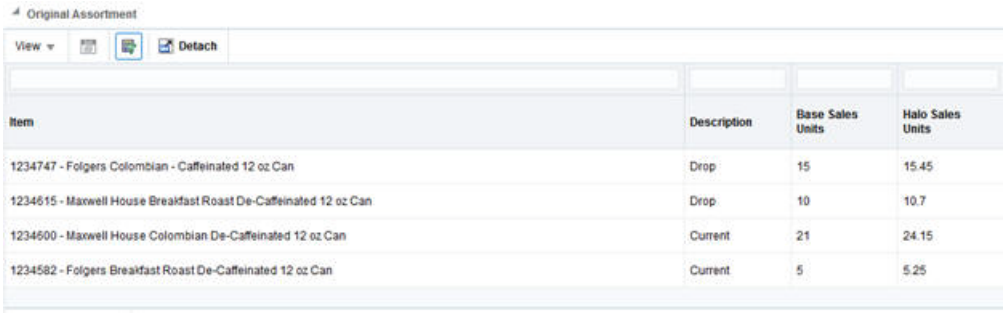


Table 8–9 Original Assortment

Field Name	Description
Item	The item number and description of the SKU in the assortment.
Description	The recommended change for a SKU in the assortment. This may either be to keep an existing SKU or drop a SKU that was in the assortment.
Base Sales Units	The average weekly sales units of the SKU at the selected location.
Halo Sales Units	The average weekly sales units of the SKU at the selected location, after adjusting for demand as a result of Halo sales due to market basket affinity rules involving the SKU.

Table 8–9 (Cont.) Original Assortment

Field Name	Description
Base Sales Amount	The average weekly sales amount of the SKU at the selected location.
Halo Sales Amount	The average weekly sales amount of the SKU at the selected location, after adjusting for demand as a result of Halo sales due to market basket affinity rules involving the SKU.
Base Sales Profit Amount	The average weekly sales profit of the SKU at the selected location.
Halo Sales Profit Amount	The average weekly sales profit of the SKU at the selected location, after adjusting for demand as a result of Halo sales due to market basket affinity rules involving the SKU.

Using Assortment Recommender Results

The assortment recommendations provided within AA aim to achieve the maximum possible value for the chosen target sales measure, such as profit or revenue. Assortment planners may use the results of this process while preparing for the next selling season or fiscal period in order to adjust their current assortments based on their business strategy. Planners may choose to accept the recommendations as-is, or choose a subset of changes in order to see how that will affect their financial targets and assortment plan. This modified assortment can then be provided to AA and reprocessed in the future to refine the recommendations.

The following example describes the workflow of an assortment planner (Anne) working to improve sales in the Coffee category.

1. Anne starts with the existing assortment of coffee items available in stores, along with current performance measures describing how each item is selling. She knows that certain items are performing poorly and wants to mark them for removal from the assortment. She also has a list of new items that can be added to the assortment. All of this information is provided as inputs to AA prior to performing the optimization process.
2. After AA has been provided with the assortments and lists of possible changes that Anne wants to make, the system will schedule and execute the optimization process using pre-defined business rules (such as the maximum number of items that can be added or dropped by the system).
3. Knowing the process executed over the weekend, Anne logs into AA on Monday morning to review the results. She selects a location and her Coffee assortment in the UI prompts, and then begins to analyze the recommendations.
4. AA has selected three items in the current assortment to be dropped and another four items that should be added. Anne first looks at the dropped items and compares them to the other items she was thinking about removing. She notes that the dropped items are not the worst-selling ones she had chosen, but AA shows them as contributing very little to halo effects in other areas, so they may have very weak market basket affinities that contribute to their removal.
5. Anne next reviews the recommendation for items to add to the assortment. Of the four items that AA has chosen, three of them have significant halo effects driven by the market basket affinities. Anne is not as sure about the last item, so she selects it and clicks on the Cannibalization button. She notes that the item has relatively low substitutable demand from other items, suggesting it is not very similar to anything in her current assortment. This makes it a good candidate for addition, as it can bring in new demand that the assortment may not have today.

6. Anne decides to accept all of the recommendations made by AA for this assortment, exports the results to Excel for later reference, and then exits the system to make the necessary changes to her assortment plan.

Market Basket Analysis

The market basket analysis screens of the application all relate to the market basket affinity rules generated by AA. The tiles and screens provide a way to view different sets of data, depending on the configurations used for your business.

Market basket rules and their associated metrics are a key insight into product demand and customer buying behaviors. Understanding which products are more likely to sell together allow your business to coordinate targeted offers and promotions for products that have strong affinities elsewhere, without having to promote those other categories. It is also possible to understand market basket differences between customer segments, allowing you to further refine your sales and promotion strategies to target specific groups of customers that will bring in the most value to your business.

There are five screens available for viewing the AA results, each containing a different subset of affinity rules with certain characteristics as described below.

Table 8–10 Market Basket

Field Name	Description
Class	Displays market basket affinities between different classes, without any consideration for whether the items were on promotion or not.
Subclass	Displays market basket affinities between different subclasses, without any consideration for whether the items were on promotion or not.
Promotion	Displays market basket affinities between different subclasses when the antecedent (if) components are promoted and the consequent (then) components are not promoted.
Customer Segment	Displays market basket affinities between different subclasses when the antecedent (if) components are promoted and the consequent (then) components are not promoted, further separated by customer segments.
All	Displays a searchable list of all available affinity rules for all types.

After you select one of the five info tiles at the top of the screen, you will be presented with additional filter prompts and a data table for displaying results.

Figure 8–6 Market Basket Filters



The filters above the table allow you to limit the data displayed in the table. AA can generate a large number of rules, and usually only the most common rules supported by the largest amount of sales transactions are used in analyses. Which filters are displayed will depend on the current screen being viewed. Table 8–11 provides an explanation of how all the filters impact the displayed results.

Table 8–11 Market Basket Filters

Filter Name	Description
Top	Selecting a different number changes how many market basket rules are displayed in the results table.
All	Selecting the check box overrides the Top filter to instead display all available market basket rules.
Classes/Subclasses	Select specific classes or subclasses for which market basket rules should be displayed. This filter will be applied as an OR condition for either the IF components or the THEN components. If there are multiple IF components, then a match is only required for one of the components.
Customer Segments	Select specific customer segments for which market basket rules should be displayed.
Select Product Hierarchy	(All Screen Only) This prompt is used to choose which level of the hierarchy you are entering search criteria for.
Show only Affinities with Promotions	(All Screen Only) This prompt allows you to choose whether the results should be limited to promoted subclass affinities.

Once you have selected your desired filter settings, you may begin to review the market basket affinities. The main table on each screen describes the product affinities that were found when analyzing the market basket data, sorted in descending order by your chosen performance metric (the Top 10 By global prompt described earlier in this chapter). In addition to the product affinities, numerous metrics are provided to describe the market basket data that contains these products.

Figure 8–7 Market Basket Affinities



Customer Segment	Promotion	If	Then	Baskets	Avg Sales from "If" Items	Avg Profit from "If" Items	Avg Sales Quantity from "If" Items	Avg Sales from Affinity Items	Avg Profit from Affinity Items	Avg Sales Quantity from Affinity Items
		Ground	Single Serve	115	$662.81	$14.59	192.571	$486.06	$62.28	176.286
		Single Serve	Ground	115	$486.06	$62.28	176.286	$662.81	$14.59	192.571

Table 8–12 Market Basket Affinities

Field Name	Description
Customer Segment	(Customer Segment screen only) Displays the customer segment for which the product affinity was found.
Promotion	(All affinities screen only) Displays whether the rule represents IF components that were on promotion.
If	Displays the classes or subclasses containing items that were purchased in the market baskets for a given product affinity. All market baskets that contributed to a rule will contain items from these product categories.
Then	Displays the class or subclass that was found to have an affinity with the IF components of the market basket rule. Market baskets containing the IF categories were found to contain items from the THEN category enough times to establish a product affinity between the IF and THEN components.
Baskets	The number of market baskets that contain items from both the IF and THEN components of the rule.
Avg Sales from If Items	The average sales retail amount from items belonging to the IF classes or subclasses.

Table 8–12 (Cont.) Market Basket Affinities

Field Name	Description
Avg Profit from If Items	The average sales profit from items belonging to the IF classes or subclasses.
Avg Sales Quantity from If Items	The average sales units from items belonging to the IF classes or subclasses.
Avg Sales from Affinity Items	The average sales retail amount from items belonging to the THEN class or subclass.
Avg Profit from Affinity Items	The average sales profit from items belonging to the THEN class or subclass.
Avg Sales Quantity from Affinity Items	The average sales units from items belonging to the THEN class or subclass.

The purpose of these summary metrics is to provide a quick view of the sales and profitability of the market baskets described by each rule. Even though a market basket affinity is present in a large number of baskets, it may not generate much profit or be of particular interest to the business user at this time. The summary metrics allow the user to quickly identify market basket affinities that are worth analyzing in more detail. Once the user has selected a rule to investigate, clicking that row in the table generates additional information about it in a panel to the right side of the screen.

Figure 8–8 Market Basket Affinity Details



Table 8–13 Market Basket Affinity Details

Field Name	Description
Support	The percent of market baskets that contain items from the IF and THEN components of the rule. Written as a formula: Rule Transaction Count/Total Transaction Count
Confidence	The ratio of the number of transactions where the entire rule is present, compared to all transactions containing IF items. Written as a formula: Rule Transaction Count/Total IF Transaction Count

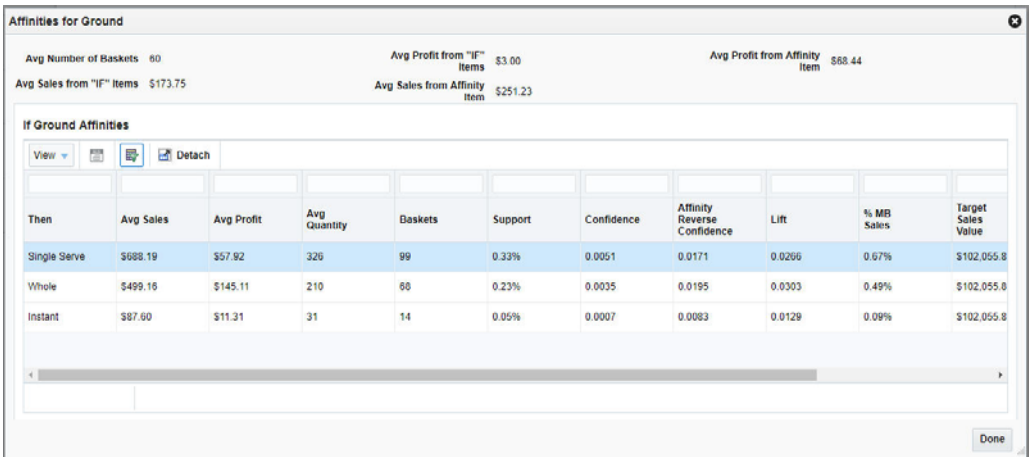
Table 8–13 (Cont.) Market Basket Affinity Details

Field Name	Description
Reverse Confidence	The ratio of the number of transactions where the entire rule is present, compared to all transactions containing THEN items. If a rule has a higher Reverse Confidence than its own Confidence value, then it indicates that the rule is better represented by the reversal of the IF and THEN components. Written as a formula: Rule Transaction Count / Total THEN Transaction Count
Lift	A statistical measure of how strong the product affinity is, as compared to random chance for the rule (higher is better). Written as a formula: Rule Transaction Count * Total Transaction Count / Total IF Transaction Count/Total THEN Transaction Count
Sales Amt Pct Of Total	The percent of total sales retail amount that the rule represents. Written as a formula: Rule Sales Amount/Average Weekly Total Sales
Avg Rule Sales Amt	The average weekly sales retail amount across all transactions containing the rule.
Avg Rule Sales Profit Amt	The average weekly sales profit across all transactions containing the rule.
Avg Rule Sales Qty	The average weekly sales units across all transactions containing the rule.
Avg Total Sales Amt	The average weekly sales retail amount across all transactions.
Avg Tot Sales Profit Amt	The average weekly sales profit across all transactions.
Avg Tot Sales Qty	The average weekly sales units across all transactions.
Avg Tot Txn Cnt	The average weekly transaction count across all transactions.

The detailed metrics for a rule provide more ways to evaluate if the rule is going to be useful for your business purpose. For example, a rule that has both high support and high confidence represents a strong product affinity that is purchased by many of your customers. Any action taken on the categories in such a rule, such as promoting a subclass on the IF side of the rule, can have significant benefits for the THEN components of the rule. Another way to use these metrics is to consider the average sales or profit that a rule brings in relative to total sales over the same period. Rules that represent a significant profit amount may be a more valuable target for your efforts than rules with lower profitability.

Another way to view the market basket affinities is to click the button found to the left of each row in the data table. This button displays a pop-up window with all of the product affinities that have been found for the IF components of the rule. This can be especially useful when your goal is to find any relationships with a specific product category, rather than seeing all of the relationships between different products.

Figure 8–9 Targeted Product Affinities Pop-Up



This screen is split into two main sections: the top section displays summary metrics for all rules associated with the IF components, and the bottom section displays all of the product affinities found for the IF components and a selection of detailed metrics for them.

Table 8–14 Targeted Product Affinities Summary Metrics

Field Name	Description
Avg Number of Baskets	The average number of baskets found on a rule containing the IF items.
Avg Sales from IF Items	The average sales retail amount for the IF items in the rules.
Avg Profit from IF Items	The average sales profit for the IF items in the rules.
Avg Sales from Affinity Item	The average sales retail amount for the THEN items in the rules.
Avg Profit from Affinity Item	The average sales profit for the THEN items in the rules.

Table 8–15 Targeted Product Affinities Detail Metrics

Field Name	Description
Then	Displays the class or subclass that was found to have an affinity with the IF components of the market basket rule.
Avg Sales	The average weekly sales retail amount for the IF and THEN items in the rules.
Avg Profit	The average weekly sales profit for the IF and THEN items in the rules.
Avg Quantity	The average weekly sales units for the IF and THEN items in the rules.
Baskets	The number of market baskets that contain items from the IF and THEN components of the rule.
Support	The percent of market baskets that contain items from the IF and THEN components of the rule.
Confidence	The ratio of the number of transactions where the entire rule is present, compared to all transactions containing IF items.
Affinity Reverse Confidence	The ratio of the number of transactions where the entire rule is present, compared to all transactions containing THEN items.

Table 8–15 (Cont.) Targeted Product Affinities Detail Metrics

Field Name	Description
Lift	A statistical measure of how strong the product affinity is, as compared to random chance for the rule (higher is better).
% MB Sales	The percent of total sales retail amount that the rule represents.
Target Sales Value	The average weekly total sales retail amount across all transactions in the selected periods.
Target Profit	The average weekly total sales profit across all transactions in the selected periods.
Target Sales Quantity	The average weekly total sales units across all transactions in the selected periods.

This view of market basket affinities provides powerful insight into the relationships between different product categories, as it may expose affinities even between seemingly unrelated categories that your customers tend to purchase together. This screen allows you to see how actions taken on the target (IF) products will impact other areas within your business. For example, any action taken to increase the sales of an IF category can lift the sales of all other products with a strong affinity to those items. Conversely, if you have a category that is performing poorly on the THEN side of a rule, you may be able to take action on the IF products to boost that category's sales indirectly.

Using Market Basket Rules

There are a variety of ways to take action within your business based on the insights that AA provides. The most common use is to develop a promotion strategy that takes advantage of product affinities to maximize halo effects across other categories. By reducing the amount spent on promotions while increasing the effect of those promotions, you can realize additional revenue and boost profit margins. For example, by analyzing the AA results, it is possible to identify multiple product categories with affinities on a target category, where some categories are more profitable than others. This can provide alternatives for promoting items that will yield more in sales or profit than a more obvious affinity like Hot Dogs and Hot Dog Buns.

Another common practice is for store planners and planogrammers to take advantage of product affinities when deciding how to arrange the products on the shelf or for adding aisle end-caps for strongly associated product categories. AA can show obvious relationships like Bread and Peanut Butter, but it may also reveal previously unknown associations like Pet Food and Beer. Using market basket rules to inform store layouts allows you to place commonly purchased groups of products close together in the store, increasing the chance of customers purchasing more from those categories.

Combining market basket analysis with customer segments further enhances the potential benefits of AA. Customer behavior information is obtained from mining transaction history, and it is correlated with customer segment attributes to inform your promotion strategies. The ability to understand market basket affinities allows marketers to calculate, monitor, and build promotion strategies based on critical metrics such as customer profitability and preferred categories.

As an example, consider a marketer that is planning promotions for the Soda and Chips categories. While analyzing the data in Affinity Analysis, the marketer finds that a particular customer segment, College Singles, has a very strong affinity for buying Soda any time they purchase Chips.Armed with this knowledge, the marketer

may decide to create a targeted offer for this customer segment only on the Chips category, knowing that they are likely to also buy Soda in the same basket. This kind of offer is more cost effective than a company-wide promotion on both the chips and soda categories, but may yield almost the same results, because the offer has focused in on the customers that will generate the most revenue in these categories.

Affinity Analysis

This section addresses MBI analysis. The functionality described here includes creating custom runs and comparing the results of multiple runs.

System runs are executed automatically once a week as part of the weekly batch cycle using settings that have been configured for a specific level of the location hierarchy. Since system runs occur weekly they can be used to examine results over a long period of time.

User runs can be run at any time by an individual user who specifies the settings for the run.

Two types of comparison are available: Compare Runs and Compare Results.

Use Compare Runs to select specific runs to compare. This comparison is useful for straightforward runs such as comparing the results of two different months.

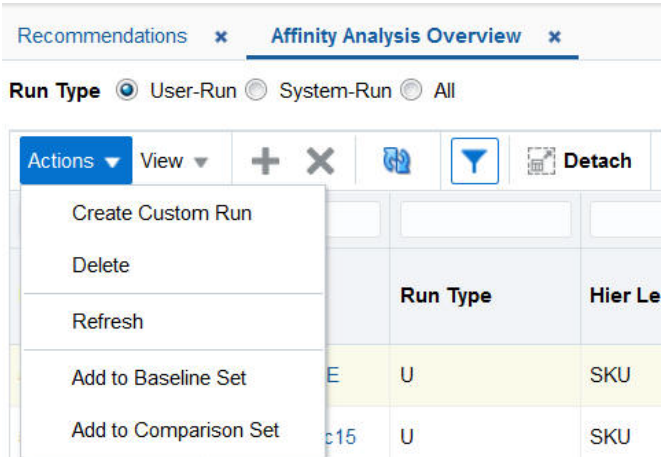
Use Compare Results to compare system runs only, with specified dates and other filters. In this case, take care in selecting runs that overlap, as some runs may be represented in the results more than once.

Affinity Analysis Overview

The Affinity Analysis Overview screen provides the tools shown in Figure 8–10. In addition to the standard application functionality, you can also

- Select from the following functionality using the Action drop-down menu:
  - Create Custom run
  - Delete
  - Refresh
  - Add to Baseline Set
  - Add to Comparison Set
- Select the type of run to view in this screen: User-Run, System-Run, or all runs.

Figure 8–10 Affinity Analysis Overview Functionality



The Affinity Analysis Overview table, shown in Figure 8–11, displays the list of completed runs and contains the following fields.

Figure 8–11 Affinity Analysis Overview Table

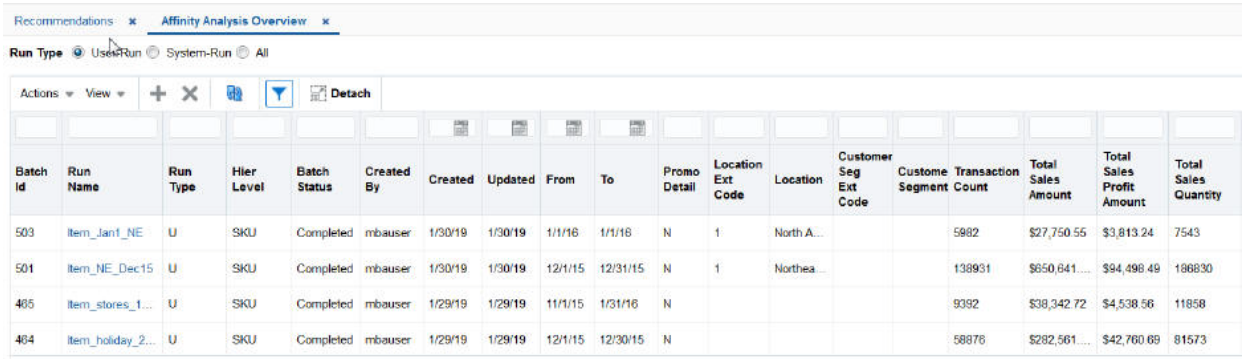


Table 8–16 Affinity Analysis

Field Name	Description
Batch ID	Identifies the run.
Run Name	User-provided name for the run.
Run Type	U = User -Run; S = System-Run.
Hier Level	Hierarchy level.
Batch Status	Indicates the status of the run. Helps to determine if the run has completed, encountered any errors, or has not yet completed.
Created By	User name of person who created the run.
Created	Date when run created.
Updated	Date when run last updated.
From	Beginning date for run data collection.
To	Ending date for run data collection.

Table 8–16 (Cont.) Affinity Analysis

Field Name	Description
Promo Detail	Y indicates that the market basket affinities between different subclasses are displayed when the antecedent (if) components are promoted and the consequent (then) components are not promoted. N indicates that no consideration is given to whether or not the items are promoted.
Location Ext Code	The numeric identifier for the location.
Location	The external identifier for the location.
Customer Seg Ext Code	The external identifier for the customer segment.
Customer Segment	The name of the customer segment which the results are for. Values include Soccer Mom, Simply Seniors, Value Seekers, Gourmet Shoppers, Livin Large.
Transaction Count	Total number of transactions processed for this run.
Total Sales Amount	Total sales amount processed for this run.
Total Sales Profit Amount	Total sales profit amount processed for this run.
Total Sales Quantity	Total sales quantity processed for this run.

When you click on a specific **Run Name**, you see a pop-up that displays Filter Results and Affinities Search Results.

Filter Results

Figure 8–12 shows the Filter Results and Affinities Search Results, which are displayed when you click on a specific Run Name. The fields are described in Table 8–17 and Table 8–18. Use the **Search** button to initiate the filtering process and use the **Reset** button to clear the filter fields.

Figure 8–12 Filter Results and Affinities Search Results Table

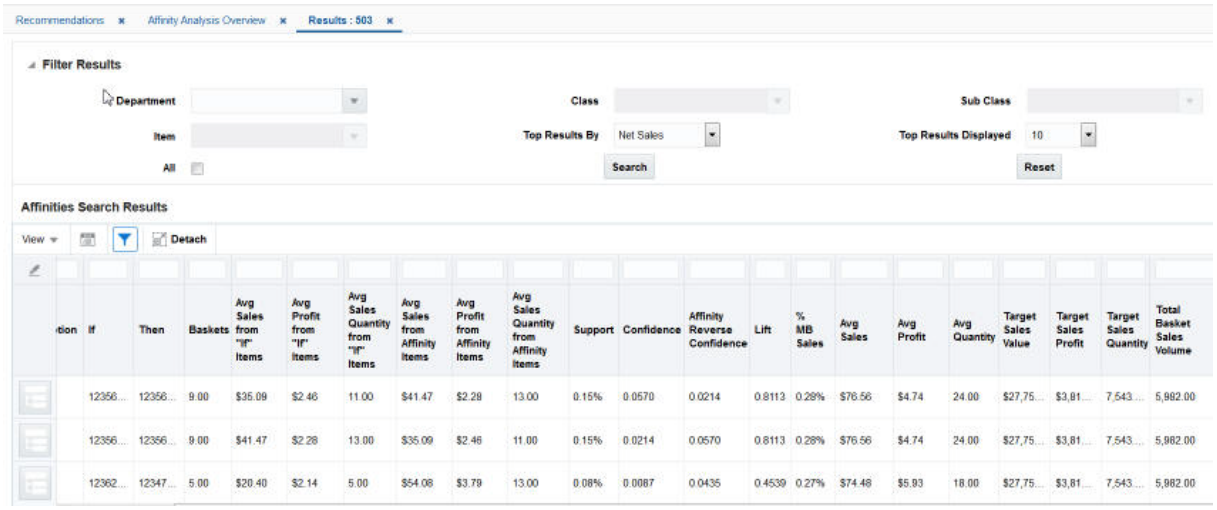


Table 8–17describes the options you can use to filter the results that will be displayed.

Table 8–17 Filter Results

Field	Description
Department	Select one or more departments for which the results should be filtered.
Item	Select the item within the department from the drop-down list.
All	Select to prevent any other filters from being used.
Class	Select one or more classes for which the results should be filtered.
Top Results By	Select a metric to see the top results by: Net Sales, Net Quantity, of Net Profit.
Sub Class	Select one or more subclasses for which the results should be filtered.
Top Results Displayed	Select a value (10, 20, 30...970, 980) to limit the number of results displayed.

Affinities Search Results

The Affinities Search Results section of the pop-up provides details about the observed associations (affinities) between sales transactions. It is expressed as "If a customer buys X, then that customer may also buy Y." The table describes the relevant values associated with the If/Then transactions and contains the following fields:

Table 8–18 Affinities Search Results

Field	Description
Customer Segment	(Customer Segment screen only) Displays the customer segment for which the product affinity was found.
Promotion	(All affinities screen only) Displays whether the rule represents IF components that were on promotion.
If	Displays the classes or subclasses containing items that were purchased in the market baskets for a given product affinity. All market baskets that contributed to a rule will contain items from these product categories.
Then	Displays the class or subclass that was found to have an affinity with the IF components of the market basket rule.
Baskets	The number of market baskets that contain items from both the IF and THEN components of the rule.
Avg Sales from "If" Items	The average sales retail amount from items belonging to the IF classes or subclasses.
Avg Profit from "If" Items	The average sales profit from items belonging to the IF classes or subclasses.
Avg Sales Quantity from "If" Items	The average sales units from items belonging to the IF classes or subclasses.
Avg Sales from Affinity Items	The average sales retail amount from items belonging to the THEN class or subclass.
Avg Profit from Affinity Items	The average sales profit from items belonging to the THEN class or subclass.
Avg Sales Quantity from Affinity Items	The average sales units from items belonging to the THEN class or subclass.
Support	The percent of market baskets that contain items from the IF and THEN components of the rule.

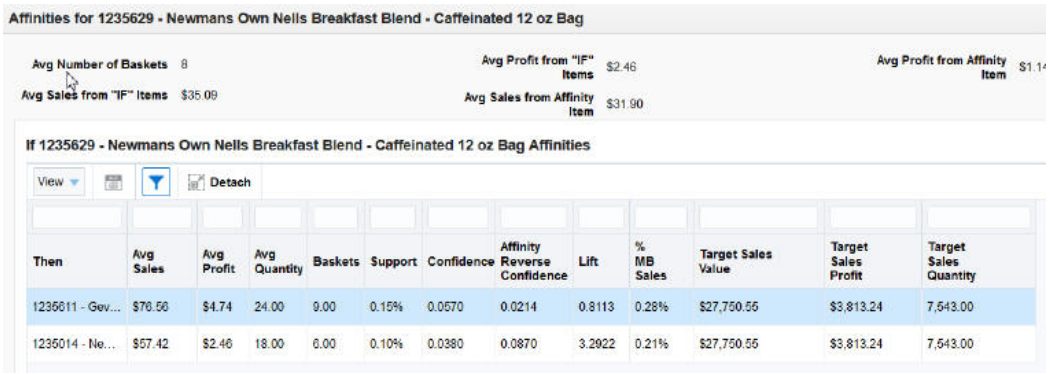
Table 8–18 (Cont.) Affinities Search Results

Field	Description
Confidence	The ratio of the number of transactions where the entire rule is present, compared to all transactions containing IF items.
Affinity Reverse Confidence	The ratio of the number of transactions where the entire rule is present, compared to all transactions containing THEN items.
Lift	A statistical measure of how strong the product affinity is, as compared to random chance for the rule (higher is better).
% MB Sales	The percent of total sales retail amount that the rule represents.
Avg Sales	The average weekly sales retail amount for the IF and THEN items in the rules.
Avg Profit	The average weekly sales profit for the IF and THEN items in the rules.
Avg Quantity	The average weekly sales units for the IF and THEN items in the rules.
Target Sales Value	The average weekly total sales retail amount across all transactions in the selected periods.
Target Sales Profit	The average weekly total sales profit across all transactions in the selected periods.
Target Sales Quantity	The average weekly total sales units across all transactions in the selected periods.
Total Basket Sales Volume	The average weekly transaction count across all transactions.

Results Pop-up

Click the right-hand icon for a specific row in the Results table to see If/Then Affinities for...pop-up. The pop-up shows all the other rules that have been discovered for the same If item. You can use this to compare the rules. In this way, you may discover another product with a stronger affinity for the item.

Figure 8–13 Affinities For



The pop-up has the following fields.

Table 8–19 Affinities For

Field	Description
Avg Number of Baskets	The average number of baskets across the runs for the combination of the If and Then items.

Table 8–19 (Cont.) Affinities For

Field	Description
Avg Sales from "If" Items	The average sales retail amount from items belonging to the IF classes or subclasses.
Avg Profit from "If" Items	The average sales profit from items belonging to the IF classes or subclasses.
Avg Sales from Affinity Item	The average sales retail amount from item belonging to the THEN class or subclass.
Avg Profit from Affinity Item	The average sales retail amount from item belonging to the THEN class or subclass.
Then	Displays the class or subclass that was found to have an affinity with the IF components of the market basket rule.
Avg Sales	The average weekly sales retail amount for the IF and THEN items in the rules.
Avg Profit	The average weekly sales profit for the IF and THEN items in the rules.
Avg Quantity	The average weekly sales units for the IF and THEN items in the rules.
Baskets	The number of market baskets that contain items from both the IF and THEN components of the rule.
Support	The percent of market baskets that contain items from the IF and THEN components of the rule.
Confidence	The ratio of the number of transactions where the entire rule is present, compared to all transactions containing IF items.
Affinity Reverse Confidence	The ratio of the number of transactions where the entire rule is present, compared to all transactions containing THEN items.
Lift	A statistical measure of how strong the product affinity is, as compared to random chance for the rule (higher is better).
% MB Sales	The percent of total sales retail amount that the rule represents.
Target Sales Value	The average weekly total sales retail amount across all transactions in the selected periods.
Target Sales Profit	The average weekly total sales profit across all transactions in the selected periods.
Target Sales Quantity	The average weekly total sales units across all transactions in the selected periods.

Panel

The side panel has two sections.To display the side panel, you must first select a row in the table. The top section reports some of the most useful values that are also displayed in the table. For details, see Table 8–13. The bottom section explains the details regarding the run you are reviewing.

Figure 8–14 Panel Top



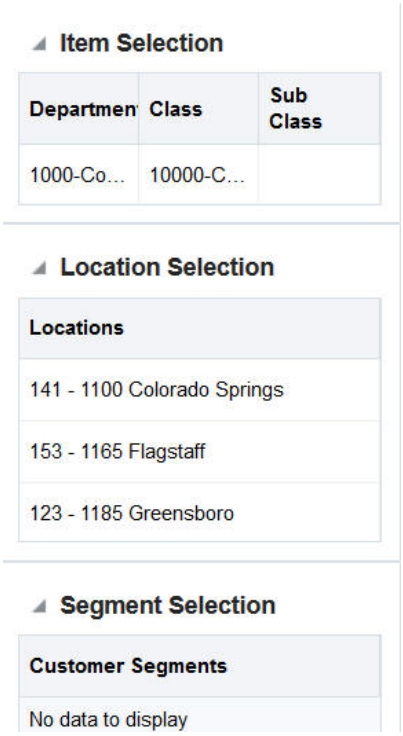
If 1234816 - Tully's French Roast Dar	
Support	0.13%
Confidence	0.0741
Reverse Confidence	0.0945
Lift	5.4780
Sales Amt Pct Of Total	0.45%
Avg Rule Sales Amt	$174.42
Avg Rule Sales Profit Amt	$0.00
Avg Rule Sales Qty	70.00
Avg Total Sales Amt	$38,342.72
Avg Tot Sales Profit Amt	$4,538.56
Avg Tot Sales Qty	11,858.00
Avg Tot Txn Cnt	9,392.00

Figure 8–15 Panel Middle



If 1234816 - Tully's French Roas	
Results Parameters	
Minimum Support	0.02%
Minimum Confidence	0.0000
Minimum Reverse Confidence	0.0000
Minimum Lift	0.0001
Minimum Transactions	2.00
Maximum Number of Rules	9,999.00
Maximum Rule Size	2.00
Date From	11/1/2015
Date To	1/31/2016

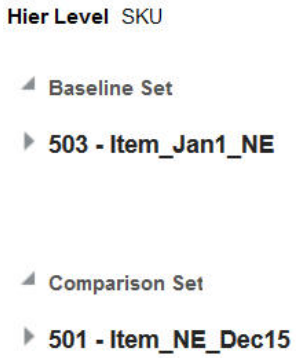
Figure 8–16 Panel Bottom



Compare Results

To compare results, select the two runs you want to compare. Highlight one run and click **Add to Baseline Set**. Highlight the second run and click **Add to Comparison Set**. You see these selections in the Compare Results right panel.

Figure 8–17 Panel Showing Run Selection



Click **View** to see the comparison. the View Results pop-up displays, which can be filtered by a combination of Department Item, Class, Top Results By, Subclass, and Top Results Displayed.

The results are displayed as a pair of rows for each item, a baseline row followed by a comparison row.

Table 8–20 Compare Results

Field	Description
If	Displays the classes or subclasses containing items that were purchased in the market baskets for a given product affinity. All market baskets that contributed to a rule will contain items from these product categories.
Then	Displays the class or subclass that was found to have an affinity with the IF components of the market basket rule.
Set	Indicates whether the row of data is for the baseline runs or the comparison runs.
Set Size	The number of Item/Then components in the affinity rule. For example, "If Coffee and Creamer Then Sugar" has a Set Size of 3.
Baskets	The number of market baskets that contain items from both the IF and THEN components of the rule.
Avg Sales from Affinity Items	The average sales retail amount from items belonging to the THEN class or subclass.
Avg Profit from Affinity Items	The average sales profit from items belonging to the THEN class or subclass.
Avg Sales Quantity from Affinity Items	The average sales units from items belonging to the THEN class or subclass.
Avg Sales from "If" Items	The average sales retail amount from items belonging to the IF classes or subclasses.
Avg Profit from "If" Items	The average sales profit from items belonging to the IF classes or subclasses.
Avg Sales Quantity from "If" Items	The average sales units from items belonging to the IF classes or subclasses.
Support	The percent of market baskets that contain items from the IF and THEN components of the rule.
Confidence	The ratio of the number of transactions where the entire rule is present, compared to all transactions containing IF items.
Affinity Reverse Confidence	The ratio of the number of transactions where the entire rule is present, compared to all transactions containing THEN items.
Lift	A statistical measure of how strong the product affinity is, as compared to random chance for the rule (higher is better).
% MB Sales	The percent of total sales retail amount that the rule represents.
Avg Sales	The average weekly sales retail amount for the IF and THEN items in the rules.
Avg Profit	The average weekly sales profit for the IF and THEN items in the rules.
Avg Quantity	The average weekly sales units for the IF and THEN items in the rules.
Target Sales Value	The average weekly total sales retail amount across all transactions in the selected periods.
Target Sales Profit	The average weekly total sales profit across all transactions in the selected periods.
Target Sales Quantity	The average weekly total sales units across all transactions in the selected periods.
Total Basket Sales Volume	The average weekly transaction count across all transactions.

Create Custom Run

To create a custom run, select **Create Custom Run** from the Action drop-down. You see the Create Custom Run pop-up with two pages - Item Selection and Results Parameters. Need more info about recommendations to create a run.

Use the Item Selection page to select the items to include in the results, including the level of the results and the hierarchy nodes to include.

Provide a meaningful name for the **Run Name** that can help identify the run.

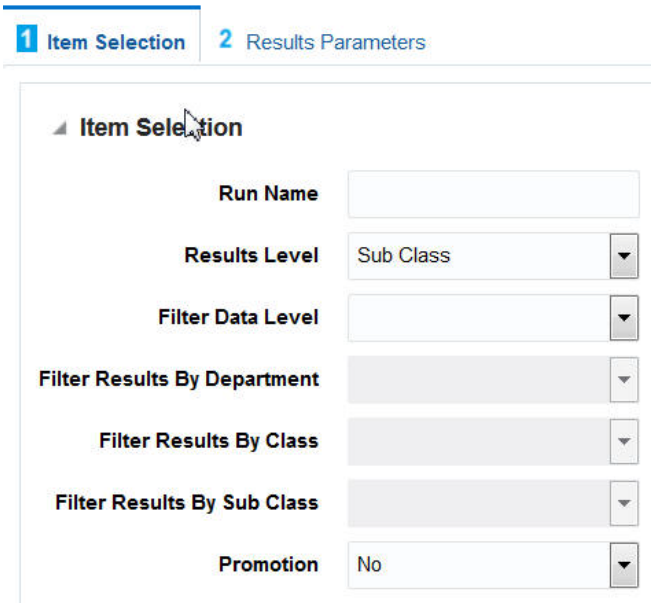
The **Results Level** defines the level of the hierarchy that the data is created for. Lower levels can required lower thresholds on the Results Parameters page (in order to find results), while higher levels can include too many results if the settings are not increased on the Results Parameters page.

The **Filter Data Level** is used to define rules regarding how to filter data. For example, if the results are only required for specific departments, then select Departments as the level and then select the departments.

A Promotion value of **Yes** only includes results in which the If component was on a promotion. A value of **No** is used when there is no concern about whether a promotion was present or not.

After entering item criteria, click **Add Item Selection**. The bottom section of the screen is populated with the results of the selection. You can also see this information in the right panel.

Figure 8–18 Item Selection Tab



The Results Parameters page, which is shown in Figure 8–19, is displayed in the right panel. The minimum value for transactions are used to prevent a situation in which there are too many results because so many transactions have been used in the calculations.

Figure 8–19 Results Parameters Tab

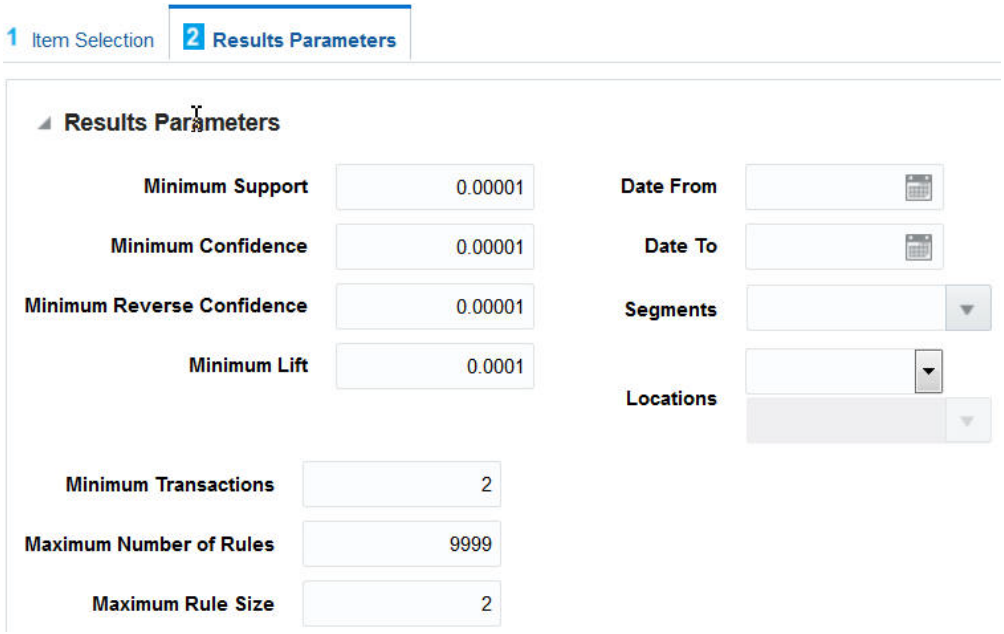


Table 8–21 explains how to interpret these values.

Table 8–21 Results Parameters

Field	Description
Minimum Support	The minimum percent of total transactions required to have this rule expressed as a decimal value. (0 - 1 range)
Minimum Confidence	The minimum confidence ratio required for a rule expressed as a decimal value. (0 - 1 range)
Minimum Reverse Confidence	The minimum confidence ratio required for the reverse rule definition, expressed as a decimal value. (0 - 1 range)
Minimum Lift	The minimum lift value required for a rule.
Minimum Transactions	The minimum value for transactions are used to prevent a situation in which there are too many results because too many transactions have been used in the calculations. This value also helps ensure some consistency if the transaction volume varies enough for Minimum Support to be used.
Maximum Number of Rules	The limit on the top number of rules to be saved from the results. Used to reduce data volume and remove the less significant rules.
Maximum Rule Size	The number of components to be included in a rule. The value of 2 is only supported for rules that include promotion or item results; otherwise, the value of 3 is supported. The additional components are always on the If side of the rule.
Dates From and To	The date range for processing sales data. This value must not span an excessive length of time, as the amount of data processed can become too large to manage reasonably.
Segments	Select any customer segments that should be included when determining the sales transactions to process.

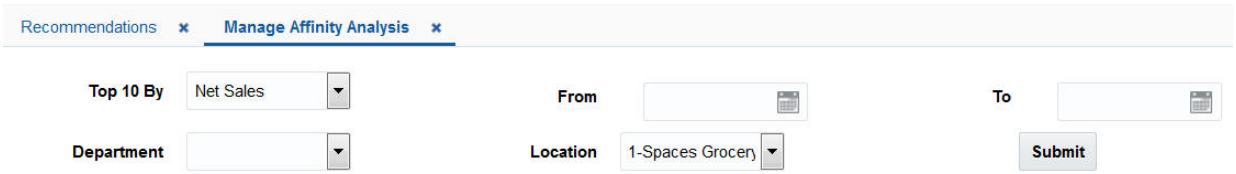
Table 8–21 (Cont.) Results Parameters

Field	Description
Location	Limit the locations that should be included when determining the sales transactions to process. This can be done by selecting any hierarchy level and then multi-selecting any locations at that level.

Manage Affinity Analysis

Use the Manage Affinity Analysis screen to view the results from all system created runs. Set the filters, select a range of calendar dates, and click **Submit**. The Info Tiles are displayed.

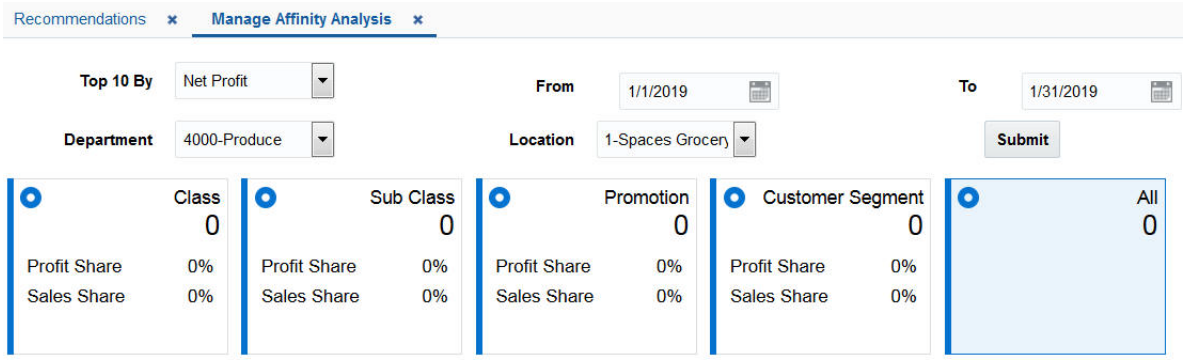
Figure 8–20 Manage Affinity Analysis Filters



Info Tiles

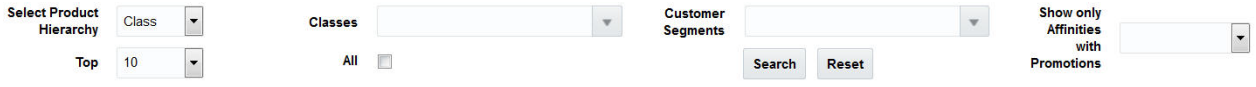
The Info Tiles display the results from all system runs. Once you specify the filters to indicate what runs you want to use, the Info tiles display an aggregate view of the results for different hierarchy levels and the types of processing that were executed during the run. Select an Info Tile and use the filters below the tiles to view the affinities for your selection.

Figure 8–21 Info Tiles



Each tile has specific filters. Here is an example of the filters for the All tile.

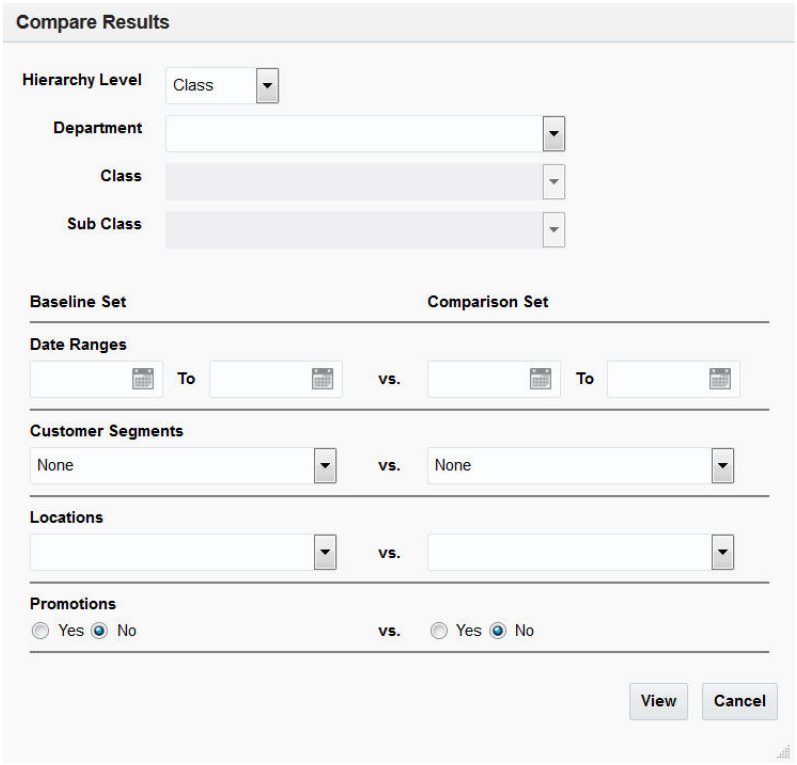
Figure 8–22 All Info Tile Filters



Compare Results

You can use Compare Results to compare system runs only, using the dates and filters you specify, as shown in Figure 8–23.

Figure 8–23 Compare Results



Click **View**. Table 8–20 describes the results of the Compare Results functionality.

9

Social Analytics

This chapter describes the Social Analytics Cloud Service module.

Introduction

Social Analytics (SA) is an enterprise-specific tool that delivers insights derived from patterns in social data and sales transactions data. It helps retailers, especially in fashion industry, to track online references to different keywords (such as a brand, color, style, and so on) and relate metrics such as sentiment or number of references to sales numbers for different categories of products. The joint analysis of social and sales data sheds light on what customers like and dislike about products, which products (or colors, or styles) may be trending, and the social metrics for the retailer as well as its competitors.

The SA cloud service consists of an Overview screen and a Social Analytics screen. The Overview screen displays the summary of insights, including trending keywords, and regional distribution and daily trends of social metrics. In the Social Analytics screen, you can visualize the daily and weekly trends in social and sales data for different keywords (such as brand, color, style, and so on) and different categories of product and look for patterns that indicate a link between social metrics and trends in sales.

Overview Screen

The Social Analytics application is one of the modules in the Decision Insight section. To access the Overview screen, log into the Retail Science Engine Cloud Services, click **Tasks**, and select **Decision Insight** from the menu.

The Overview screen consists of three cards. You can navigate through the cards using the small round buttons below the SA tile. The information displayed in each card is based on all active topics that you have defined in Oracle Social Relationship Management (SRM).

Besides displaying the summary of insights in the three cards, you can also navigate to the Social Analytics screen by clicking **Social Analytics** at the top of any of the cards.

Card 1

The first card shows the top three trending keywords and trending sources, based on all of your active topics in the SRM application. Trending keywords are those with the highest number of references, and trending sources are the media outlets and blogs that mention the keywords most frequently. You can find a longer list (top ten) of trending keywords and trending sources in the right panel.

Card 2

The second card displays the summary of activity and sentiments across all regions, that is, the regions with the highest and lowest activity, positive sentiment, and negative sentiment, based on all your active topics in the SRM application. Activity is the total number of conversations across all topics over a specific date range, and sentiment (which is either negative, neutral, or positive for each conversation), is determined based on the overall tonality score of the conversations.

You can display the activity and sentiments breakdown by region for the most recent week by hovering over the map in the right panel. The applicable date range (most recent week) is shown on the top of the panel. For each region, rate of sales and best/worst categories of the product are displayed. Rate of sales is defined as the percentage of change in sales across all categories of product in the most recent week (compared to the previous week). Best (worst) categories of product are defined as the categories that have highest (lowest) percentage of change in sales.

Card 3

The third card displays a summary of activity and sentiments, that is, the dates with the highest and lowest activity, positive sentiment, and negative sentiment, based on all your active topics in the SRM application, over the most recent week. You can display the daily activity and daily share of positive and negative sentiments in the right panel.

For a detailed explanation of social metrics (such as sentiment, mentions, and so on) see the Oracle SRM documentation.

Social Analytics Screen

To access the Social Analytics screen, click **Social Analytics** on the top of any of the cards in the Overview screen. The Social Analytics screen consists of one tab (Trend Analysis).

Trend Analysis Tab

You can use the Trend Analysis tab to visualize the social and sales trends and look for patterns that indicate a link between social metrics and trends in sales.

On the left side, there are three graphs that show the daily or weekly activity, sentiments, and sales trend. To switch between daily and weekly view, use the radio buttons on the top right of the screen.

The top graph shows the daily or weekly activity for the most recent eight weeks and for the topics(s) and keyword selected on the right side. For example, if you select the topic "competition" and select the keyword "dress", the graph will show number of daily or weekly conversations that match the filters defined for "competition" (in the SRM application) and also contain the keyword "dress".

The middle graph shows the daily or weekly percentage of positive and negative sentiments for the most recent eight weeks and for the topic(s) and keyword selected on the right side. To switch between positive and negative sentiments, use the radio buttons next to the graph.

The bottom graph shows the daily or weekly unit and amount of sales for each of the selected product categories on the right. Alternatively, you can display the aggregate sales over all applicable categories by checking the check box above the list of categories on the right. The applicable product categories displayed on the bottom

right are determined based on the name of sub-classes within the category (that is, by looking for categories where the name of at least one sub-class matches the selected keyword).

To switch between unit and amount in the bottom graph, use the radio buttons next to the graph.

To update the graphs after selecting product categories, click **Update Graphs** on the bottom left.

Figure 9–1 shows the Trend Analysis tab and the three steps for specifying topics, keyword, and product categories.

Figure 9–1 Trend Analysis



Based on the selected product categories (for sales data) and selected topics (for social data), the correlation table on the right shows the correlation numbers between unit or amount of sales and activity or positive/negative sentiments. Large correlation numbers (close to 1) indicate that trends in sales numbers and social metrics are correlated (that is, they tend to go up and down together). For example, a strong correlation between activity for a specific color and sales units for a specific product category can be an indication that there is growing demand for that particular color.

10

Offer Optimization

This chapter describes the Offer Optimization (OO) module. In general, this chapter describes both modules that are offered: Promotions and Markdowns Cloud Service (PMO) and Offer Optimization Cloud Service (OO). Throughout the document, for the ease of describing the components, the term "Offer Optimization" will be used to refer to these two combined services. However, when referring to a specific component that is only available in one particular service, the complete service name will be used.

Introduction

Offer Optimization is used to determine the optimal pricing recommendations for promotions, markdowns, or targeted offers. Promotions and markdowns are at the location level (for example, price-zone), and targeted offers consider the location and can be specific to each customer. The timing and depth of promotions, markdowns, or targeted offers are important levers for managing the inventory over the life cycle of the product.

The module can be useful in the following scenarios:

- Bring inventory to the desired level, not only during the full-price selling period but also during the clearance period.
- Maximize the total gross margin amount over the entire product life cycle.
- Assess in-season performance.
- Update recommendations each week. This facilitates decision-making based on recent data, including new sales, inventory, price levels, planned promotions. and other relevant data.
- Provide targeted price recommendations at the segment-level.

This section provides information to help you understand the functionality of Offer Optimization.

Administering Offer Optimization

For information about the administration of Offer Optimization, see *Oracle Retail Science Cloud Services Administration Guide*.

Users and Roles

The following four user login roles are available and can be assigned at the time of user creation:

- Pricing Analyst–main business user responsible for day-to-day pricing optimization activities. This user is responsible for working through price recommendations.
- Pricing Manager (or Analytical Super User)–responsible for analytical configuration, testing, and model diagnosis.
- Targeted Offer User–this user is responsible for working through the targeted price recommendations.
- Buyer–user who is responsible for a department or departments, who approves or rejects the OO run, and who is responsible for the translation of data between OO and Oracle Retail Price Management (RPM) and Oracle Retail Customer Engagement (CE).
- Pricing Administrator–responsible for the general system setup and configuration tasks related to the business.

Table 10–1 provides a description of the privileges. Subsequent tables describe privileges that are allowed based on the user role and the run status.

Table 10–1 Privileges

Privilege	Description
Create new OO run	Set up and execute optimizations. All steps on the main train can be carried out, including creating a new run via the copy run functionality.
Modify the name and description of an existing OO run	Open an existing run and modify the name and description of that run.
Modify the General Rule of an existing OO run	Ability to modify the General Rule for that run.
Modify the Custom Rule of an existing OO run	Ability to modify or specify the Custom Rule for that run.
View existing OO run	Open an existing run and visit any step in the main train. The user can view data and search/aggregate/filter but cannot change anything.
Optimize an OO run	Optimize a new or modified run or a failed run.
Submit OO runs	Approve run results.
Modify Promotion or Markdown Recommendations	Ability to accept, reject, or override promotion or markdown price recommendations for that run.
Modify Targeted Recommendations	Ability to accept, reject, or override targeted recommendations for that run.
Delete saved runs	Delete run setup information and results from the database.

Table 10–2 User Roles and Permissions

Privilege	Data Access	Analyst	Buyer	Administrator	Analytical Superuser	Targeted Offer User
Create new OO run	All runs	X			X	
View existing OO run	All runs	X			X	
Modify the name and description of an existing OO run.	Runs they created	X	X	X	X	

Table 10–2 (Cont.) User Roles and Permissions

Privilege	Data Access	Analyst	Buyer	Administrator	Analytical Superuser	Targeted Offer User
Modify the General Rules of an existing OO run	Runs they created	X			X	
	Batch runs				X	
Modify the Custom Rules of an existing OO run	Runs they created or batch runs				X	
Optimize or re-optimize OO run	Runs they created	X			X	
	All runs with technical failures			X		
Modify Promotion or Markdown Recommendation	Runs they created or batch runs	X			X	
Modify Targeted Recommendations	Runs they created or batch runs					X
Recalculate OO run	Runs they created or batch runs	X				
Review OO runs	Runs they created or batch runs	X			X	
Submit OO runs	All runs		X			
Delete saved OO runs	Runs they created	X			X	
	All runs			X		

Table 10–3 Run Status Constraints

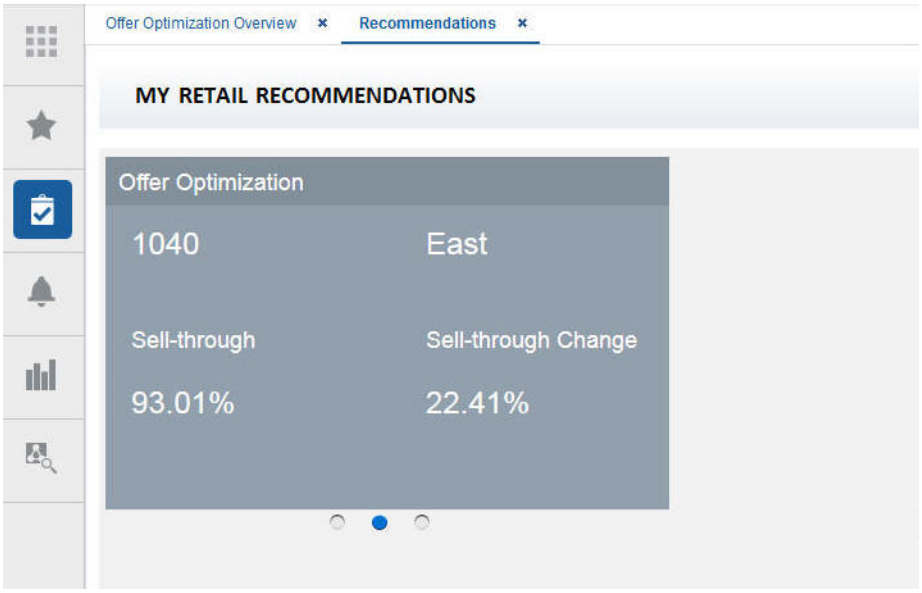
Role	Privilege	Object	Status
Users	cannot modify existing OO run	when run is	running or submitted
Users	cannot optimize OO runs	when run is	running, submitted, or not autorun
Users	cannot optimize OO runs	unless run is	reviewed
Users	cannot delete saved OO runs	when run is	running or submitted
Administrator	cannot delete saved OO runs	when run is	submitted
Users	cannot optimize batch runs	when run is	any status

Overview of the User Interface

The OO user interface, at a high level, consists of the following:

- My Retail Recommendations Tab. When you log into the application, this tab provides a quick view of the latest run that has completed optimization. If you subscribe to other ORASE applications, then the corresponding tiles for each application are shown. The Offer Optimization tile consists of a link to the Offer Optimization Overview screen and the last completed run (that is, the run that is in Ready for Review status) information. This tile provides three cards that show the absolute numbers and percentage change compared to the current price policy for revenue, sell-through, and gross margin. The user can click Offer Optimization to navigate to the Overview tab. Alternatively, the user can click the Run ID or Run Name to open the run tab, which opens the Results screen.

Figure 10–1 My Retail Recommendations Tab



- Task Pane. Located on the left side of the application. Use this to navigate between the applications.
- Offer Optimization Overview Tab. You can access run management in this screen.
- Offer Optimization Run Tab. The main tab for application work such as setting up the business rules, optimizing a run, or viewing and editing the results. There are three main steps in this tab: Setup, Rules, and Results.
- Contextual BI information. This tab is located on the right side of the application. It provides information in graphical format that can be used in setting up runs and interpret the optimization results.
- Oracle Digital Assistant (or Chatbot). This tab is located on the right side of the application. This functionality provides the user with voice or conversation-based assistance. The user can ask the chatbot queries (see "Integration with Oracle Digital Assistant") and Offer Optimization provides answers.

Icons

The following icons are used in user interface. Certain icons have slightly different definitions, depending on the context.

Table 10–4 Icons

Icon	Description
	Copy a run.
	Create. Create a row. Create a run.
	Detach
	Delete. The object deleted depends on the context.
	Edit. Change status.

Table 10–4 (Cont.) Icons

Icon	Description
	Embedded help.
	Export to Excel. Export all.
	Freeze.
	Go to top.
	Go up.
	Go to bottom.
	Go down.
	Move selected item to another list.
	Progress indicator. The indicator changes mode when the application is processing data.
	Query by example.
	Refresh.
	Remove all items form list.
	Remove selected items from list.
	Required.
	Select date.
	Wrap.

Buttons

Buttons are used to perform certain actions and for navigation.

Table 10–5 Buttons

Button	Description
Action	Provides access to the following actions: Save, Optimize, Re-Optimize, Reviewed, Save, Recalculate, and Submit.
Back	Used to return to the previous train stop.
Cancel	Used to close a dialog box without making a selection.
Clear Selection	Clears the selected rows list.
Next	Moves to the next train stop.

Table 10–5 (Cont.) Buttons

Button	Description
Save	Saves the existing settings.
Select All	In Business Rules - Product Groups, selects the entire set of rows displayed.
Accept All	In Results and Analysis, accepts all the price recommendations.

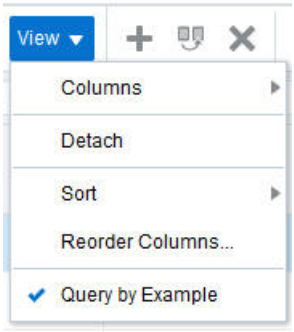
Pull-Down Menus

The user interface has three pull-down menus that provide access to a variety of functionality.

The View pull-down menu can be used to adjust how the display is organized. The options provided are:

- Columns provide the ability to manage columns. You can choose to hide certain columns from display. Or you can re-order how the columns are shown.
- Detach a table from the interface. This maximizes the table to fit your browser.
- Sort provides the ability to specify up to three columns to sort the data in ascending or descending order for each column.
- Reorder Columns provides the ability to specify the order for the columns displayed.
- Query by Example can be turned off or on by unchecking or checking. This is the top row of empty text boxes that is shown over any table by default. The user can filter the display based on specified values for each column.

Figure 10–2 Query by Example



The Actions pull-down menu provides functionality that you use to manipulate the application data. Some of the following functionality is also provided by the icons described in Table 10–4.

Figure 10–3 Actions Pull-Down Menu

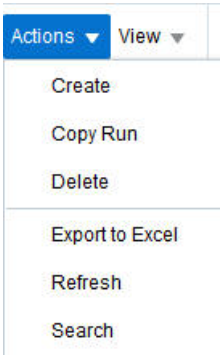


Table 10–6 Actions Pull-Down Menu Functionality

Action	Description
Create Run	In Overview, takes user to Optimization Setup train stop.
Copy Run	In Overview, used to create a duplicate of an existing run.
Delete Run	In Overview, used to delete a run.
Export to Excel	Used to export the set of rows displayed to Excel.
Add	In Business Rules–Product Groups, used to add data.
Edit	In Business Rules–Product Groups, used to edit data.
Delete	In Business Rules–Product Groups, used to delete data.
Create	In Fixture and Product Data, used to create.

Charts

Certain stages have associated charts available on the right-hand side of the display that list data in a tabular format.

Process Train

The process train displays the stages of Offer Optimization. The current stage is highlighted. You can also use the Back button and the Next button to move through the train.

Embedded Help

Embedded help, which you access by clicking the Question Mark icon, provides additional information about the details required by certain fields.

Process Indicator

At the top of the user interface, the red horizontal line is a process indicator (that goes from left to right) that you can use to monitor the status of a user action such as clicking Next to go to the next stage.

Search

In certain cases, you can customize your search, using advanced search capabilities to specify the search criteria.

Workflow

OO is used to optimize a buyer's business strategy for a department by generating optimal price recommendations. The buyer has a business strategy and wants to determine when to offer a promotion, the depth of the promotion, which items to be considered as part of a promotion, and similar questions regarding markdown, keeping in mind the business rules and inventory levels. The application runs through the different price combinations, considering the gain when offering a promotion or markdown and whether to offer any customer-specific promotion (if the user subscribed to targeted offers) in an efficient manner that best achieves the optimization objective. Once the optimization is complete, the run goes into a Ready for Review status. An analyst can review the recommendations at each item-level and decide to accept, reject, or override any recommendation. Once the analyst finishes the review, the run status is changed to Reviewed. A reviewed run goes to the user with the buyer role for approval. If the buyer likes the recommendations then the buyer can change run status to Submitted. Here, Submit means that the price recommendations are sent to a price execution system such as CE or RPM. OO also provides the user with the ability to view and work on batch runs that are executed at regular interval (for example, week). Users can view and modify the recommendations and send them for approval to the buyer once the review is complete.

Optimization

An optimization can be carried out at the configured processing (or run) location or a price-zone, merchandise level, and calendar level. The user has option of configuring the use of either price-zone or a node in the location hierarchy (but not both in the same instance). Once the optimization is complete, the recommendations can be generated at a lower level than the processing level (called recommendation levels for merchandise level and location or price-zone level). The location and merchandise level can be any level in the location hierarchy and merchandise hierarchy, respectively. Alternatively, price-zone can be used to define a set of stores (and/or online locations) and items. An example of a price zone is women's apparel in all university-based stores, grouped into one price-zone. The usual levels for the run are Region or Price-zone, Department, and Week, and, for the recommendation, the levels are Region or Price-zone and Style/Color. If Targeted Offers is available, then the recommendations will also be generated at the Customer-Segment level, along with location or price-zone and merchandise level. The Promotions and Markdowns are always at the location or price-zone and merchandise level.

The optimization is set up as follows:

- Optimization is at done at one level below the run's merchandise level. For example, if the run merchandise level is Department, then each job is at the Class level.
- Inventory is rolled to the desired recommendation merchandise level. Further, the inventory is aggregated across all the locations to the run's location level.
- Price recommendations are generated at the configured recommendation level for merchandise, location or price zone, and calendar level and the customer segment level.

For example., if the run location level is Price-zone, the run merchandise level is Department, the run calendar level is Week, and the recommendation level for merchandise is Style/Color, then the recommendations are generated at the Price-zone, Style/Color, Week, and Segment (when TO) levels.

Offer Optimization Overview

The Offer Optimization Overview is the dashboard for the OO runs. In this tab, you can see a list of all existing runs, along with details that describe each run. The list includes runs created by other users, which you can open in read-only mode. You can create a run, copy a run, open a run, or delete a run.

This screen has three components, Search, Tiles, and the Table of Runs, as shown in Figure 10–4.

The search panel is collapsed by default. You can bring up the Search panel by clicking Search under the Actions Menu. The search panel lets you filter the runs to be seen on a regular basis. It filters the run based on Department, Location, Created By, and Created Date. You can filter the runs to satisfy all the above criteria or any subset of the above criteria.

The Tiles dashboard provides quick access to information on run status. It shows the runs that are in the following status: Submitted, Reviewed, Ready for Review, Worksheet, Failed, and All. The first three tiles show how many runs are in that status, as well as a pie-chart displaying the price recommendations that have been accepted, rejected, and overridden over all the runs with that status.

The table of runs includes the columns described in Table 10–7. This provides relevant information on all the runs in a tabular format; you can perform operations such as create, copy, and delete runs in this table. For each run with a successfully completed optimization, certain KPIs such as Revenue, Gross Margin, and Sell-Through are shown in absolute numbers as well as percent gain. Note that each run can be in its own currency.

This table also provides the information about the batch runs. Batch runs are prefixed with a "Batch" in the run name, and the run status will always say "Auto run". You cannot send the batch runs to optimization. Rather, the batch runs are run at a regular frequency specified by the batch process. For example., each batch run extracts the latest sales data, updates the relevant inputs, and sends the run for optimization. If the batch run is successfully executed, the results and analysis will be shown. If not, it will not show any results and analysis. Note that the run statuses are not automatically updated; you must click Refresh to update the table.

Figure 10–4 Offer Optimization Overview

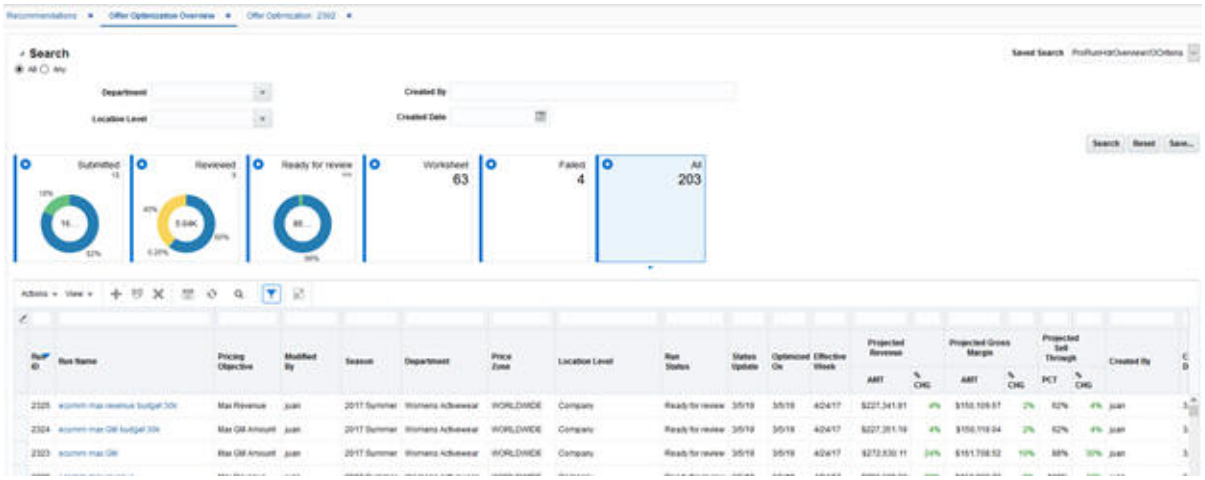


Table 10–7 Optimization List

Field	Description
Run ID	System-generated unique identifier for the run.
Run Name	The Run Name is the user-provided name.
Pricing Objective	A description of the objective used for the run.
Season	A description of the Season.
Department	A description of the Department.
Price Zone	A description of the price zone used for the run.
Location	The location selected for optimization.
Run Status	The current status of the run. Values include Worksheet, Running, Ready for Review, Reviewed, Submitted, Auto Run, and Failed. See Table 10–8 for details about each status.
Status Update	The date when the status was last updated.
Created Date	The date when the run was first created.
Optimized On	The date when the run was optimized.
Effective Week	The week ending date for which the run was executed.
Created By	Login ID of the user who created the run
Projected Revenue	Projected Revenue (AMT) for the run is the revenue generated by optimization in the run's currency, and Projected Revenue (%CHG) is the percent revenue gain compared to the current price policy.
Projected Gross Margin	Projected Gross Margin (AMT) for the run is the gross margin generated by optimization in the run's currency and Projected Gross Margin (&CHG) is the percent gross margin gain compared to the current price policy.
Projected Sell Through	Projected Sell Through (PCT) for the run is the sell-through generated by optimization, and Projected Sell Through (%CHG) is the percent gain in sell-through over the current price policy.

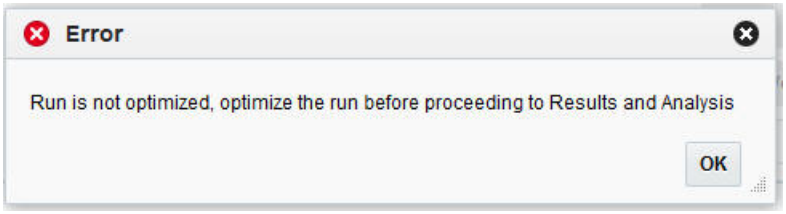
Table 10–8 Run Status

Status	Description
Worksheet	User is setting up the run, which has not yet been sent for optimization.
Running	The run has been sent for execution, and the process of collecting data and optimization is currently running. When the user selects re-calculate, the run enters into running status as well, but here it will be recalculating all the metrics based on the user price overrides (no re-optimization).
Ready for Review	The optimization process is complete, and the run results are ready for review.
Reviewed	The analyst finished reviewing the price recommendations and changes the run status to Reviewed so that it goes to the Buyer for approvals.
Submitted	The run results have been reviewed and approved for execution. Run results are ready for export to other applications.
Auto run	This status indicates that the run is a batch run and will be executed as part of batch process.
Failed	Run execution was aborted due to technical or data errors.

When the run is in Ready for Review status, clicking on the run name in the overview screen will automatically take you to the Results stage.

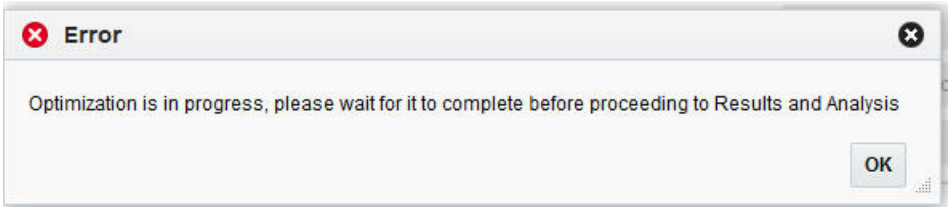
When the run is in Worksheet status, you cannot go to Results and Analysis stage since the optimization is not complete. You cannot proceed to the last stage and will see the error message shown in Figure 10–5.

Figure 10–5 Run Not Optimized



When the run is in Running status, you cannot go to Results nor make any more edits. The UI is locked from further edits until the optimization is completed. If you try to go to Results, the error message shown in Figure 10–6 is displayed.

Figure 10–6 Running Status



When the run is in Submitted status, you cannot accept, reject, or override any price recommendations. The UI is locked from making any further changes or edits.

Create a Run

To create a new run, click the Create Run icon. You are taken to the Setup stage in the Run tab.

In order to create a run, you must have either OO analyst permission or OO super user permission.

Copy a Run

To copy an existing run, highlight that run in the displayed list and click the Copy Run icon. You are taken to the Setup stage in the Run tab. The fields are populated with copies of all inputs and data. Any existing results are not included. The list of runs is updated with a run that has a new ID. The Run Name is prefixed with "Copy Of." All other fields match the existing run.

Open a Run

To edit an existing run, highlight that run in the displayed list and click the run name. Any user can open a run; however, edit mode is only allowed for runs that a user created with appropriate roles and only when the run does not have a status of Running, Reviewed, or Submitted. All other runs are opened in read-only mode.

Delete a Run

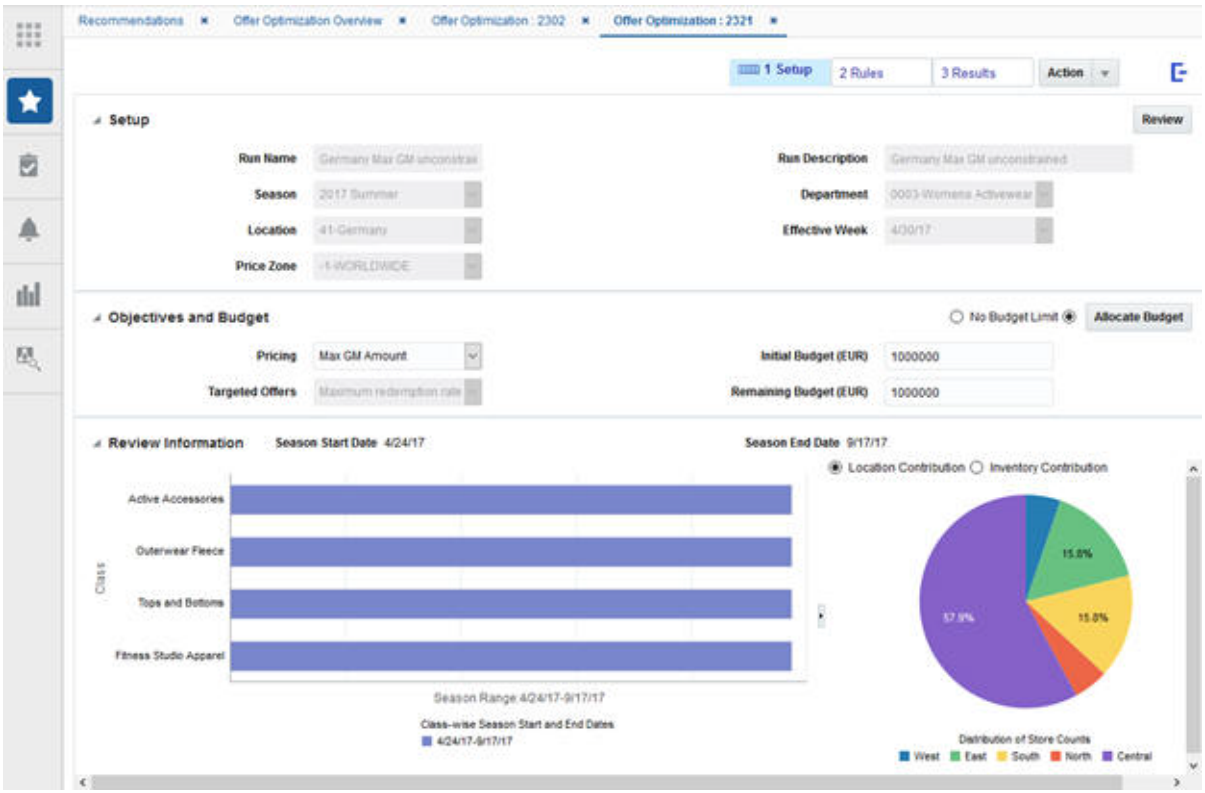
Runs can be deleted, depending on the status of the run and the permissions a specific user has. A user with Administrator permissions can delete any run except for one that has a status of Running or Submitted. Users can delete any runs they have created that do not have a status of Running, Reviewed, Submitted, or Auto run. They cannot delete runs created by another user. More than one run can be selected for deletion at the same time.

Offer Optimization Run Tab

Every run is opened in an Offer Optimization Run tab, as shown in Figure 10–7. The title of the tab is Offer Optimization: <Run ID>. The Run tab opens whenever you create a new run or click an existing run. This is the main tab where you can specify the business rules and goals for the run. It provides a series of three stages that you progress through in order to set up, run, and analyze the results of the optimization run.

- Setup. Used to select a season, location, and department. It is also used to select the objectives and specify the budgets.
- Rules. Used to view or change business rules.
- Results. Used to view results, and override, approve, or revisit prior steps in order to make changes.

Figure 10–7 Offer Optimization Run Tab



Setup

The Setup stage is used to pick a Season, Department, and Location or Prize Zone. It is organized into three major panels: the Setup section, the Objectives and Budget section, and the Review Information section.

This tab is read-only if the run has been submitted to the system, if the user does not have permission to edit runs.

You can go to the next stage only after the run is saved. Once you click Save, you can go to the next stage. If you want to send the run for optimization, click Optimize (or Re-optimize) the run for optimization. Use the Action menu to access Optimize (Re-optimize) and Save. If you send the run for optimization at this point, the default business rules are used.

Setup

The Setup section, shown in Figure 10–8, is used to select the Season, Department, Location or Price Zone, and Effective Week for the run. Click Save to save changes to the Name and Description fields. You can also use this button to save changes made in the Objectives and Budget section. You do not need to click save if you simply want to review the information for the selected season, location, department combination. You can click Review, which populates the bottom Review Information section. This is helpful if you are interested in understanding how much inventory is remaining for the selections in the Setup. However, once the run is saved, the review panel is automatically populated.

Figure 10–8 Setup

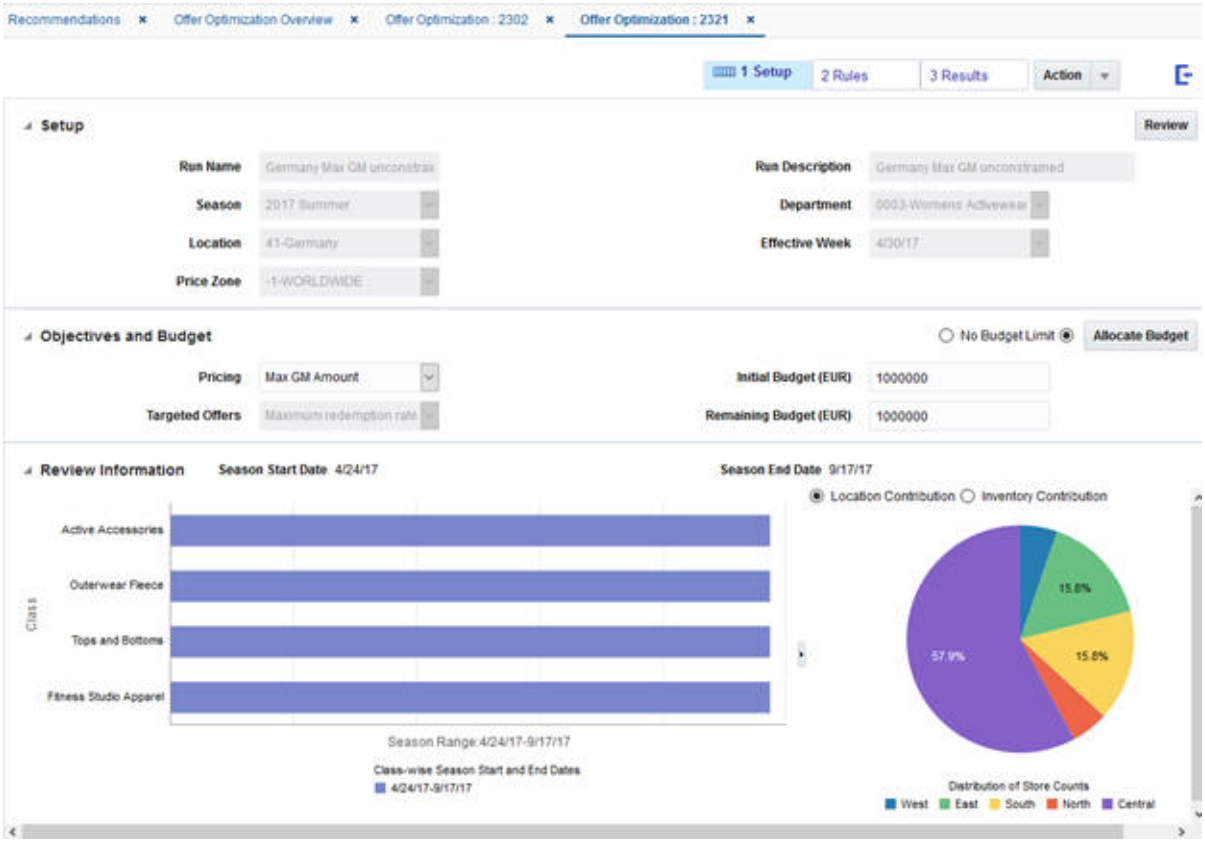


Table 10–9 Setup

Field	Description
Name	The name identifying the run. The name must be 80 characters or less and use alphanumeric characters only. It is case sensitive.
Description	A brief description of the run. The value must be 200 characters or less and use alphanumeric characters only.
Season	You can choose from the drop-down list of seasons available for the run.
Department	You can choose from the drop-down list of departments available for the run.
Location	You can choose from the drop-down list of locations available for the run.
Effective Week	You can choose from the drop-down list of effective weeks available for the run. Effective weeks in the past compared to the last pull of sales data (whenever the batch process was successfully executed) are not displayed.

Objectives and Budget

This section, shown in Figure 10–8, can be used to specify the objective for Pricing and Targeted Offers problems. You can specify whether there is a Budget Limit or No Budget Limit for the Pricing problem. Four objectives are available for the user: Max Revenue, Max GM Amount, Max Revenue with Salvage, or Max GM Amount with Salvage for the pricing problem. Targeted Offers support only one objective, maximize redemption rate. You must click Save under the Action button to save any changes made in this section.

Table 10–10 Objectives and Budget

Field	Description
Pricing Objective	You can choose between Max Revenue (default), Max GM Amount, Max Revenue with Salvage, and Max GM Amount with Salvage from the drop-down set of values.
Targeted Offers Objective	Supports only Maximize Redemption Rate.
No Budget Limit	Indicates no budget limit for promotions and markdowns for the run.
Initial Budget	The budget available at the beginning of the season for this department and location.
Remaining Budget	The remaining budget available at some point in the season for this department and location. For batch runs, this value is updated based on the promotions and markdowns taken so far. For a new run, this value is automatically populated with the value entered for the Initial Budget.
Allocate Budget	Use the radio button before this button to toggle to put a limit on the budget for the run. This button becomes available only when you specify a non-zero value for the Remaining Budget. When you click this button, it brings up a pop-up that you can use to specify Class-level budget limit for promotions and markdowns.

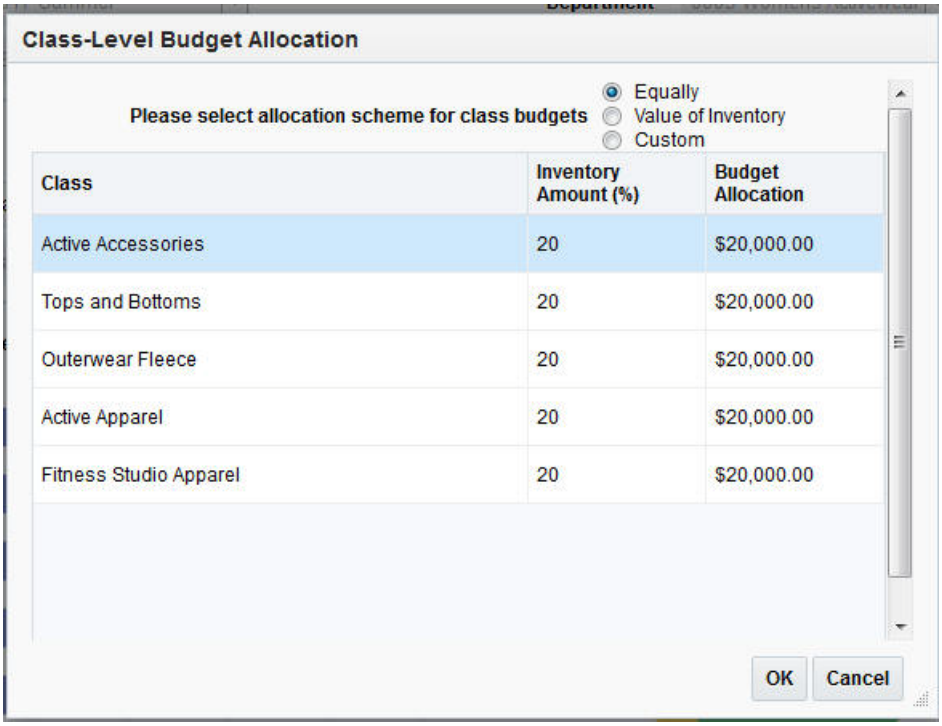
The Allocate Budget button is enabled only when you enter a non-zero value in the Remaining Budget field and save the run. If you enter a non-zero value but do not save, then selecting Allocate Budget will display the error message shown in Figure 10–9.

Figure 10–9 Allocate Budget Error



As shown in Figure 10–10, you can click Allocate Budget to bring up an utility that you can use to allocate the Department-level budget to classes based on three options. Equally distributes the budget equally among all classes. Value of Inventory allocates budget to each class in proportion to the value (current price * inventory) of that class. Classes with higher inventory and higher valued items receive more budget limitations. Custom allows you to specify the percent for each class. When you select this radio button, the default percentages for each class shown is based on the value of the inventory.

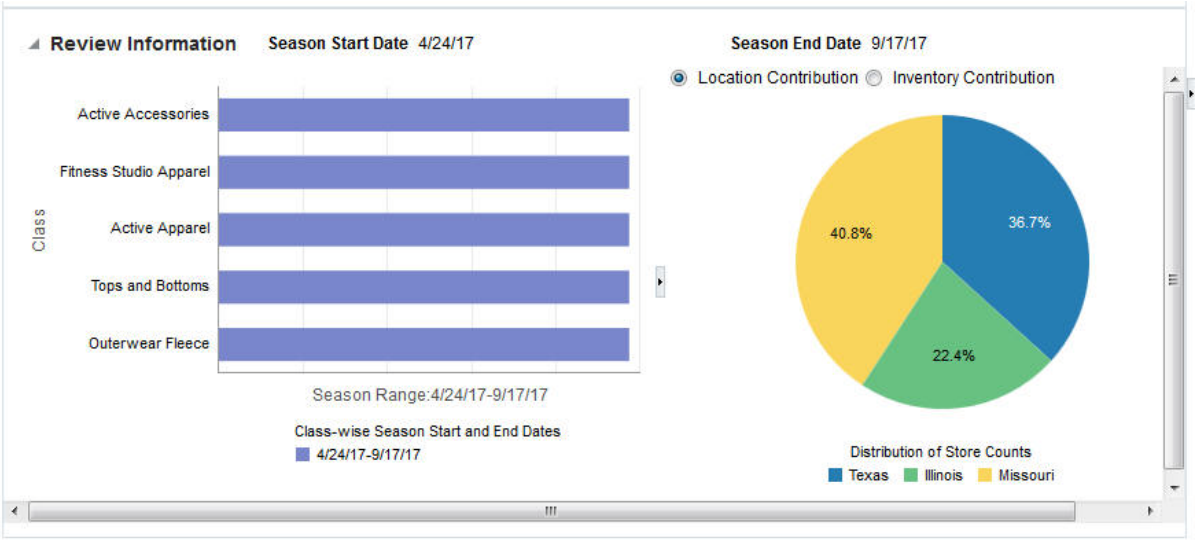
Figure 10–10 Allocation Scheme



Review Information

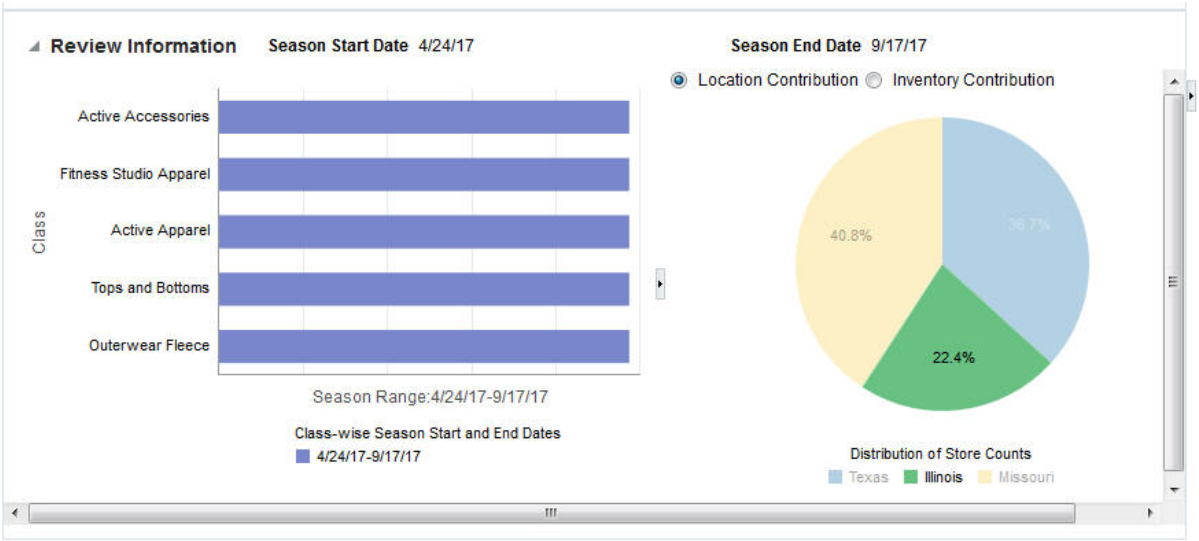
This section provides information on the start and end dates by class, a season's start and end date, and pie-charts of store counts by the immediate child location of the location selected in the setup. For example, if a region such as US Central was selected in the setup, then this displays the contribution by District (Texas, Illinois, Missouri). You can toggle between the Location Contribution pie-chart and the Inventory Contribution pie-chart. The inventory contribution shows the distribution of unsold inventory by the child location. This information is useful as it provides the status of unsold inventory at different locations.

Figure 10–11 Review Allocation Information



When you hover over a component in the chart, the corresponding legend item becomes highlighted. You can also hover over a legend item to see the corresponding chart component highlighted. For example, in Figure 10–12, Illinois is highlighted.

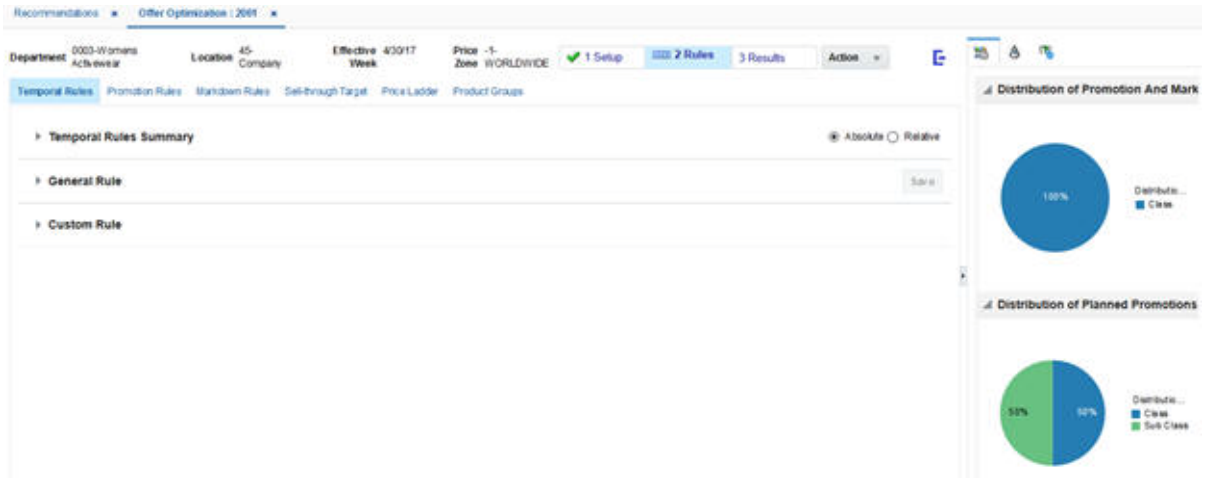
Figure 10–12 Highlight Chart Component



Rules

Business Rules correspond to the different constraints that are imposed on a pricing optimization problem. The Rules stage is divided into five different tabs: Temporal Rules, Promotion Rules, Markdown Rules, Sell-Through Target, Price Ladder and Product Groups. As shown in Figure 10–13, each rule screen is generally divided into three panels: Summary at the top, General Rule in the middle, and Custom Rule at the bottom. On the right, in the contextual area, the pie chart shows the percentage of rules specified at different merchandise levels.

Figure 10–13 Rules



The Summary panel displays the summary of all rules specified at different merchandise levels. By default, all rules are specified at the Class-level. When the run is saved (or optimized), the default values are used to populate at the Class-level. This panel is read-only and not editable.

In the General Rule section, you can specify the constraint at the run-level, which means that the rule is applied to all locations and all classes in the run. All classes obtain the same values mentioned in the General Rule section. The Save button is disabled if you have not made any changes in the related portion of the screen.

Custom Rule is divided into Custom Rule Selector and specific UI-related to the rule. This section can be used to modify the rule at a specific merchandise level. In the Custom Rule Selector portion, you select the merchandise level to apply the constraint.

- The set of available merchandise levels shown in the drop-down list depends on the selected department in the Optimization Setup screen.
- The merchandise level that can be selected is up to three levels, based on Class, Sub-class, or Style. The UI displays all the child nodes for the selected merchandise level. For example, if you select only a Class, then the child nodes that will be displayed are the subclasses for that class.
- Based on the selected levels, the bottom half of the screen is dynamically populated. Custom Rules are always specified at subclass or lower levels.
- The save button is disabled if you have not made any changes in the related portion of the screen.

When the rules are sent to optimization, the rules are applied based on the concept that lower-level rules override higher-level rules. That is, the rules screens allow you to manage rules by exception. Consider the following example. Suppose the General Rule says that the first min discount percent cannot be lower than 10 percent. This means that all items in all Classes are assigned this 10 percent criterion. If you decide to add a Custom Rule for a specific Style, say, 20 percent, then all the items in all Classes are assigned the 10 percent rule and the items in this particular Style are assigned the 20 percent rule.

Temporal Rules

This screen, shown in Figure 10–14, is used to specify the business rules associated with temporal aspects of the pricing problem. There are two ways you can specify the temporal rules. The first approach is called "Absolute Calendar" and the second

approach is called "Relative Calendar". The key difference between the two approaches is that the retailer uses actual dates to specify the rules in the former and in the later uses relative weeks from start (that is, from when the item started selling based on certain criteria).

By default, Absolute Calendar is used to display the rules screens. You can make a change to the relative calendar only on the Temporal Rules screen. You can switch and use Relative Calendar only once, and the functionality is further disabled for that run. A pop-up alerts you that "This selection is final and cannot be reversed for this run."

When you specify using the Absolute Calendar (default option selected):

- Default values at the class level are populated based on the interface.
- The interface allows you to specify the promo, markdown, or no-touch absolute calendar at the class level. That is, for each period, the kind of pricing change that is allowed is specified through this interface.
- Note that you cannot edit these class-level defaults in the UI. You can only override the class-level defaults by specifying the rules at lower levels of merchandise.
- You can only modify no-touch at the beginning and at the end of life for all classes in the UI.

When you specify using the Relative Calendar:

- Default rules at the class level are specified based on the default length of the promotion (60 percent) and default length of the markdown (40 percent).
- You can edit the class-level defaults (for all classes) by modifying in the General Rule panel.
- This calendar option is useful when the retailer introduces merchandise frequently into the store and a fixed or absolute calendar cannot be applied to all items that are in different phase (for example, regular vs. markdown) of the life.

The Absolute Calendar is useful when most of the items are introduced and exit at almost the same time. The Relative Calendar is useful when the retailer introduces items at frequent intervals and having a fixed calendar for all items is restrictive and pushes items into promotion or markdown too quickly within the items' life.

Figure 10–14 Temporal Rules

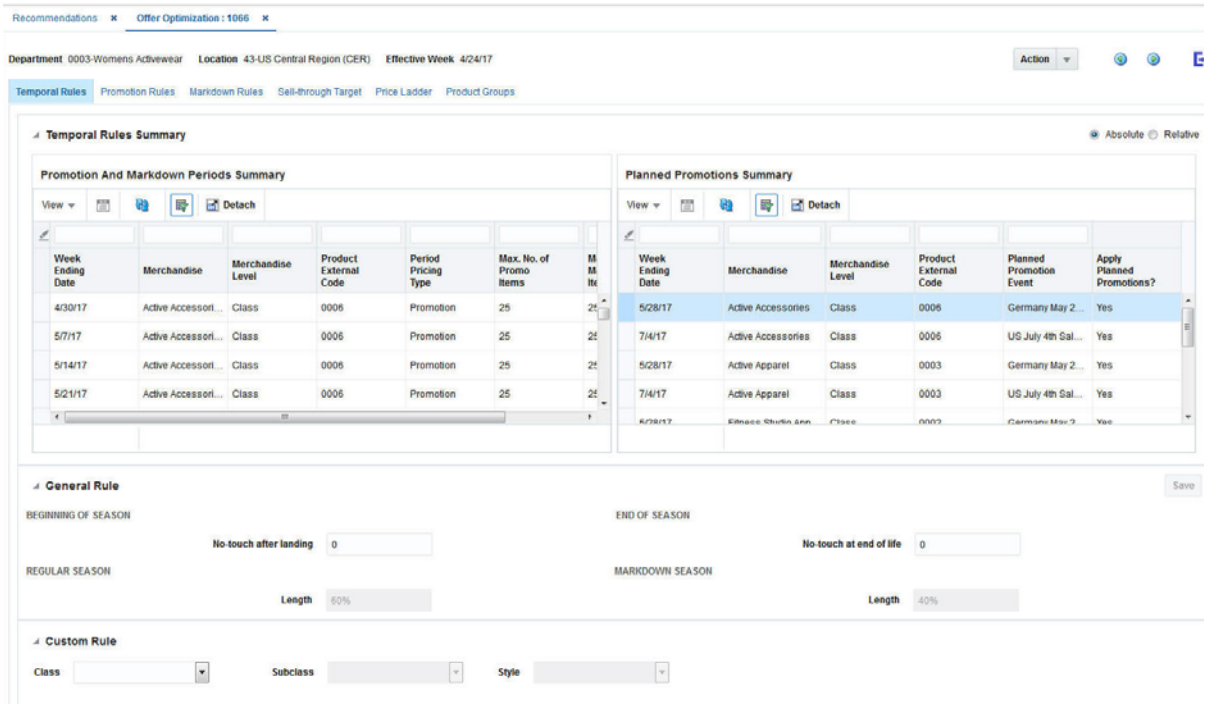


Figure 10–15 Custom Rules

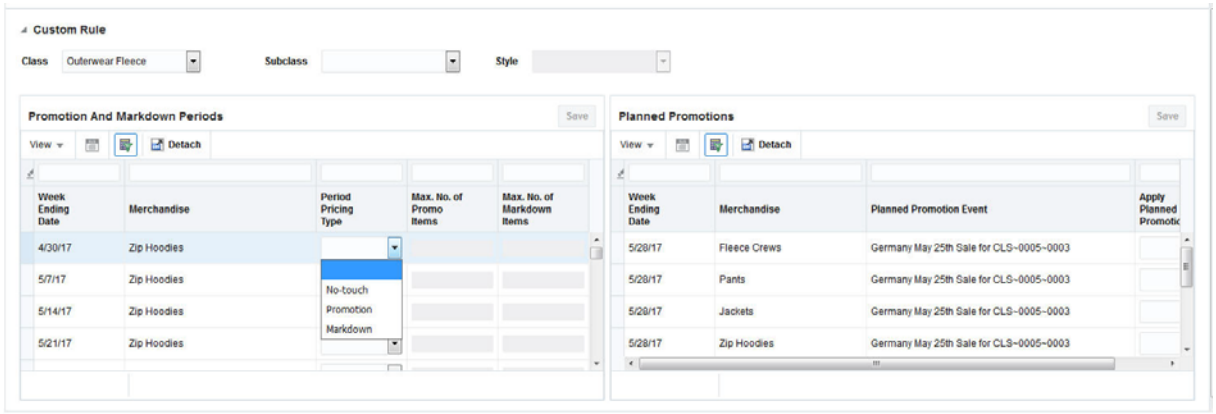


Table 10–11 Temporal Rules

Field	Description
Week Ending Date	The week ending date for the period. Not editable field.
Weeks From Start	This field is non-editable and runs from Week 1 to Maximum no. of weeks available for the selected merchandise level. Depending on the no-touch values, the Weeks from Start are adjusted to begin at 1+No touch after Landing and end at N-No touch at end of life. Available with only Relative Calendar option.
Merchandise	Description of the merchandise.
Merchandise Level	Merchandise level of the rule. It can be Class, Subclass, or Style/Color.

Table 10–11 (Cont.) Temporal Rules

Field	Description
Product External Code	Merchandise level's external code that can be identified by the retailer.
Period Pricing Type	Indicates whether the period is marked as Promotion Allowed or Markdown Allowed or No-touch.
Max No. of Promo Items	This feature is editable and provides the ability to limit how many items can be promoted for that period. Default value is the number of items at that merchandise level and name. When a period is marked as Promotion Allowed, then the Max No. of Markdown Items is disabled and set to 0. When a period is marked as No-touch, then this field is disabled and set to 0.
Max No. of Markdown Items	This feature is editable and provides the ability to limit how many items can be marked down for that period. Default value is the number of items at that merchandise level and name. When a period is marked as Markdown Allowed, then the Max No. of Promotion Items is disabled and set to 0. When a period is marked as No-touch, then this field is disabled and set to 0.
Planned Promotion Event	Name of the planned promotion event. UI does not provide the ability to modify or edit the events. They are loaded as part of data interface loads.
Apply Planned Promotion?	Indicates whether to apply the planned promotion event to the items in that merchandise level. This field provides the ability for the user to turn off planned promotions if a particular merchandise level (and selected stores) is not considered for planned promotions. Default value is that planned promotions are always turned on.
No Touch After Landing	Default value is 0 weeks and is applied at the item level (and for all the locations). Indicates that price of an item cannot be changed from the initial full price until 0 weeks have been completed after the item start selling (also referred to as model start date).
No Touch at End of Life	Default value is 0 weeks and is applied at the item level (and for all the locations). Indicates that price of an item cannot be changed in the last 0 weeks before the exit date.
Length of Regular Season	Default value is 60 percent and is applied at the item level (and for all the locations). Indicates that if the total season length of an item is 25 weeks, 15 weeks are allocated for promotions (also called regular season). Regular season occurs towards the beginning of the season. This field is enabled only for relative calendar and not for absolute calendar.
Length of Markdown Season	Default value is 40 percent and is applied at the item level (and for all selected locations). Indicates that if the total season length of an item is 25 weeks, 10 weeks are allocated for markdowns (also called clearance season). Clearance season occurs towards the end of the season. This field is enabled only for relative calendar and not for absolute calendar.
Min No. of Promotion Items	This feature is editable and provides the ability to specify minimum number of items to be promoted for that period. Default value is 0. When a period is marked as Markdown Allowed, then the Min No. of Promotion Items is disabled and set to 0. When a period is marked as No-touch, then this field is disabled and set to 0.

Table 10–11 (Cont.) Temporal Rules

Field	Description
Min No. of Markdown Items	This feature is editable and provides the ability to specify minimum number of items to be marked down for that period. Default value is 0. When a period is marked as Promotion Allowed, then the Min No. of Markdown Items is disabled and set to 0. When a period is marked as No-touch, then this field is disabled and set to 0.

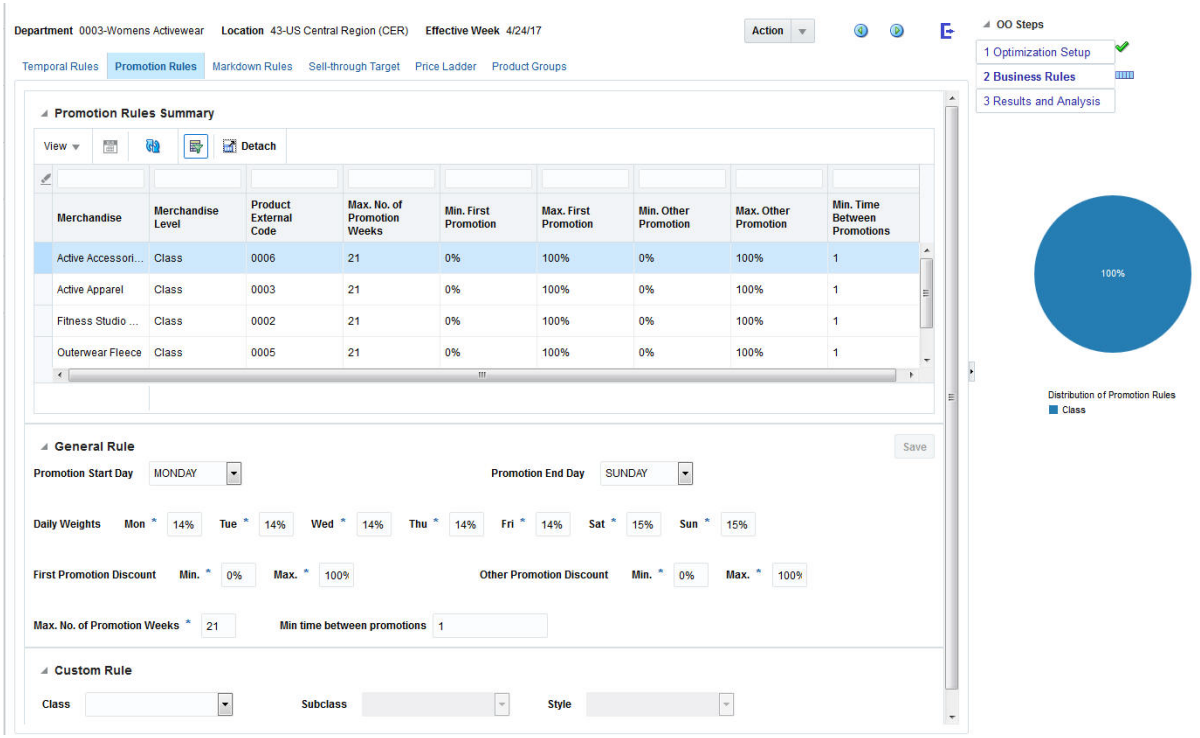
Promotion Rules

In this screen, as shown in Figure 10–16, you can specify the rules associated with promotions, specifically the reactive promotions (and not the planned promotions).

The default values that are populated for this run, such as General Rule, are based on global parameters available in RSE_CONFIG. The retailer can choose to modify the global values during implementation.

Contextual area shows the percentage of rules specified by the merchandise-level, that is, percentage specified at Class-level vs. Subclass vs. Style/Color level.

Figure 10–16 Promotion Rules



The fields available are described in Table 10–12.

Table 10–12 Promotion Rules

Field	Description
Merchandise	Description of the merchandise.
Merchandise Level	Merchandise level of the rule. It can be Class, Subclass, or Style/Color.

Table 10–12 (Cont.) Promotion Rules

Field	Description
Product External Code	The merchandise level's external code that can be identified by the retailer
Max No. of Promotion Weeks	Default is total number of weeks in the season. This allows the user to limit how many weeks are available for promotions for a particular merchandise.
Min First Promotion Discount	Default value is 0 percent based on a global parameter specified on RSE_CONFIG. The minimum discount that is allowed for the first promotion in the season.
Max First Promotion Discount	Default value is 0 percent based on a global parameter specified on RSE_CONFIG. The maximum discount that is allowed for the first promotion in the season.
Min Other Promotion Discount	Default value is 0 percent based on a global parameter specified on RSE_CONFIG. The minimum discount that is allowed for the subsequent promotions (except the first) in the season.
Max Other Promotion Discount	Default value is 100 percent based on a global parameter specified on RSE_CONFIG. The maximum discount that is allowed for the subsequent promotions (except the first) in the season.
Min Time Between Promotions	Default value is 1 week based on a global parameter specified on RSE_CONFIG. The minimum time (in periods or weeks) between consecutive promotions in the season. Suppose there are 4 months in the promotion season, and the min time is 4 weeks, then roughly 1 promotion week is available per month.
Promotion Start Day	Default value is Monday. This says when the promotion can start within a week.
Promotion End Day	Default value is Sunday. This says when the promotion can end within a week.
Daily Weights	Default values for each day of the week is populated based on the global values specified in RSE_CONFIG or values coming through an interface. These values represent the default values for all customer segments and for all time periods (e.g. weeks). This helps spread the weekly forecast into daily forecasts depending on when the promotion is applicable (e.g., Thu-Sun) for that week.

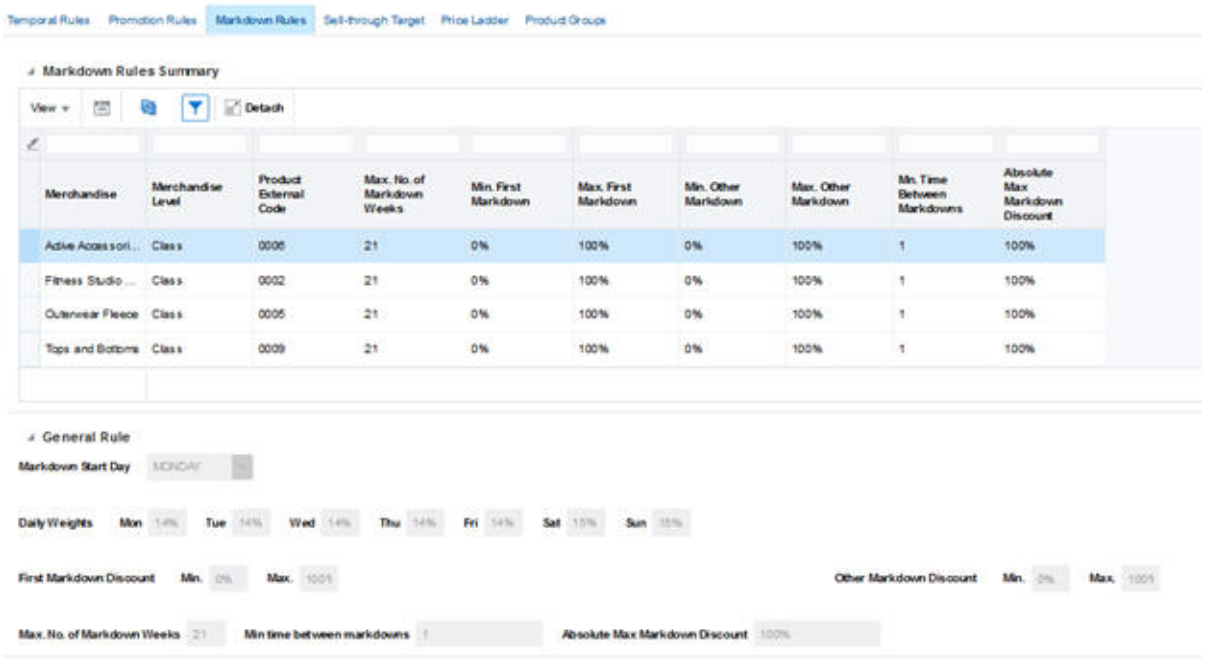
Markdown Rules

In this screen, as shown in Figure 10–17, you can specify the rules associated with markdowns.

The default values that are populated for this run, such as General Rule, are based on global parameters available in RSE_CONFIG. The retailer can choose to modify the global values during implementation.

Contextual area shows the percentage of rules specified by the merchandise-level, that is, percentage specified at Class level vs. Subclass vs. Style/Color level.

Figure 10–17 Markdown Rules



The fields available are described in Table 10–13.

Table 10–13 Markdown Rules

Field	Description
Merchandise	Description of the merchandise.
Merchandise Level	Merchandise level of the rule. It can be Class, Subclass or Style/Color.
Product External Code	The merchandise level's external code that can be identified by the retailer.
Max No. of Markdown Weeks	Default is total number of weeks in the season. This allows the user to limit how many weeks are available for markdowns at particular merchandise.
Min First Markdown Discount	Default value is 0 percent based on a global parameter specified on RSE_CONFIG. The minimum discount that is allowed for the first markdown in the season.
Max First Markdown Discount	Default value is 100 percent based on a global parameter specified on RSE_CONFIG. The maximum discount that is allowed for the first markdown in the season.
Min Other Markdown Discount	Default value is 0 percent based on a global parameter specified on RSE_CONFIG. The minimum discount that is allowed for the subsequent markdown (except the first) in the season.
Max Other Markdown Discount	Default value is 100 percent based on a global parameter specified on RSE_CONFIG. The maximum discount that is allowed for the subsequent markdowns (except the first) in the season.

Table 10–13 (Cont.) Markdown Rules

Field	Description
Min Time Between Markdowns	Default value is 1 week based on a global parameter specified on RSE_CONFIG. The minimum time (in periods or weeks) between consecutive markdowns in the season. Suppose there are four months in the clearance season, and the min time is four weeks, then roughly one markdown week is available per month.
Markdown Start Day	Default value is Monday. This says when the markdown starts within a week. Markdowns are assumed to last the entire full week.
Daily Weights	Default values for each day of the week is populated based on the global values specified in RSE_CONFIG or values coming through an interface. These values represent the default values for all customer segments and for all time periods (e.g. weeks). This helps spread the weekly forecast into daily forecasts depending on when the markdown is applicable (e.g., Thu-Wed).
Max Absolute Markdown Discount	Default value is 100 percent, based on a global parameter specified on RSE_CONFIG. The maximum discount (of full price) that is allowed for any markdown in the season.

Min Other Markdown and Max Other Markdown Discount need further explanation. Consider the following example. Suppose the user sets these values to 25 percent and 50 percent, respectively, for an item with full price of $100. If the first markdown was 20 percent, then the item's ticket price drops to $80. Because of the Min Other and Max Other Markdown Discount, the second markdown will be limited to a value between $80 *(1-0.25) = $60 (or 40 percent off) and $80 * (1-0.50) = $40 (or 60 percent off). Of course, if there was a Max Absolute Markdown Discount of 50 percent, then the second markdown would be further constrained to be between $60 and $50.

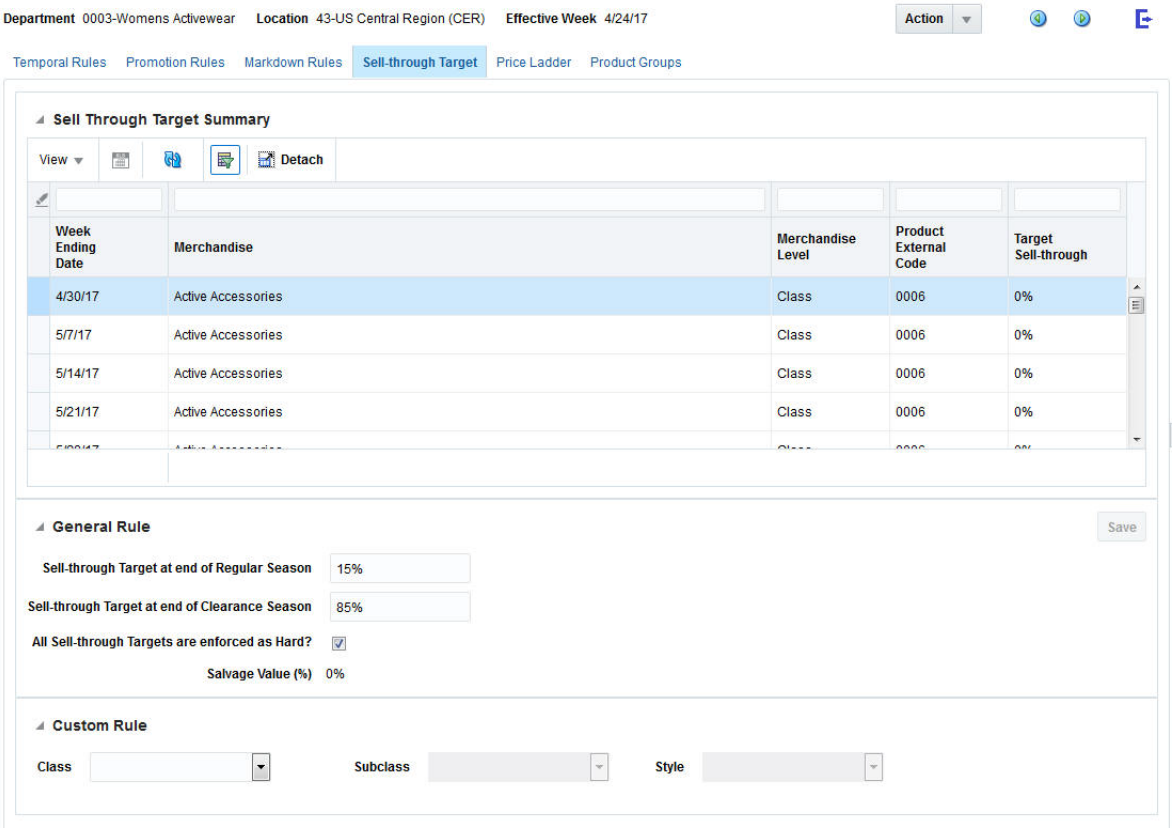
Sell-Through Target

In this screen, as shown in Figure 10–18, you can specify the sell-through targets by week and by merchandise. The default values that are populated for this run, such as General Rule, are based on global parameters available in RSE_CONFIG. The retailer can choose to modify the global values during implementation.

Contextual area shows the percentage of rules specified by the merchandise-level, that is, percentage specified at Class level vs. Subclass vs. Style/Color level.

Depending on the calendar option chosen in the Temporal Rules Screen, this screen will use either the Absolute Calendar or the Relative Calendar option. The Absolute Calendar is useful when most of the items are introduced and exit at almost the same time. The Relative Calendar is useful when the retailer introduces items at frequent intervals and having a fixed calendar for all items is restrictive and pushes items into promotion or markdown too quickly during the items' life.

Figure 10–18 Sell-Through Target



The fields available are described in Table 10–14.

Table 10–14 Sell-Through Target

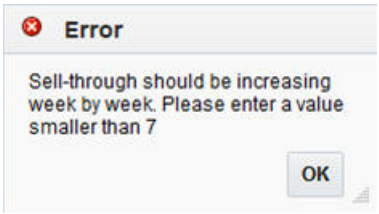
Field	Description
Week Ending Date	The week ending date for the period. Not editable field.
Weeks From Start	This field is non-editable and it runs from Week 1 to Maximum no. of weeks available for the selected merchandise level. Depending on the no-touch values, the Weeks from Start are adjusted to begin at 1+No touch after Landing and end at N-No touch at end of life. Available with only Relative Calendar option. Available with only Relative Calendar option.
Merchandise	Description of the merchandise.
Merchandise Level	Merchandise level of the rule. It can be Class, Subclass or Style/Color.
Product External Code	The merchandise level's external code that can be identified by the retailer.
Target Sell-through	Sell-through target at the end of the specified period for the merchandise as percentage (0% - 100%). If enforced as Hard, then this sell-through target must be met otherwise the optimization will not return any solution.
Sell-through Target at end of Regular Season	Default value is 0 percent based on global parameter specified on RSE_CONFIG. The sell-through target that needs to be achieved by the end of the regular season.

Table 10–14 (Cont.) Sell-Through Target

Field	Description
Sell-through Target at end of Clearance Season	Default value is 0 percent based on global parameter specified on RSE_CONFIG. The sell-through target that needs to be achieved by the end of the clearance season. This value has to be higher than the target at the end of the regular season otherwise it is ignored.
All Sell-through Targets are enforced as Hard?	Default value is soft (or unchecked) based on global parameter specified on RSE_CONFIG. If enforced as Hard, then all the sell-through targets must be met otherwise the optimization will not return any solution.
Salvage Value %	Non-editable field. Default value is 0 percent based on global parameter specified on RSE_CONFIG. The Salvage value for an item is calculated using this percentage as, Salvage value percentage multiplied by the full price of the item.

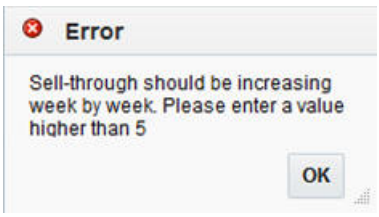
When you enter custom sell-through target values by week, the system validates that the percentage must be non-decreasing for the selected merchandise level. It returns an error indicating the same. In Figure 10–19, the user entered a value higher than 7 for Week Ending Date 5/28/17.

Figure 10–19 Error: Value Greater Than Seven



Similarly, if the user entered a value lower than 5 then that also throws a validation error.

Figure 10–20 Error: Value Less Than Five

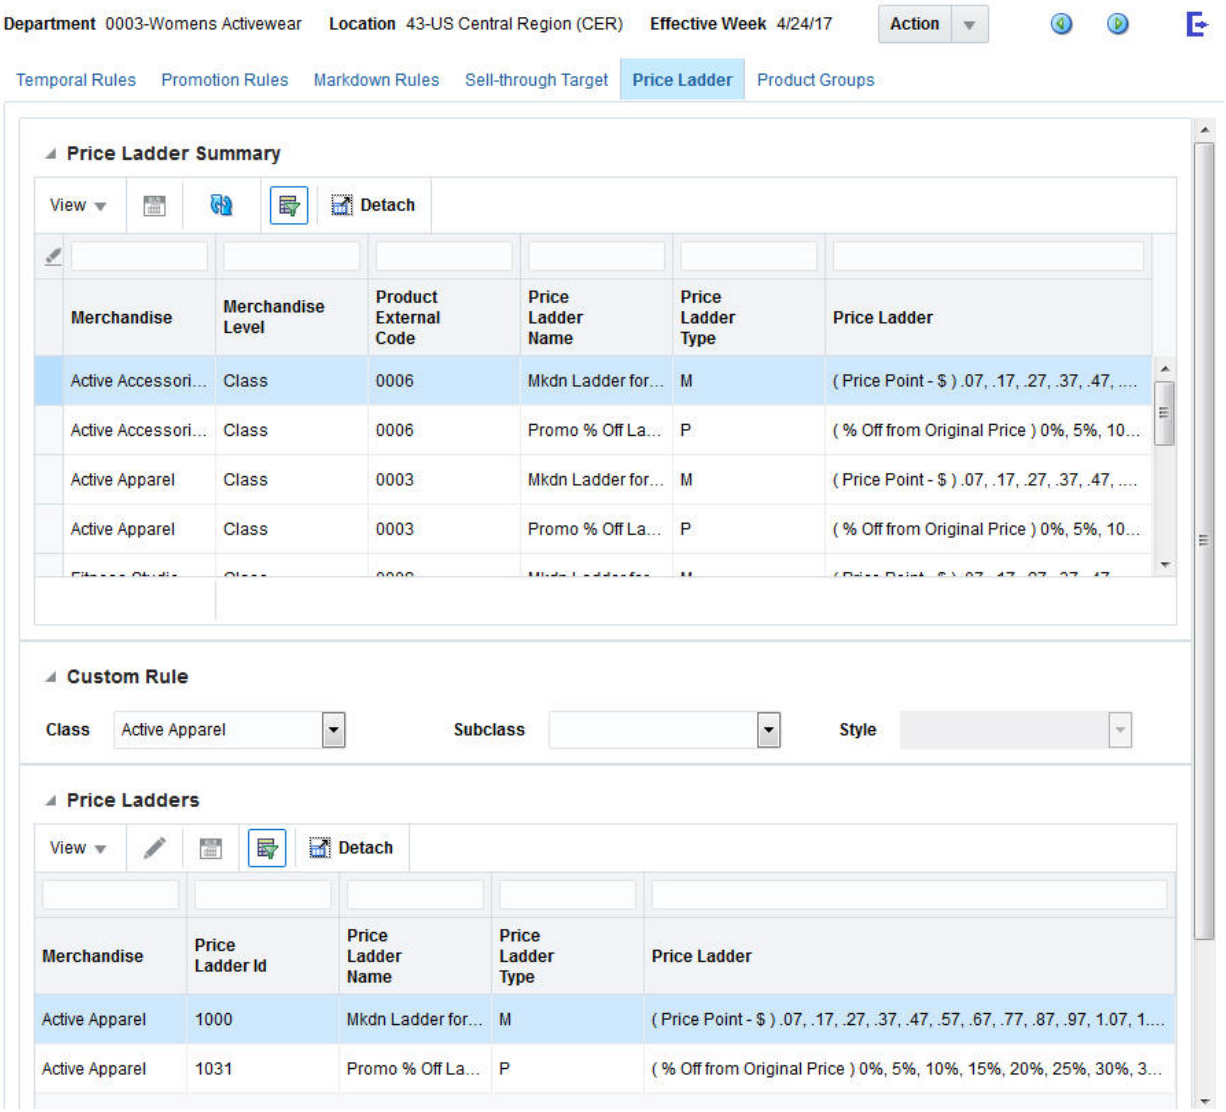


Price Ladder

In this screen, as shown in Figure 10–21, you can assign price ladders that have been loaded and are available for the selection location and merchandise levels. You can load as many price ladders at different merchandise levels (Class and above). By default, the system assigns one price ladder for promotion and one for markdown at the Class level that is closest to the class under consideration.

The contextual area shows the percentage of rules specified by the merchandise level, that is, the percentage specified at Class level vs. Subclass vs. Style/Color level.

Figure 10–21 Price Ladders



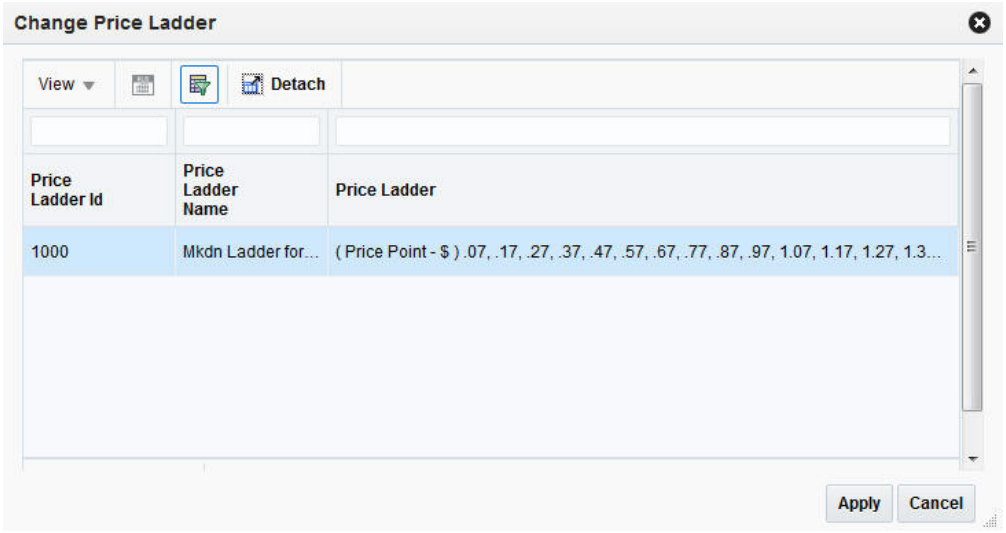
The fields available are described in Table 10–15.

Table 10–15 Price Ladders

Field	Description
Merchandise	Description of the merchandise.
Merchandise Level	Merchandise level of the rule. It can be Class, Subclass or Style/Color.
Product External Code	The merchandise level's external code that can be identified by the retailer.
Price Ladder Name	Name for the price ladder.
Price Ladder Type	Type indicates whether price ladder is for Promotions (P) or Markdowns (M).
Price Ladder Values	Values available for this price ladder. Price Ladder can be a percentage ladder (off full price or current ticket price) or a price-point ladder. Currency for the run shown depends on the locale loaded during the implementation.

As shown in Figure 10–22, in the Custom Rule, you can pick one price ladder for Promotion and one for Markdown, based on the set of available ladders at that merchandise level. You can select one row and click Edit to override the price ladder for that merchandise and price ladder type. You can select any price ladder in the Change Price Ladder pop-up and click Apply.

Figure 10–22 Change Price Ladder



Product Groups

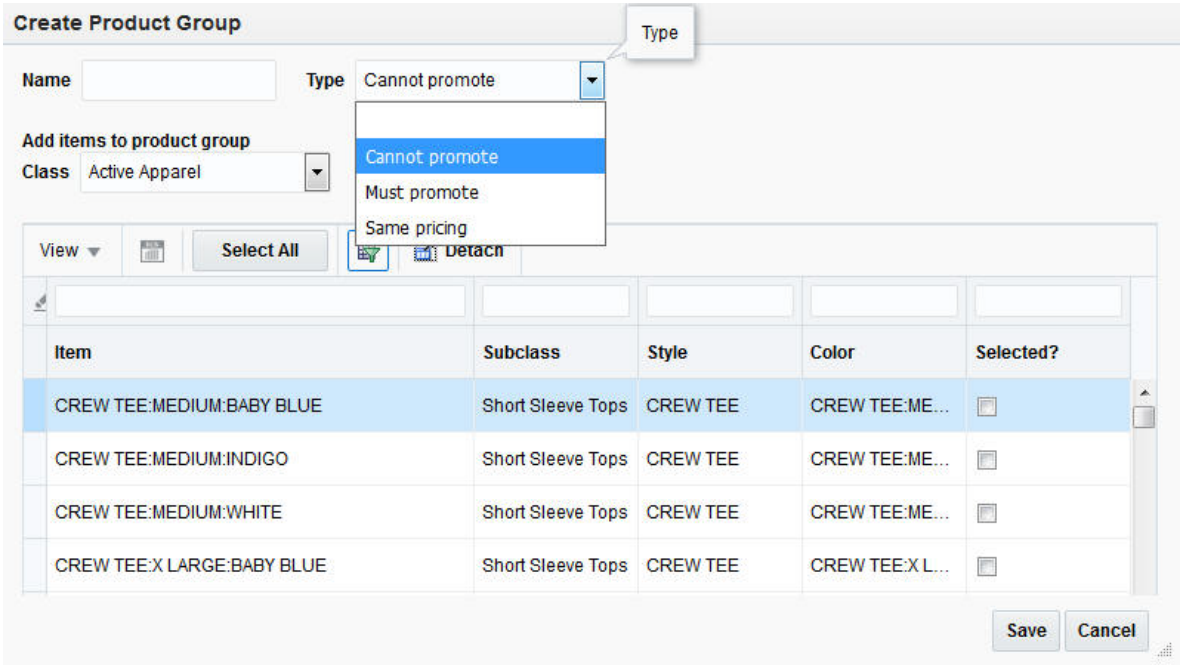
In this screen, shown in Figure 10–23, you can group products within a class such as Cannot Promote, Must Promote, or Same Pricing.

- In the Cannot Promote group, at most one product of the group can be promoted in a given week (or time period). For example, competing brands such as Levis vs. Calvin Klein cannot be promoted at the same time.
- In the Must Promote group all the items must be promoted if any one item of the group is promoted in a given week (or time period).
- In the Same Pricing group, all the items receive same percentage off recommendations over the life of the item. Notice that the final absolute recommended price can be different for items within this group.

You can click Create New to display a Create Product Group interface. You must enter a group name and the group type from the drop-down list. You can select a product, one at a time, and add it to the group. Alternatively, you can filter the products using Query by Example, use Select All to add the products to the group, and then click Save. There is no conflict checking or validation in place, so you must be careful when putting products in two conflicting groups (for example, must promote and cannot promote).

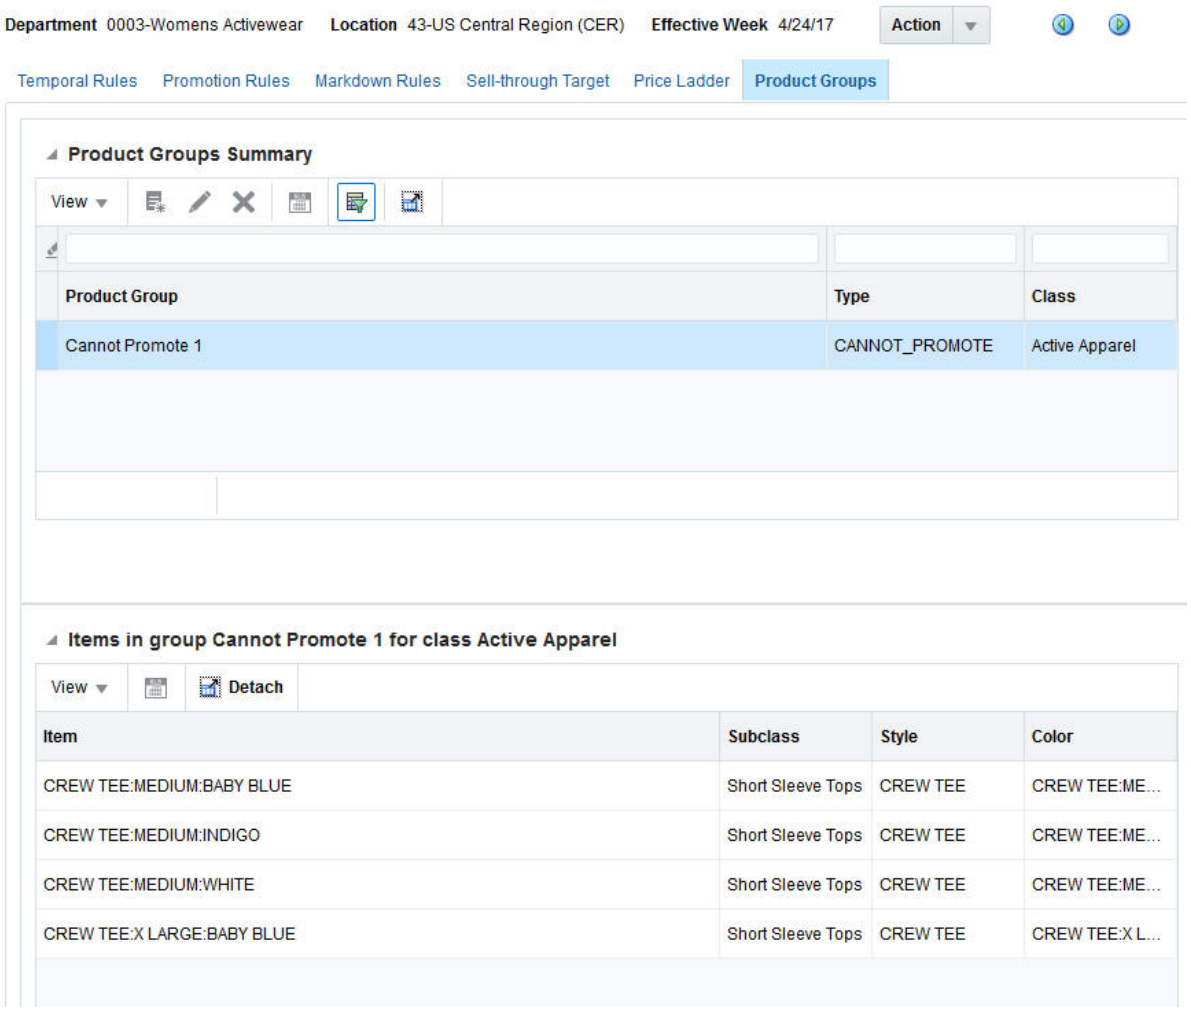
The contextual area shows the percentage of groups specified by the group type, that is, the percentage of groups categorized as Cannot Promote vs. Must Promote vs. Same Pricing.

Figure 10–23 Create Product Group



All the groups created are shown in the Product Groups Summary. You can select any group, and the set of products within that group are displayed. You can edit or delete any group shown in the Summary panel.

Figure 10–24 Product Groups Summary



Results

After the optimization run is complete, the results are displayed in the Results stage. In this stage, you can accept, reject, or override any price recommendation. After you make substantive changes, you can change the appropriate settings or recalculate all the projected metrics.

The results are displayed in five tiles: Revenue Tile, Promotion Tile, Markdown Tile, Targeted Tile, and Offers Tile.

Possible Actions

From the Results stage, you can:

- Recalculate the projected metrics after making price overrides.

  Revisit earlier stages in order to make changes to various settings and then re-optimize the run.

- Review the run and change the status to Reviewed.
- Submit the run that writes the recommendations to Export Interface and also sends the results to a downstream price execution system such as CE or RPM.

- Leave the run in the optimization list without taking any actions. You cannot delete the run; you must be in Overview in order to delete a run.

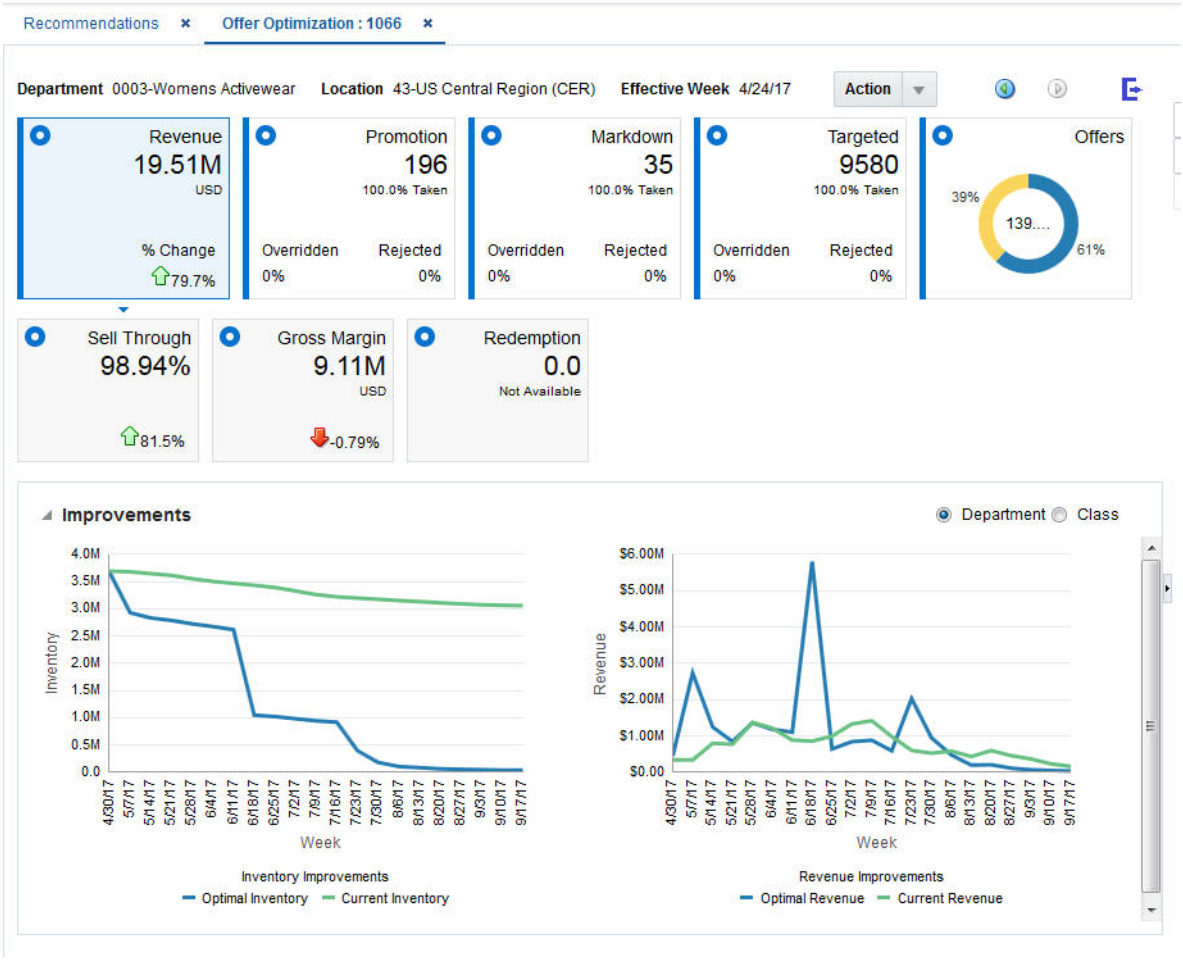
Revenue Tile

As shown in Figure 10–25, the revenue tile displays the overall summary of the run in forms of tiles as well as in BI charts. Information displayed in this tile is non-editable.

- Revenue. This is the projected revenue over the life of the items, assuming that you accept all recommendations. This is displayed in the local currency defined for the location of the run.
- Revenue % Change. This is the percentage gain or loss in revenue compared to staying at the current price for all items over their entire life.
- Sell Through. This is the projected sell through at the end of the life, assuming that you accept all recommendations. This is shown as a percentage of initial inventory (or inventory at the beginning of the season).
- Sell Through % Change. This is the percentage gain or loss in end of life sell-through, compared to staying at the current price for all items over their entire life.
- Gross Margin. This is the projected gross margin over the life of the items, assuming that you accept all recommendations. This is shown in the local currency defined for the location of the run.
- Gross Margin % Change. This is the percentage gain or loss in gross margin, compared to staying at the current price for all items over their entire life.

The bottom panel displays the improvements as two BI charts: Inventory vs. Time and Revenue vs. Time. By default, the information in the chart is at the Department level. However, you can do a deep dive into a particular class by toggling the radio button to Class and selecting a class from the drop-down list of classes.

Figure 10–25 Revenue Tile



Promotion Tile

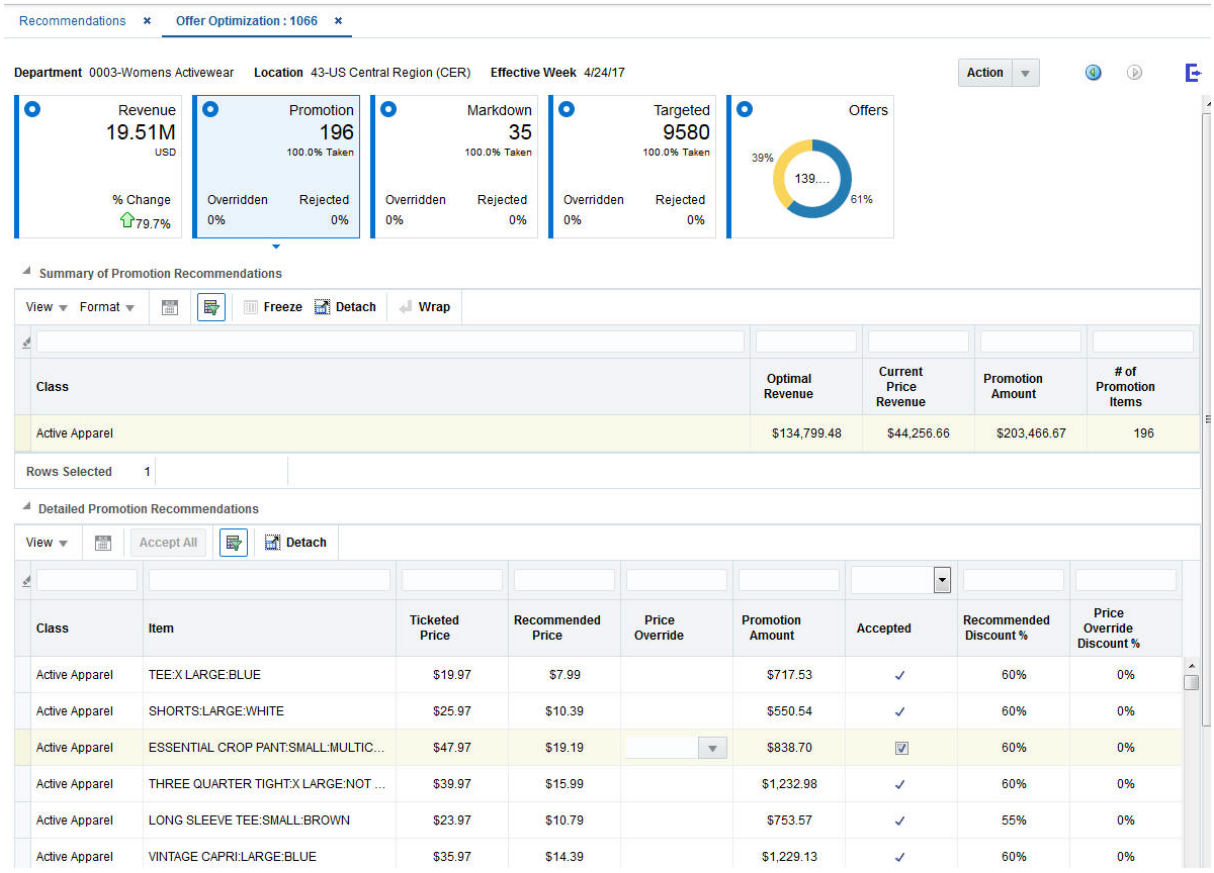
This tile displays the promotion recommendations for the effective week selected. In the tile, it shows the total number of promotion recommendations (items), percentage of recommendations taken or accepted, percentage of recommendations overridden, and percentage of recommendations rejected.

In the summary panel, the Class-level overview of the promotion recommendations is displayed. It contains the following metrics (non-editable):

- Optimal Price Revenue. Projected optimal price revenue for this week over all items with promotion recommendation at the class level.
- Current Price Revenue. Projected revenue at current price for this week over all items with promotion recommendation at the class level.
- Promotion Amount. This is calculated as the amount spent due to the discount offered multiplied by the projected sales units at the class level. This is the amount given towards promotions at the class level.
- # of Promotion Items. This is the total number of items at class level that have received a promotion recommendation.

You can select a class in the top summary panel; then the detailed panel is populated. The Detailed Panel displays all the item recommendations for the selected effective week.

Figure 10–26 Promotion Tile



In the detailed panel, you can accept, reject, or override the price recommendation for the effective week. It contains the following information:

- Class. Name of the class.
- Item. Description of the item.
- Ticket Price. Current ticket price of the item.
- Recommended Price. Price recommended by the optimization.
- Price Override. You can click the box and it displays a price ladder assigned for that item. You can choose a price from the price ladder. The price override is applied only to the effective week. Any price overrides are not reflected in the projected metrics until the user selects re-calculate button.
- Promotion Amount. This is calculated as the amount spent due to the discount offered multiplied by the projected sales units at the item level. This is the amount given towards promotions.
- Accepted. You can uncheck or check the box to indicate whether to accept or reject the recommendation. When the recommendation is rejected, it reverts to the current ticket price.
- Recommended Discount %. This field is non-editable. This shows the percentage discount offered with respect to the full price for the item selected, based on the recommended price.

- Price Override Discount %. This field is non-editable. This shows the percentage discount offered with respect to the full price for the item selected, based on the Price Override column.

You can see the full price path for an item in the contextual area, as shown in tFigure 10–27. Each price type is shown in different color: Black - Regular Price, Grey - Override, Blue - Promotion, Red - Markdown, Green - Planned Promotion. You can also see the Projected Sales and Inventory Path and Projected Budget Usage for Promotions and Markdowns (Figure 10–28).

Figure 10–27 Promotion Price Path

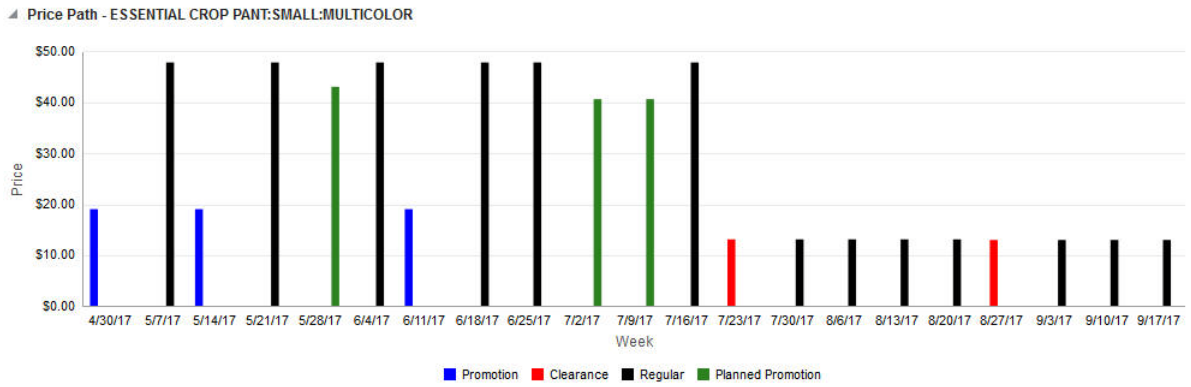
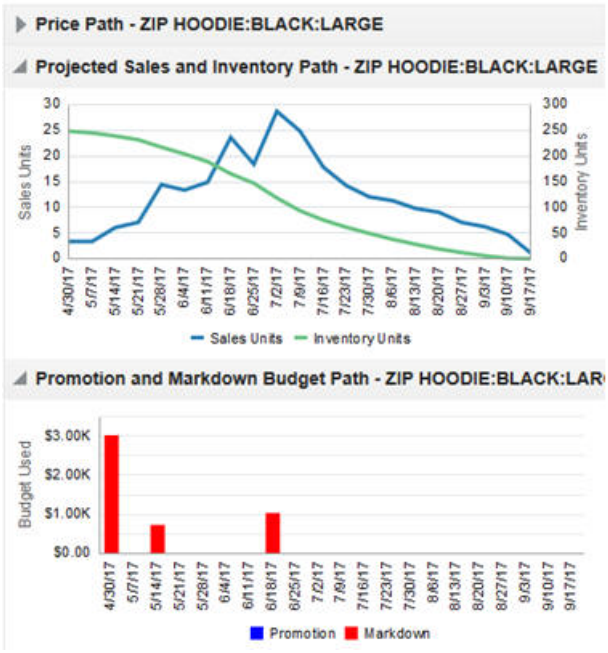


Figure 10–28 Projected Sales and Inventory Path and Budget Usage Path



Markdown Tile

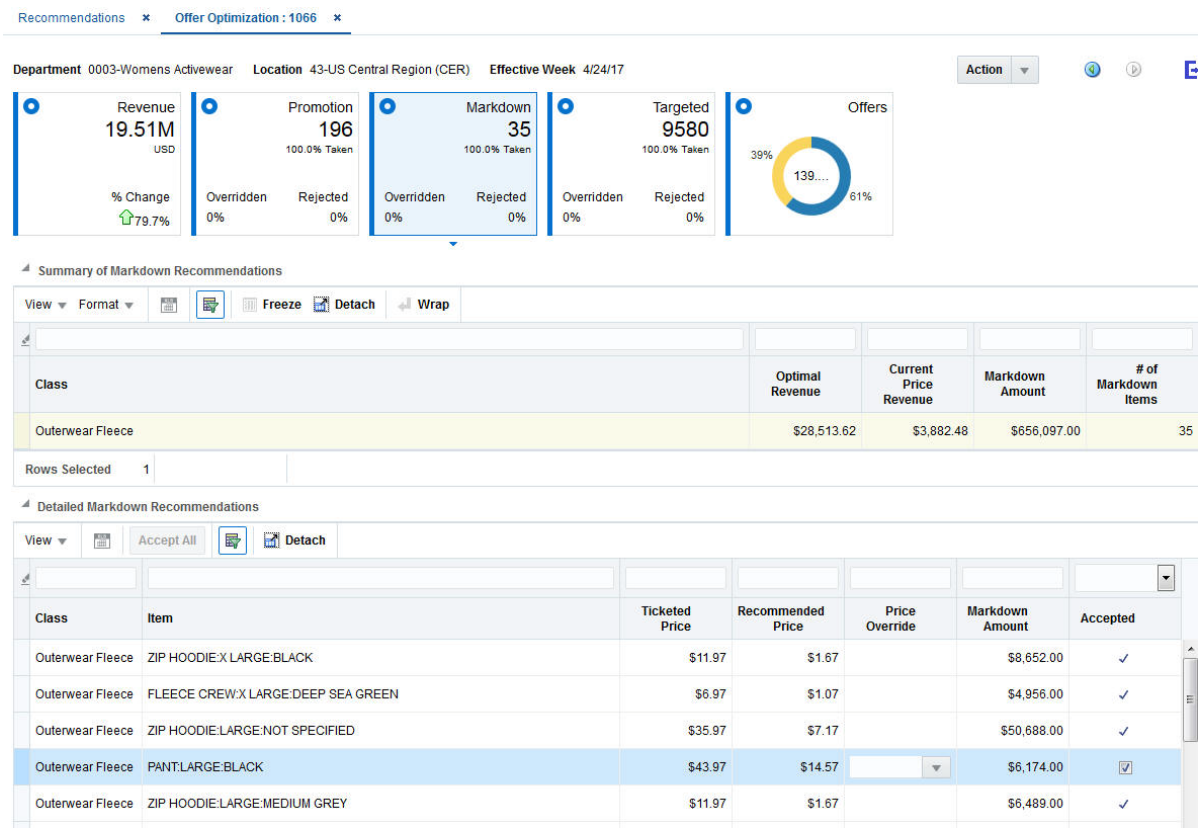
The Markdown Tile displays the markdown recommendations for the effective week selected. In the tile, it shows the total number of markdown recommendations (items), percentage of recommendations taken or accepted, percentage of recommendations overridden, and percentage of recommendations rejected.

In the summary panel, it displays the class level overview of the markdown recommendations. It shows the following metrics (non-editable):

- Optimal Price Revenue. Projected optimal price revenue for this week over all items with promotion recommendation at the class level.
- Current Price Revenue. Projected revenue at current price for this week over all items with promotion recommendation at the class level.
- Markdown Amount. This is calculated as the amount spent due to the discount offered multiplied by the projected sales units at the class level. This is the amount given towards promotions at the class level.
- # of Markdown Items. This is the total number of items at the class level that have received a markdown recommendation.

You can select a class in the top summary panel, and the detailed panel is populated. The Detailed Panel displays all the item markdown recommendations for the selected effective week.

Figure 10–29 Markdown Tile



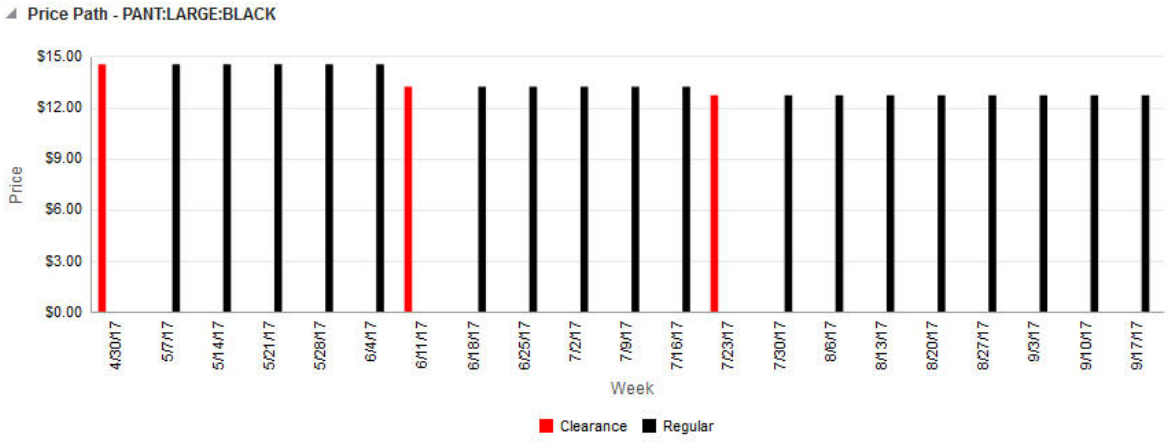
In the detailed panel, the user can accept, reject or override the price recommendation for the effective week. It contains the following information:

- Class. Name of the class.
- Item. Description of the item.
- Ticket Price. Current ticket price of the item.
- Recommended Price. Price recommended by the optimization.

- Price Override. You can click the box and it displays a price ladder assigned to that item. You can select a price from the price ladder. The price override is applied to the effective week and propagated appropriately to future weeks. Any price overrides are not reflected in the projected metrics until the user selects the re-calculate button.
- Markdown Amount. This is calculated as the amount spent due to the discount offered multiplied by the projected sales units at the item-level. This is the amount given towards promotions.
- Accepted. You can uncheck or check the box to denote whether to accept or reject the recommendation. When the recommendation is rejected, it reverts to the current ticket price.

You can see the full price path for an item in the contextual area, as shown in Figure 10–30. Each price type is shown in a different color: Black - Regular Price, Grey - Override, Red - Markdown, Green - Planned Promotion. You can also see the Projected Sales and Inventory Path and Projected Budget Usage for Promotions and Markdowns (Figure 10–28).

Figure 10–30 Markdown Price Path



Targeted Tile

The Targeted Tile displays the targeted recommendations for the effective week selected. In the tile, it shows the total number of targeted recommendations (items), percentage of recommendations taken or accepted, percentage of recommendations overridden, and percentage of recommendations rejected.

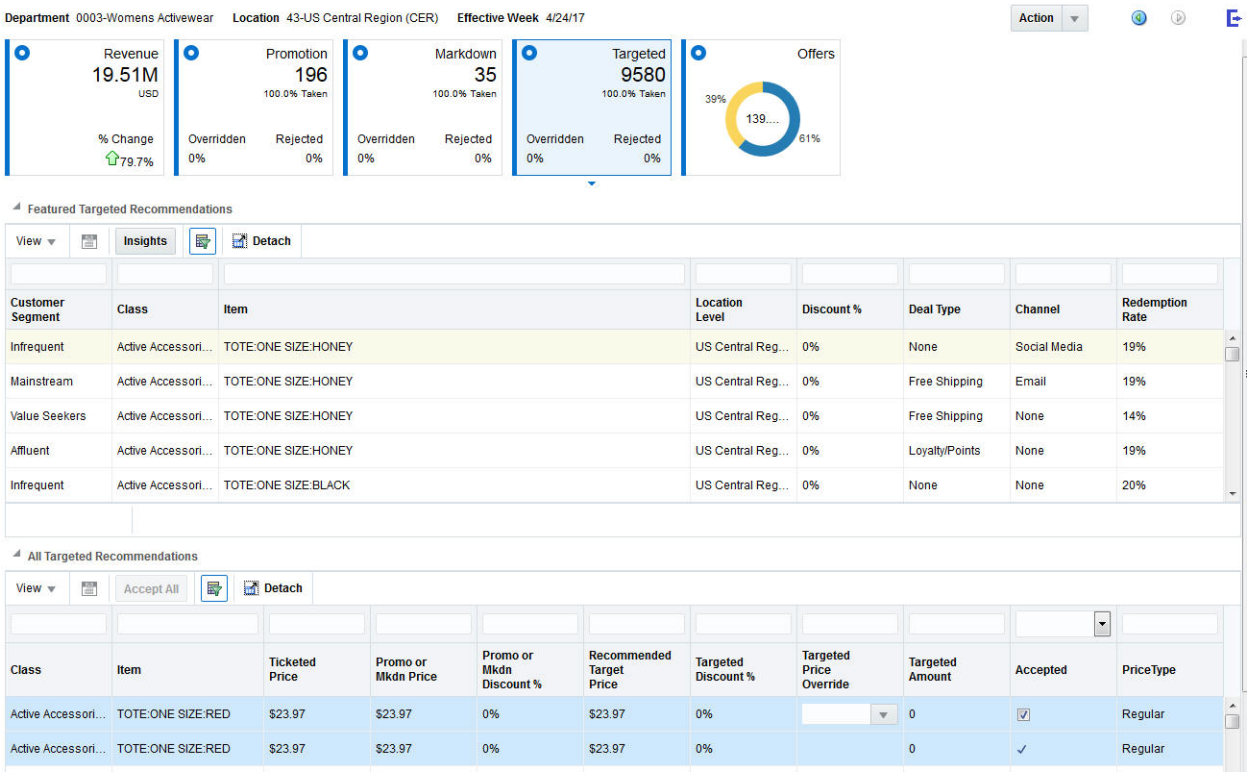
In the Featured Targeted Offers panel, it displays Item, Deal type, and Channel offers by class and segment that result in the best redemption rate. It displays the following metrics (non-editable):

- Customer Segment. Name of the customer segment.
- Class. Name of the class.
- Location. Name of the selected location.
- Discount %. Recommended discount percentage for the item.
- Deal Type. Deal type (for example, Z% Off, Loyalty Points) associated with this item that result in a high redemption rate for this segment.

- Channel. Channel (for example, Email, Mobile) associated with this item that result in a high redemption rate for this segment.
- Redemption Rate. Predicted redemption rate for this offer.

You can select a class segment in the top featured targeted offers panel. The detailed panel is then populated with all the item level price recommendations for this effective week. You can override the targeted price recommendations in this panel.

Figure 10–31 Targeted Tile



In the detailed panel, you can accept, reject, or override the targeted price recommendation for the effective week. It contains the following information:

- Class. Name of the class.
- Item. Description of the item.
- Ticket Price. Current ticket price of the item.
- Promo or Markdown Price.
- Promo or Markdown Discount %.
- Recommended Target Price. Price recommended by the optimization.
- Targeted Discount %.
- Targeted Price Override. You can click the box, and it displays a price ladder assigned for that item. You can select a price from the price ladder.
- Targeted Amount. This is calculated as the amount spent due to the discount offered multiplied by the projected sales units at the item-level. This is the amount given towards targeted offers.

- Price Type. Denotes whether the price recommendation is regular, markdown, promotion, or targeted.
- Accepted. You can uncheck or check the box to denote whether to accept or reject the targeted recommendation. When recommendation is rejected, it reverts to location-level price recommendation or override.

You can see the product image (if loaded) and the current price, recommended price, and targeted price for the item in the contextual BI area.

Figure 10–32 Targeted Price Recommendation

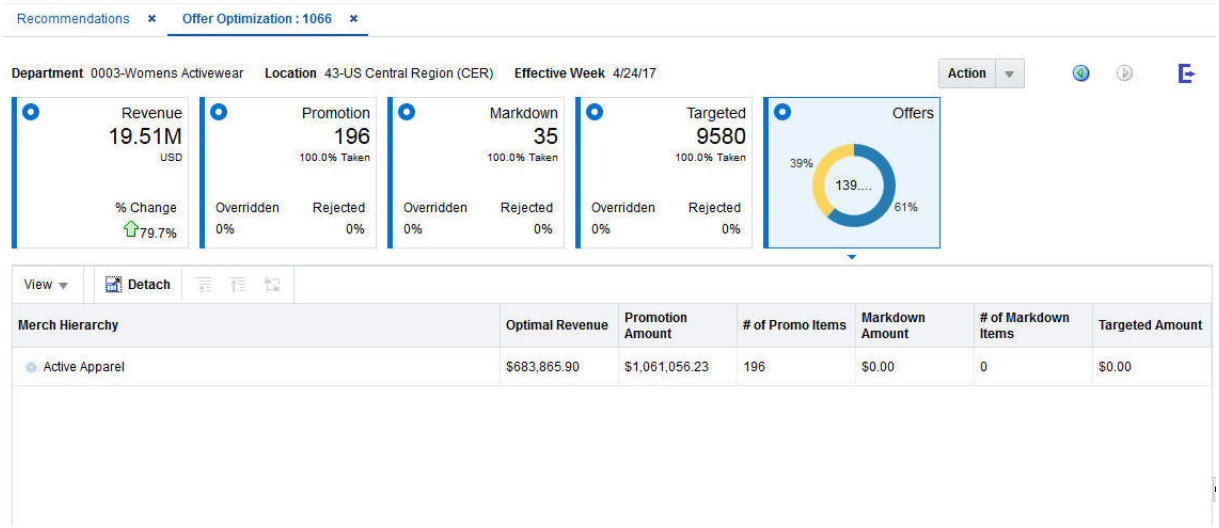


Offers Tile

The Offers Tile displays all the recommendations for the effective week selected. In the tile, it displays the percentage of the amount used for promotions, the percentage of amount used for markdowns, and the percentage of amount used for targeted offers.

In the tree table below, it shows the number of promotion recommendations, the number of markdown recommendations, and the corresponding amounts by merchandise hierarchy.

Figure 10–33 Offers Tile



Export

The accepted recommendations for runs that are in Submitted status are sent for export. Once the run has entered a Submitted status, you cannot modify any aspects of the run. If you decide to send another run for the same Department, Location, Season, and Effective Week, then the new recommendations will override the old recommendations. See the interface guide for details on the export interface.

Innovation Workbench

See the Innovation Workbench chapters for details. Apart from the core hierarchy tables, the following validation and result tables that are applicable for offer optimization are exposed in read-only mode. You can slice and dice the results table for reporting purposes and examine the validation table to determine why a particular run has ended in Failed status. The tables are:

- PRO_RUN_SANITY_CHECK_RSE_VW contains information on the errors/alerts/warnings generated.
- PRO_RUN_RECOM_OPT_RESULT contains all the results associated with promotions, markdowns, and targeted pricing recommendations.

Integration with Oracle Digital Assistant

Offer Optimization is integrated with a voice assistant such as Oracle Digital Assistant to handle complex workflow and business process questions. The functionality integrates voice assistants, machine learning techniques such as NLP, and a pricing solution such as Offer Optimization. This reduces the workflow for a typical user of such software and helps in handling complex business process use cases. For example, buyers or managers who manage price recommendations for retail apparel might be interested in use cases such as:

- Chain/country-level view. For example, a user can ask to see the performance of a department or all departments at the country level.

  Exception management. For example, a user can ask to see the locations for a department where the margin improvements are lower than 5%.
- Approval process. A few use cases are possible. A user can approve all batch runs or select specific runs for approval. For example, a user can reject the price recommendations for runs that have a discount higher than 40% or approve all batch runs that show an improvement of 5% for revenue.
- Recalculate the key performance metrics. For example, a user can reject recommendations that have a discount higher than 70% and recalculate the metrics.
- Q&A. This integration helps in answering a user's most frequently asked questions. For example, a user can ask about promotions and the Q&A tab brings up the relevant questions and answers related to the term "promotions."

A user can enable the Oracle Digital Assistant by clicking the microphone symbol. A user can say a particular question (utterance) to a voice-enabled device, and the utterance is converted into text and mapped to an intent/entity using existing Natural Language Processing algorithms such as Oracle Intelligent Bots. Once the intent/entity has been identified, it is sent to a pricing solution such as Offer Optimization, which then performs the appropriate processes and returns information.

Some examples of the conversation for the above use cases are provided below. The user utterance is prefixed with "User" and the response from ODA is prefixed as "Bot."

Sample Conversation: Country-Level View for All Departments

User: "I would like to see how my stores in US are performing."

Bot: "Okay. It seems as if you want to take a look at the country-level metrics for stores. Please specify the location from the following: <America, Europe, E-commerce US>?"

User: "America."

Bot: "Do you want to review a specific department or all departments?"

User: "All departments."

Bot: "Aggregated metrics for the specified location(s) and all departments are: Revenue is 2756451.736, Change in Revenue is -1.49%, Gross Margin is 1868901.003, Change in Gross Margin is -3% , Sell Through is 0.6861, Change in Sell Through is 1.28%."

Sample Conversation: Country-Level View for Selected Departments

User: "I would like to see how my stores in US are performing."

Bot: "Okay. It seems as if you want to take a look at the country-level metrics for stores. Please specify the location from the following: <America, Europe, E-commerce US>?"

User: "America."

Bot: "Do you want to review a specific department or all departments?"

User: "Specific Department."

Bot: "Which department are you interested in?"

User: "Womens Activewear"

Bot: "Aggregated metrics for the specified location(s) and department(s) are: Revenue is 2756451.736, Change in Revenue is -1.49%, Gross Margin is 1868901.003, Change in Gross Margin is -3% , Sell Through is 0.6861, Change in Sell Through is 1.28%."

Sample Conversation: Exception Management

User: "I want to see the locations that do not meet the approval criteria."

Bot: "Okay. It seems as if you are interested in finding exceptions. First, which department are you interested in?"

User: "Womens Activewear."

Bot: "What metric are you looking for: Revenue Improvement, Margin Improvement, Sellthrough Improvement, or Sellthrough?"

User: "Sellthrough."

Bot: "Please specify the threshold (for example, 15%) for identifying the poorly performing locations."

User: "10%."

Bot: "The location for Department Womens Activewear with Sellthrough below 10% is Company."

Bot: "Would you like to see the run summary for this Department and the location?"

User: "Yes."

Bot: "The metrics for run Copy_Of_test are: Revenue is 2756451.736, Change in revenue is -1.49%, Gross Margin is 1868901.003, Change in gross Margin is -3% , Sell through is 0.6861, Change in sell through is 1.28%."

Bot: "Would you like to approve the run - Yes/No?"

Depending on what the user answers here, the flow will be different. If the user says "Yes," then the user is taken to the Approval Use Case. If the user says "No," then the user is taken to Recalculate Use Case.

Sample Conversation: Approval Use Case–Single Run

User: "I want to review the metrics and would like to take care of run approvals today."

Bot: "Okay. It seems as if you are interested in approving runs. Do you want to approve all batch runs or a single location?"

User: "Single location."

Bot: "Which location are you interested in?"

User: "Company."

Bot: "Which department are you interested in?"

User: "Womens Activewear."

Bot: "The metrics for the run Weekly Run 04242017 are: Revenue is 2756451.736, Change in Revenue is -1.49%, Gross Margin is 1868901.003, Change in Gross Margin is -3% , Sell Through is 0.6861, Change in Sell Through is 1.28%."

Bot: "Would you like to approve the run - Yes/No?"

User: "Yes."

Bot: "The run has been approved, and the offers have been sent to the customers."

Bot: "Please let me know how can I help you with the Offer Optimization application."

Sample Conversation:Approval Use Case–All Batch Runs

User: "I want to review the metrics and would like to take care of run approvals today."

Bot: "Okay. It seems as if you are interested in approving runs. Do you want to approve all batch runs or a single location?"

User: "All batch runs."

Bot: "Which department are you interested in?"

User: "Womens Activewear."

Bot: "All five batch runs have been approved."

Bot: "Please let me know how can I help you with the Offer Optimization application."

Sample Conversation: Recalculate Use Case

User: "I want to review the metrics and would like to see whether to make any changes to the recommendations."

Bot: "Okay. It seems as if you are interested in accepting or rejecting recommendations. First, which location are you interested in?"

User: "Company."

Bot: "Which department are you interested in?"

User: "Womens Activewear."

Bot: "The metrics for run Copy_Of_test are Revenue is 2756451.736, Change in revenue is -1.49%, Gross Margin is 1868901.003, Change in gross Margin is -3% , Sell through is 0.6861, Change in sell through is 1.28%."

Bot: "Would you like to approve the run - Yes/No?"

User: "No."

Bot: "Would you like to ACCEPT/REJECT available recommendations? - ACCEPT/REJECT."

User: "ACCEPT."

Bot: "Would you like to ACCEPT promotion/markdown recommendation? - PROMOTION/MARKDOWN."

User: "Promotion."

Bot: "What is the maximum percent discount criterion to accept Promotion recommendations?"

User: "40%."

Bot: "Promotion Price recommendations below 40% have been accepted."

Bot: "Would you like to recalculate the metrics again - Yes/No?"

User: "Yes."

Bot: "The run has been submitted for recalculating the metrics. To see the new metrics, please query the bot."

Bot: "Please let me know how I can help you next with the Offer Optimization application."

11

Size Profiles

This chapter describes the Size Profiles Cloud Service module.

Overview

The Size Profiles (SP) is a module under Retail Profile Sciences and is used to estimate the distribution of demand across different sizes (size profile) for different merchandise and location levels.

Size profile is estimated at different levels of merchandise and location. The lowest level of estimation is style-color (for merchandise) and store (for location). Size profiles may also be estimated at higher levels of aggregation on both the merchandise and location dimensions (for example, at subclass-store or style/color-store cluster) depending on the retailer's requirements. Irrespective of the level of estimation, the final output for size profiles is at lowest level (style-color/store) and size profiles are shown at the lowest level in the UI.

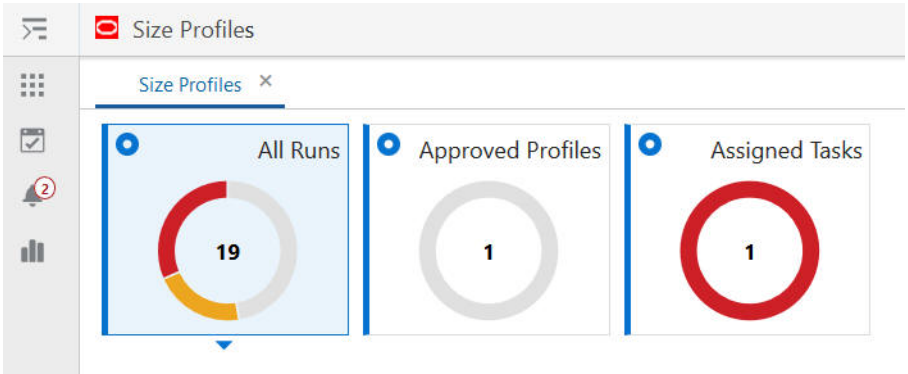
At a high level, the workflow for SP consists of creating a run, reviewing, overriding, and approving the output (that is, size profiles), and submitting the output to another user for further review or to a downstream application.

Navigation

You can use the Info tiles at the top of the screen, shown in Figure 11–1, to navigate to:

- All Runs, which displays all previously created runs.
- Approved Profiles, which displays the approved profiles for all runs.
- Assigned Tasks, which displays all profiles assigned to the user for review.

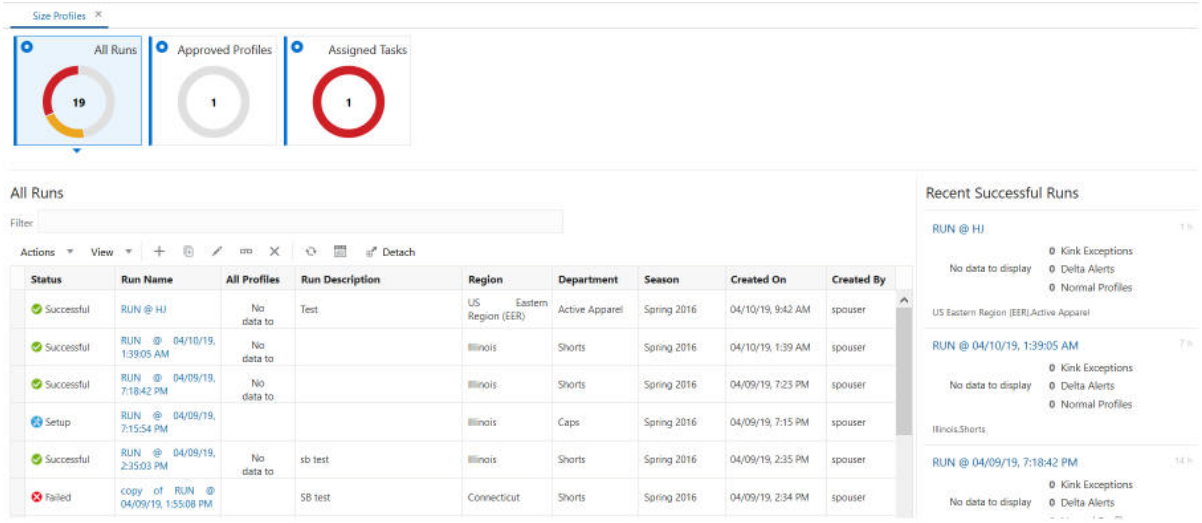
Figure 11–1 Info Tiles



All Runs

The All Runs screen, shown in Figure 11–2, lists all previously created runs, regardless of their status.

Figure 11–2 All Runs



Click on the **Run Name** on any run with a status of Setup, In progress, or Successful, in order to open that run. The Output tab opens, as shown in Figure 11–7.

When you select a run with a status of Setup, you can modify the information and parameters for creating a run and complete the Create Run process, as described in "Creating a Run".

When you select a run with a status of Successful or In Progress, you see the run output tab, as described in "Run Output".

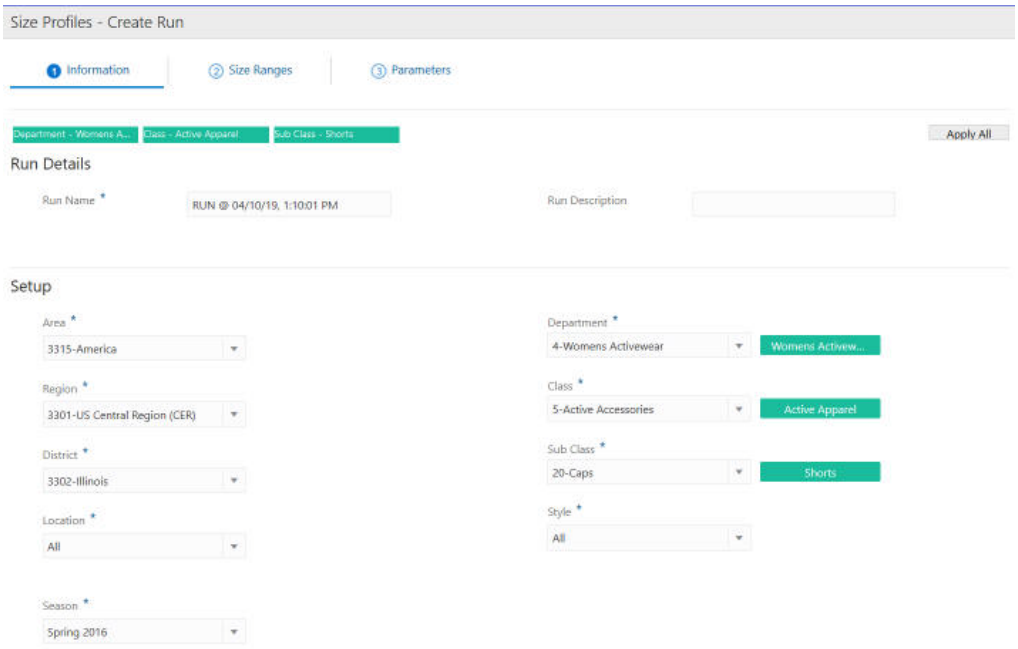
You can filter any of the columns by entering text into the filter bar.

The Recent Successful Run list is displayed in the contextual area on the right. You can also click on a run in this list in order to open it.

Creating a Run

To estimate size profiles, you create a run and specify the settings to be used for the estimation. To begin the process, either click the **Create** icon at the top of the overview table or select a row in the overview table and click the **Duplicate** icon to duplicate a run. In either case, the Create Run screen is displayed, as shown in Figure 11–3.

Figure 11–3 Create Run Showing Information Tab



The Create Run screen has the following features:

- Three tabs: the "Information Tab", the "Size Ranges Tab", and the "Parameters Tab".
- The **Save** button, which saves the changes made in any of the tabs.
- The **Cancel** button, which discards the changes in the current tab and closes the Create Run screen.
- The **Submit** button, which saves all changes, begins the execution of the run, and closes the screen. At this point, the status of the run changes from Setup to In Progress.
- The contextual area on the right, which displays a summary of user selections in each of the three tabs.

Information Tab

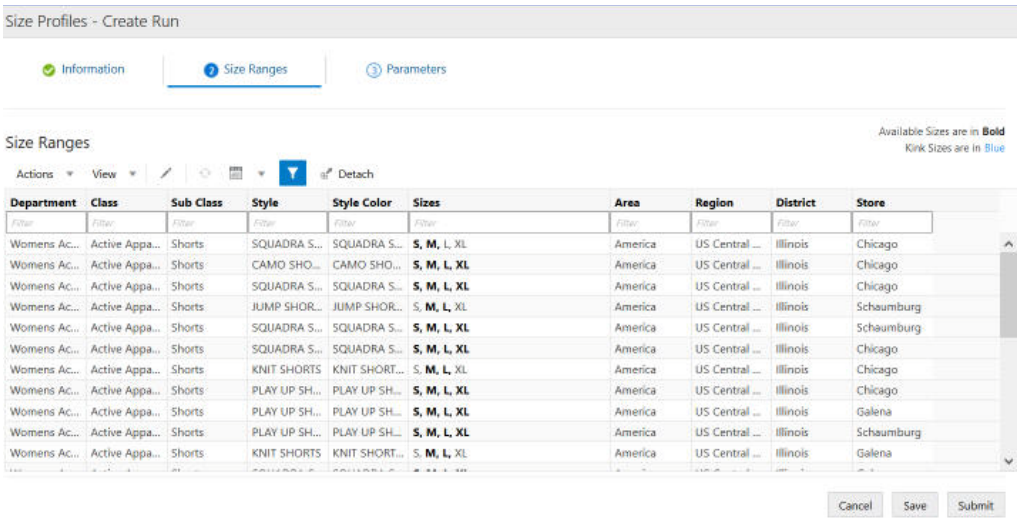
Use the Information tab, shown in Figure 11–3, to provide the run name and description in the Run Details section. Use the drop-down lists in the Setup section to select the products and locations that the size profile should be generated for. Select the season code for the time period for the analysis.

Size Ranges Tab

A size range consists of a set of all available sizes for each style-color. For example, the size range for the style-color of Men's Pullover Black and Store of Boston Downtown is S, M, L, XL. The size range data is provided by retailer.

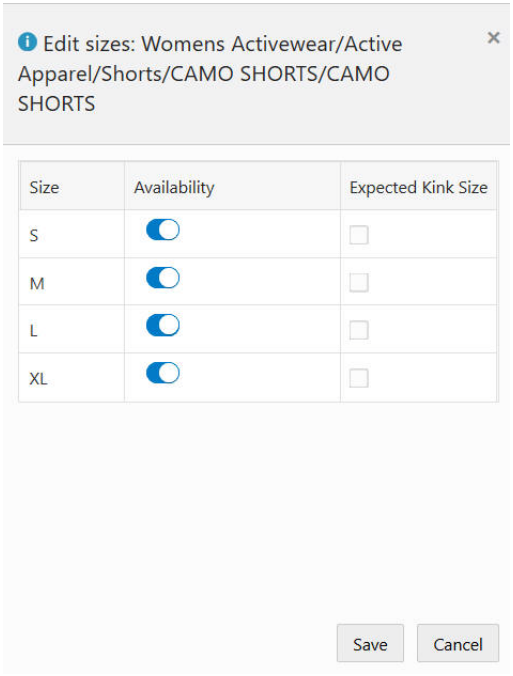
In the Size Ranges tab, shown in Figure 11–4, size range data is shown for each style-color and each store. If a certain size is not available in a size range, it will be shown as greyed out. If a size has been specified as an expected Kink Size (which means it is an odd size that may have significantly fewer or more sales than other sizes), it will be shown in blue.

Figure 11–4 Create Run Showing Size Ranges Tab



To edit the available sizes and kink size flags, select a row and click the **Edit** icon. This opens a pop-up, shown in Figure 11–5

Figure 11–5 Edit Available Sizes



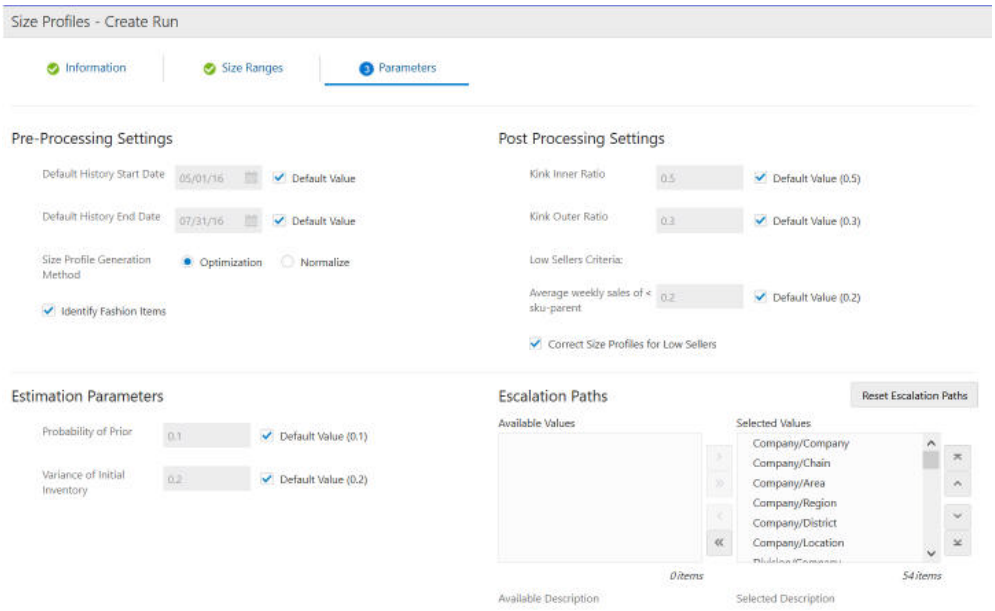
Parameters Tab

The Parameters tab displays the following parameters, which are used by the size profile generation algorithm. The escalation path is used to find the optimal size profiles associated with each style-color and store. When a size profile at a given level of the escalation path does not meet the post-processing criteria, it automatically gets rejected and the size profile associated with a higher level is examined. This process continues until a valid size profile is fond at one of the levels along the escalation path.

- Pre-Processing Settings
- Post-Processing Settings
- Estimation Parameters
- Escalation Paths

You can edit the values for these parameters, as shown in Figure 11–6.

Figure 11–6 Create Run Showing Parameters Tab



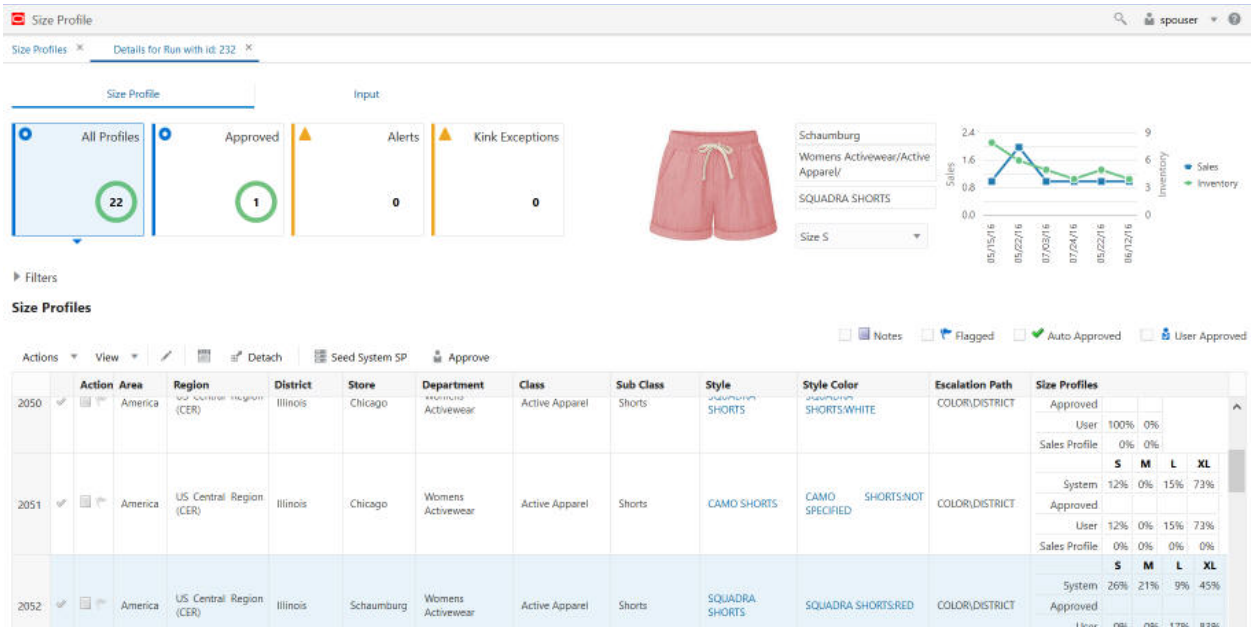
Reviewing a Run

Once the run has completed successfully, you can review the size profile output. You can override and approve or un-approve size profiles at style-color/store level in the Output screen.

Run Output

To open a run with a status of Complete, click the **Run Name** in either the All Runs table or from the list of recent runs, as shown in Figure 11–2. The Run Output screen opens in a new tab. This screen has two tabs, Size Profiles and Inputs, as shown in Figure 11–7.

Figure 11–7 Run Output



Size Profiles Tab

The Size Profiles tab displays the generated size profiles for each style-color/store for a specific run. You can filter the profiles based on their approved and alert status using the three Info tiles at the top of the screen for All Profiles, Approved, Kink Exceptions, and Alerts.

Below the tiles, you can use the collapsible filter section to filter the size profiles based on product and location hierarchies. You can also use the check-box filters above the table to filter profiles by status (auto/user approved) as well a notes/flag.

Size Profiles Table

The size profiles table is multi-row selectable, using the Ctrl and Shift buttons. All rows can be selected at once by clicking on the top-left cell in the table. This table has the following columns:

- The first column (with no header) shows the status of the profile. Values are user approved ( ), auto approved ( ), not approved ( ).
- The action column has the note and flag icons. Values are some notes have been added ( ), no notes added yet ( ), flagged ( ), not flagged ( ).

  You can click on the note icon to view or edit the notes in a pop-up. You can also click the flag icon to flag or un-flag a size profile.
- The columns that show the product and location hierarchy (for example, Zone, Store, Class, Style).
- The escalation path that indicates the level at which the size profile has been generated.

The size profiles chart shows the following values:

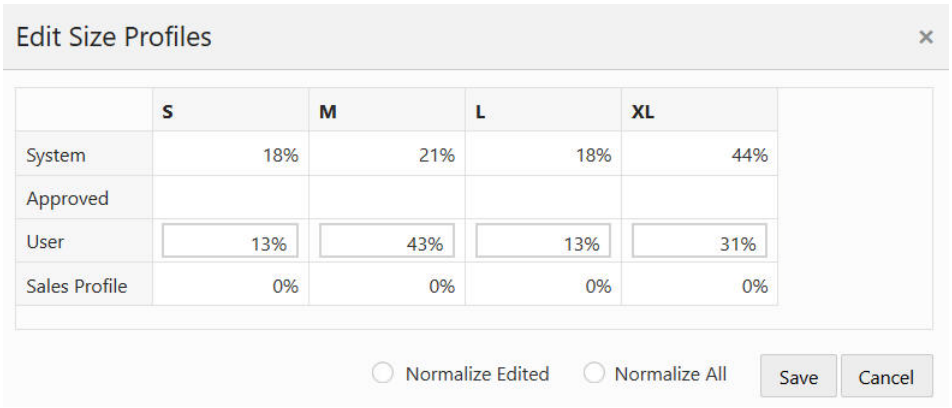
- System: the size profile generated by the system (using the estimation algorithm)
- User: the values edited by the user. If the user has not edited the values or has not seeded system values in order to populate user values, this row will be empty.

- Approved: the approved values, if the size profile was approved. If the size profile has not been approved, this row will be empty.
- Sales Profile: shows the profile of the observed sales (without estimating the lost sales and any other post-processing).

The following actions can be done in the Size Profiles table:

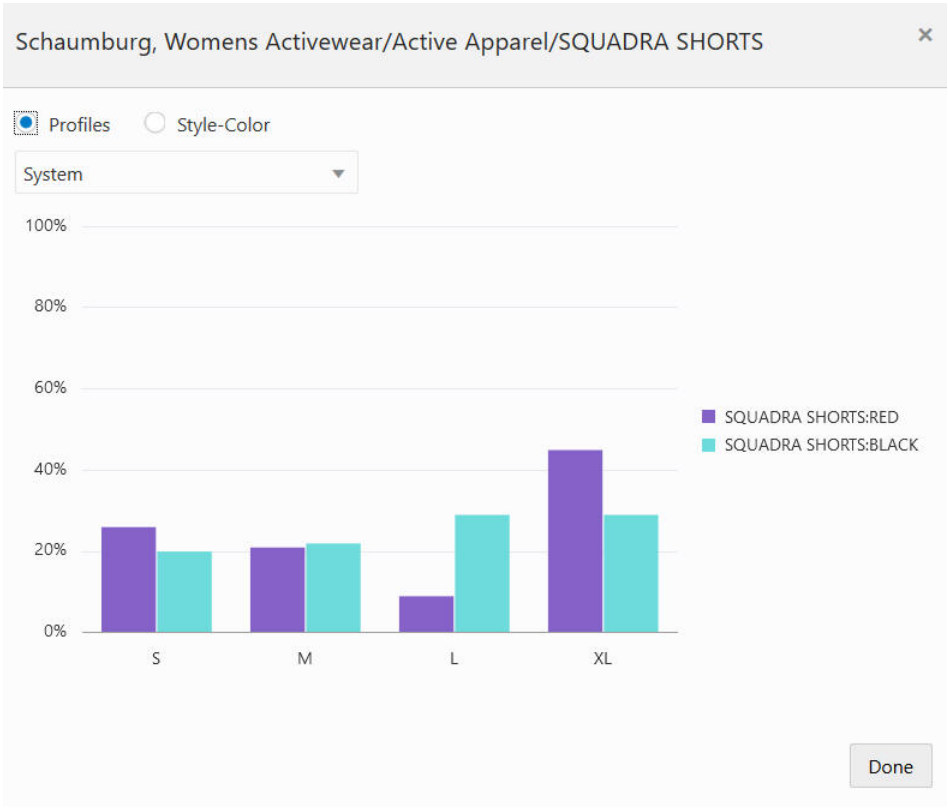
- Seed the user values with the System size profile by selecting one, multiple, or all rows and clicking on the Seed System SP button at the top of the table.
- Edit a profile in the edit pop-up by selecting a single row and clicking on the Edit icon at the top of the table. If a size profile is approved, edit is not allowed. To edit, the user must first un-approve the profile. The edit pop-up is shown in Figure 11–8. The edited values must sum up to 100 percent. If they do not, you will not be able to save. To normalize the values, you can use one of the two options (Normalize All or Normalize Edited). The first option adjusts all values so that they sum up to 100 percent. The second option adjusts only the edited values so that the sum of all the values becomes 100 percent.

Figure 11–8 Edit Size Profiles



- Approve or un-approve a profile by selecting one, multiple, or all rows and clicking on the Approve/Un-approve button at the top of the table.
- View and edit the notes on each profile using the Note icon in the Action column, which opens a pop-up.
- Flag or un-flag a profile by clicking on the Flag icon in the Action column.
- View the graphs for each profile by clicking on a hyper-linked style (which opens the graph pop-up and shows the Profile tab by default as shown in Figure 11–9), or by clicking on a hyper-linked color (which opens the pop-up and shows the Color tab by default as shown in Figure 10).

Figure 11–9 Profiles View



- View the weekly sales and inventory and product image at the top of the screen for a selected row in the table.

Note: If a profile (for a given style-color/location) has already been submitted or sent for review, the seed, approve/un-approve, and edit actions are not available. None of these actions can be performed for the profile. Note and flag can still be modified.

Note: If a user approves a profile for a given style-color/store that has previously been approved in another run, it will override the older run. As a result, the previously approved profile will become unapproved.

Input Tab

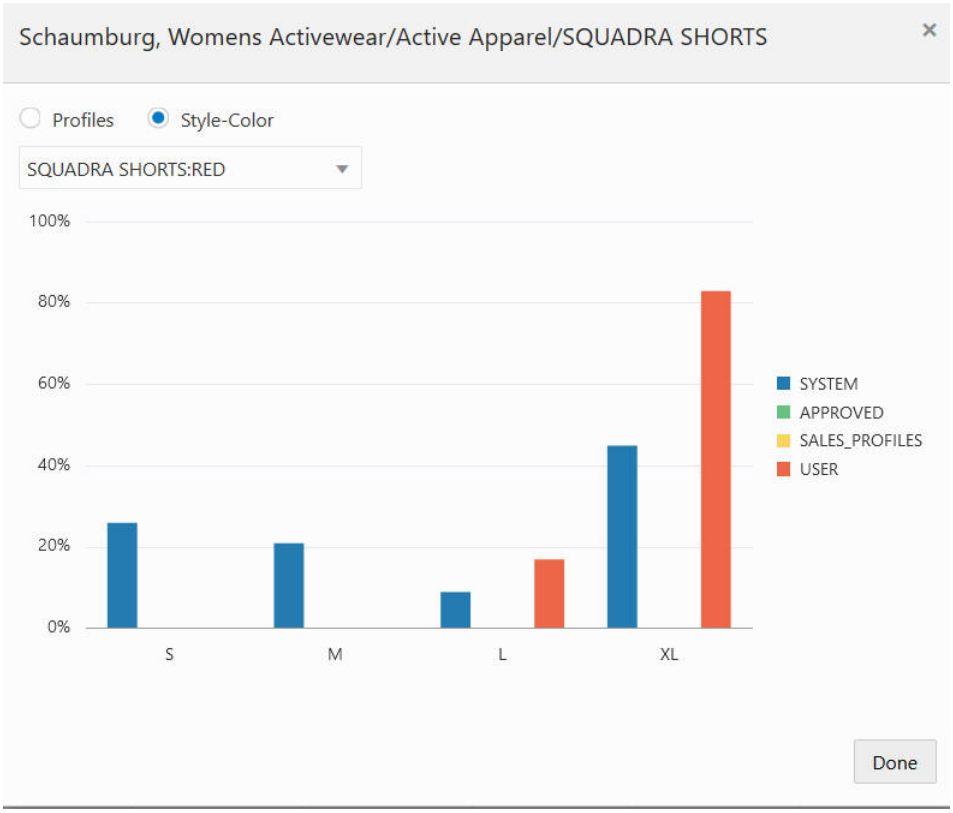
The input tab shows a summary of all inputs for the run in read-only mode.

Graphics

The colors view shown in Figure 11–10, shows the four profiles (system, user, approved and sales profiles) for each color that is available within the selected style. You can select a different color from the drop-down list.

The profiles view shown in Figure 11–9 shows one of the profiles (system, user, approved or sales profiles) for all colors under the style. You can select a different profile from the drop-down list.

Figure 11–10 Colors View



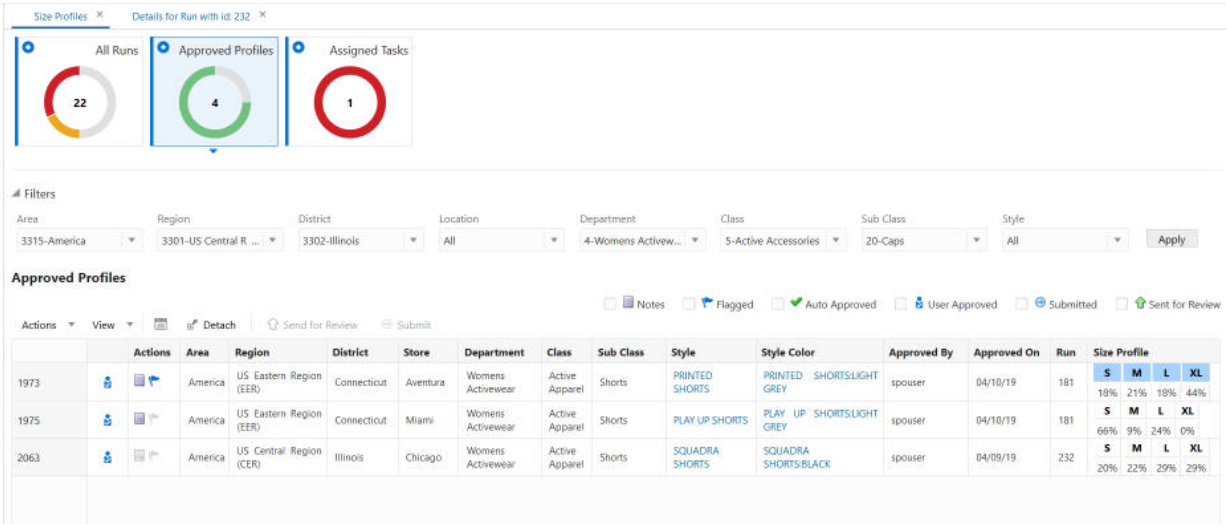
Submitting a Run

You can see all the approved profiles across all runs in the Approved Profiles screen. You can also select one or multiple profiles and send them for further review to another user. Alternatively, you can submit one or multiple profiles to a downstream application. Depending on your user role, you may only be able to send the approved profile for review, but not submit it for final export.

Approved Profiles

In the approved profiles screen, shown in Figure 11–11,you can view approved profiles across all runs.

Figure 11–11 Approved Profiles



The Filters section is the same as the one in the Run Output screen. You can use it to filter the size profiles based on product and location hierarchies. You can also use the check-box filters above the table to filter profiles by status (auto/user approved, submitted, and sent for review) as well a notes/flag.

Approved Profiles Table

The approved profiles table is multi-row selectable, using the Ctrl and Shift buttons. All rows can be selected at once by clicking on the top-left cell in the table. This table has the following columns:

- The first column (with no header) shows the status of the profile: Value are user approved ( ), auto approved ( ), submitted ( ), sent for review ( ).
- The action column has the note and flag icons. Values are some notes have been added ( ), no notes added yet ( ), flagged ( ), not flagged ( ).
- You can click on the Note icon to view or edit the notes in a pop-up. You can also click the Flag icon to flag or un-flag a size profile.
- Columns that show the product and location hierarchy (for example, Zone, Store, Class, Style).
- The escalation path that indicates the level at which the size profile has been generated.
- Approved by, Approved on, and Run (on which the profile was approved).
- Size profiles shows the approved profile.

The following actions can be performed in this screen:

- View and edit the notes on each profile, using the Note icon in the Action column that opens a pop-up.
- Flag or un-flag a profile by clicking on the Flag icon in the Action column.
- Send a profile for review to another user role, using the **Send for Review** button above the table. It will open a pop-up in which you can select the reviewer role.

  If a profile has been submitted, it cannot be sent for review again.The button will be disabled for submitted profiles.

- Submit the profile to be used by other applications using the **Submit** button above the table.

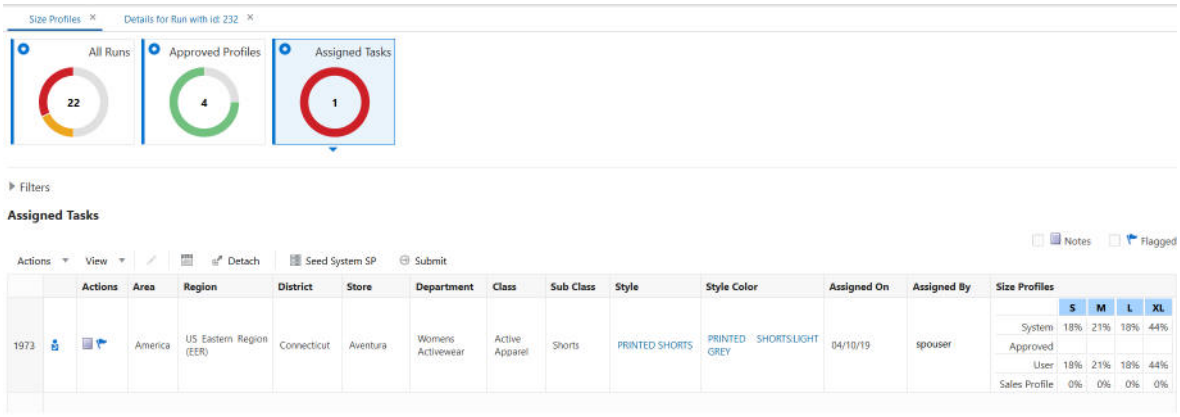
  If a profile has been sent for review, it cannot be submitted again. So the button will be disabled for 'sent for review' profiles.

- View the graphs for each profile by clicking on a hyper-linked style (which opens the graph pop-up and shows the Profile tab by default as shown in Figure 11–9), or by clicking on a hyper-linked color (which opens the pop-up and shows the Color tab by default as shown in Figure 11–10).

Assigned Tasks

In the Assigned Tasks screen, shown in Figure 11–12, you can see any profile that has been assigned to you for review by another user.

Figure 11–12 Assigned Tasks



The filters section is the same as the one in the Run Output screen. You can use it to filter the size profiles based on product and location hierarchies. You can also use the check-box filters above the table to filter profiles by status (auto/user approved, submitted, and sent for review) as well a notes/flag.

Assigned Tasks Table

The assigned tasks table is multi-row selectable using the Ctrl and Shift buttons. All rows can be selected at once by clicking on the top-left cell in the table. This table has the following columns:

- The first column (with no header) shows the status of the profile: Values are user approved ( ), auto approved ( ) as well as icons for notes ( ) and flag ( ). User can click on the note icon to view the notes in the opened pop-up. Note is not editable. The flag icon indicates whether the profile has been flagged by the sender user or not. The profile cannot be flagged or un-flagged in the assigned tasks screen.
- The columns that show the product and location hierarchy (for example, Zone, Store, Class, Style).
- The escalation path which indicates the level at which the size profile has been generated.
- Assigned on indicates the date the profile was approved on.
- Assigned by indicates the user who approved the profile.

Size profiles chart that shows the following values:

- System: the size profile generated by the system (using estimation algorithm)
- User: the values edited by the user. Initially this row shows the values edited by the sender user (if any edits has been made).
- Approved: the values that were approved by the sender user.
- Sales Profile: shows the profile of the observed sales (without estimating the lost sales and any other post-processing).

The following actions can be performed in this screen:

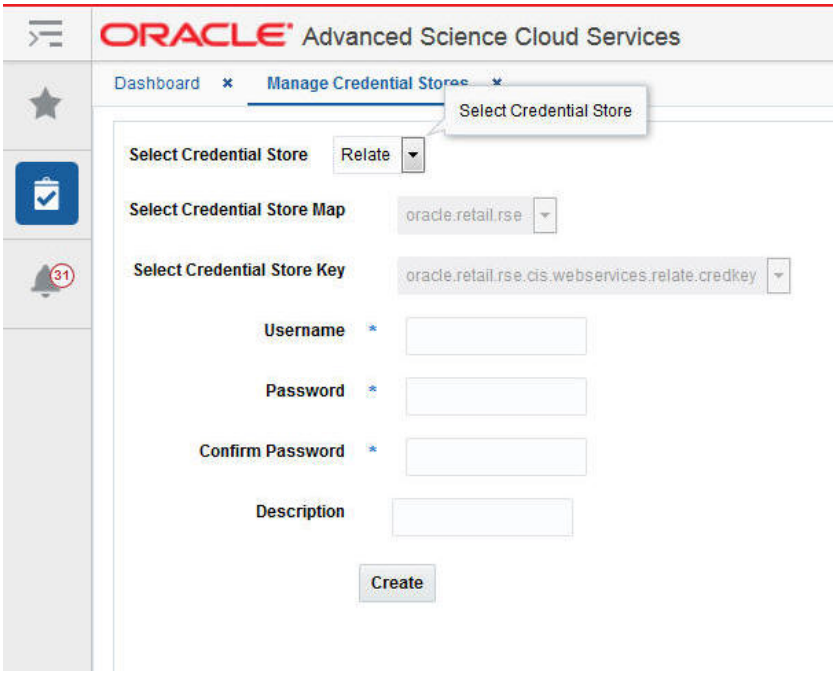
- Seed the user values with System size profile, by selecting one, or multiple, or all rows and clicking on the **Seed System SP** button at the top of the table.
- Edit a profile in the edit pop-up, by selecting a single row and clicking on the Edit icon at the top of the table.
- Submit the profile to be used by other applications, using the **Submit** button above the table.
- View the notes on each profile, using the in-line Note icon. Notes are not editable.
- View the graphs for each profile by clicking on a hyper-linked style (which opens the graph pop-up and shows the Profile tab by default as shown in Figure 11–9), or by clicking on a hyper-linked color (which opens the pop-up and shows the Color tab by default as shown in Figure 11–10).

12

Data Management

Oracle Customer Engagement (Relate) web service related credentials are managed in credential stores using the interface shown in Figure 12–1. The administrator can use this screen to configure the name and password, which is used for generating the authentication key that is sent as part of the message to Oracle Customer Engagement.

Figure 12–1 Manage Credential Stores



The Credential Stores dialog box requires the following information:

- Label - describes the credential store map and key
- Username - user name provided by Oracle Customer Engagement (Relate)
- Password - password provided by Oracle Customer Engagement (Relate)
- Confirm Password - prompt to confirm user password
- Description - used by the administrator to describe the credential store used for Customer Segment integration with Oracle Customer Engagement.

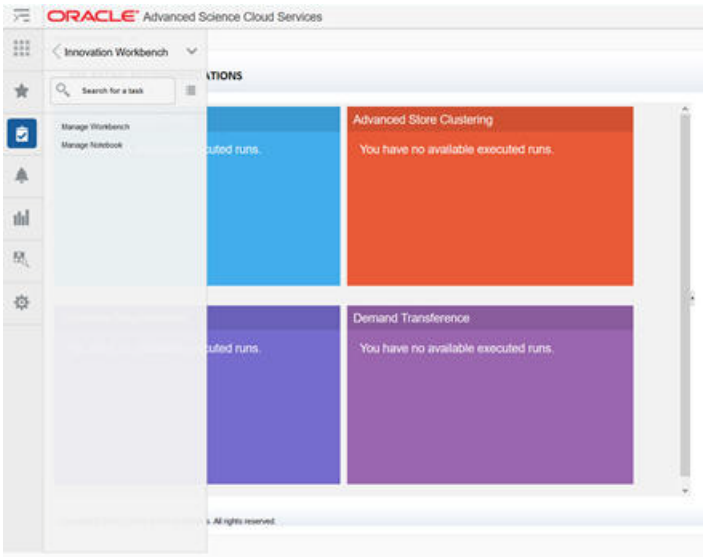
13

Innovation Workbench

Innovation Workbench (IW) is a service that enables you to access read-only ORASE data using Oracle APEX and implement algorithms and analytics as an extension. This extension is a workbench for advanced analytics users that can be used to add new implementations via Oracle Advanced Analytic (Oracle R/ODM) algorithms that are implemented as SQL/PLSQL functions.

Innovation Workbench allocates a workspace to a retailer. This is a logical work area that is associated with a pre-allocated retailer database schema.

Figure 13–1 Accessing Innovation Workbench



Process

To access the workbench, in the Task menu shown in Figure 13–1, click **Innovation Workbench.** You see a list of workspaces in a new browser tab that are pre-allocated to the retailers you have access to.

Workspace List

The Workspace List displays all the workspaces that you were initially assigned to when workspaces were created at the time of installation and configuration. You can select a workspace from the list to review database objects using Oracle Application Express, shown in Figure 13–2.

Figure 13–2 Oracle Application Express

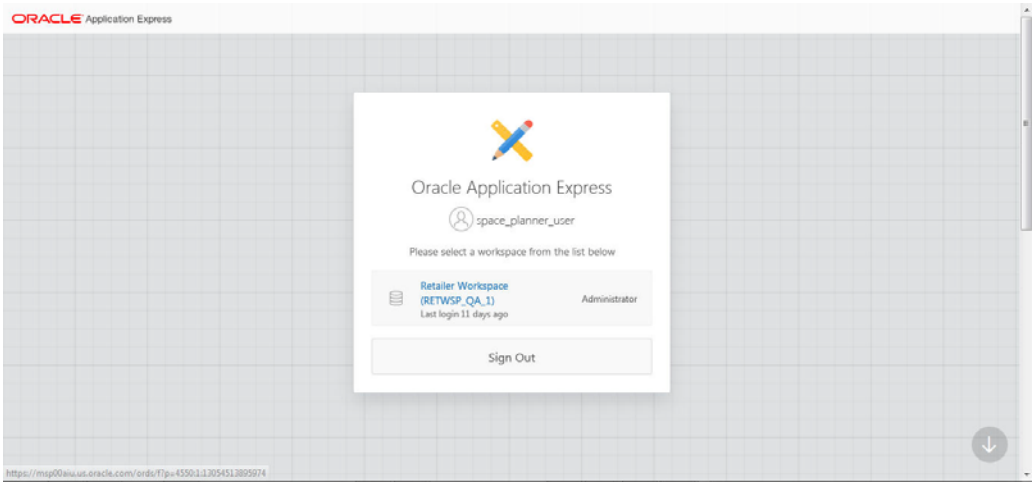
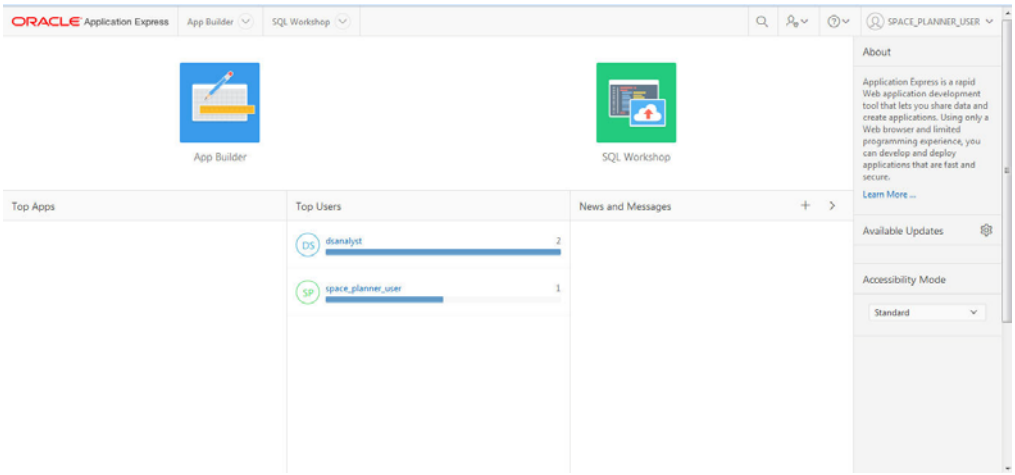


Figure 13–3 shows the Oracle APEX workspace, which displays SQL Workshop and App Builder.

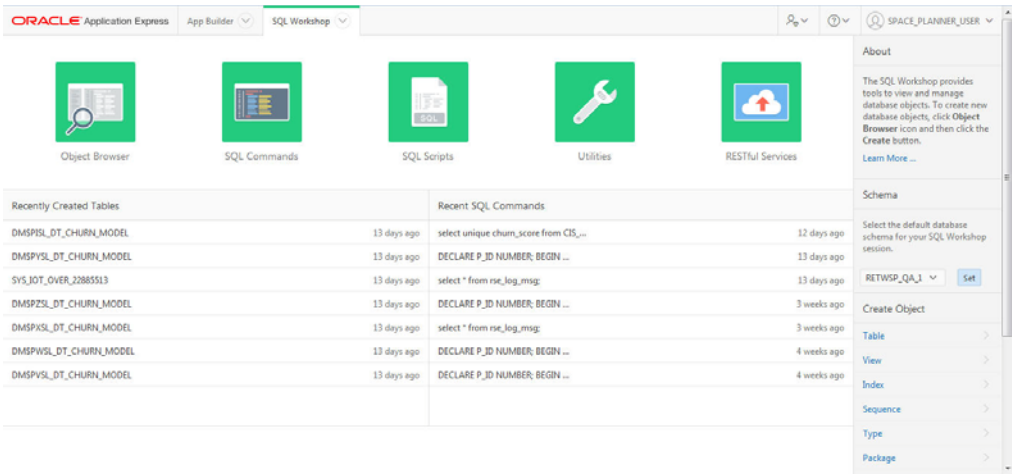
Figure 13–3 Oracle APEX Workspace



SQL Workshop

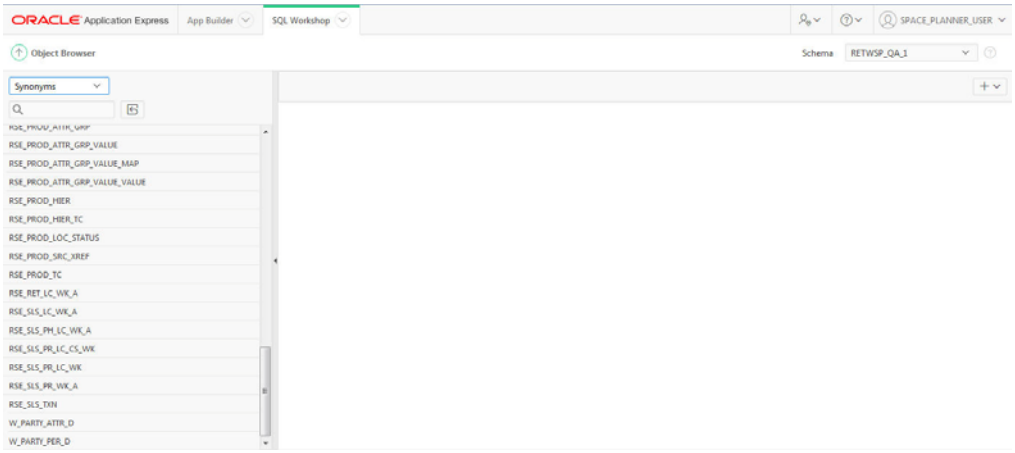
The SQL Workshop shown in Figure 13–4 provides tools to view and manage database objects.

Figure 13–4 SQL Workshop



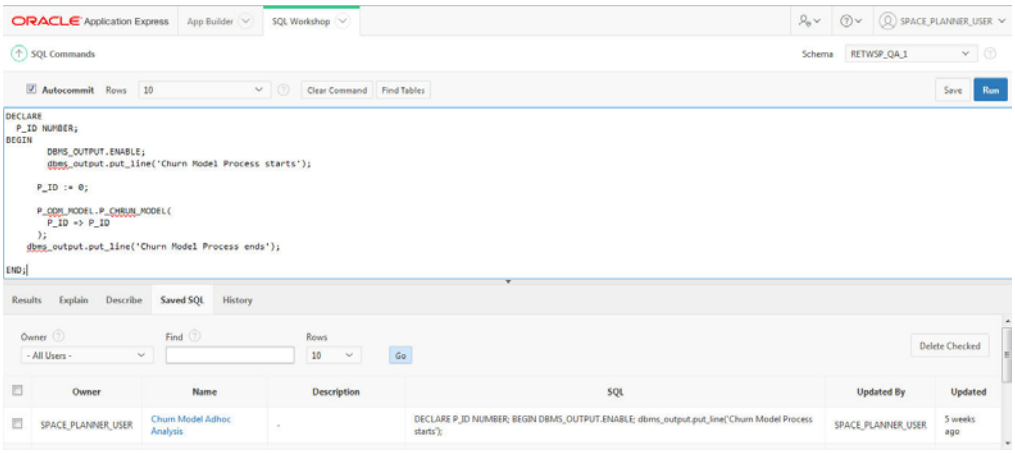
The Object Browser shown in Figure 13–5 lets you to use a tree control to view object properties and create new objects. The read-only ORASE database objects that are aggregated and cleansed are available for you to review and organize for data mining. Figure 13–4 shows the schema associated with the Retailer Workspace schema.

Figure 13–5 Contextual Area



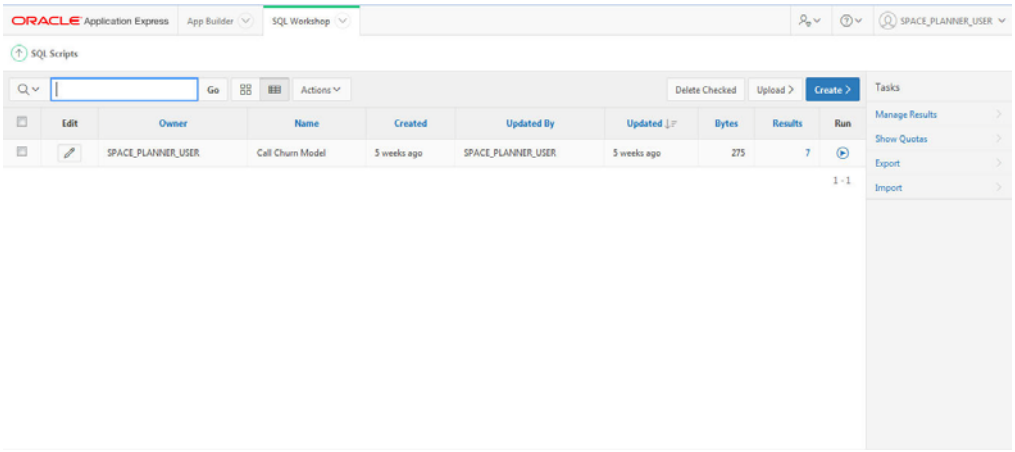
You can use the SQL Commands tool to enter ad-hoc SQL. The Query Builder can be used to create join queries using drag and drop. Figure 13–6 shows how to execute the ad hoc PL/SQL block.

Figure 13–6 SQL Commands



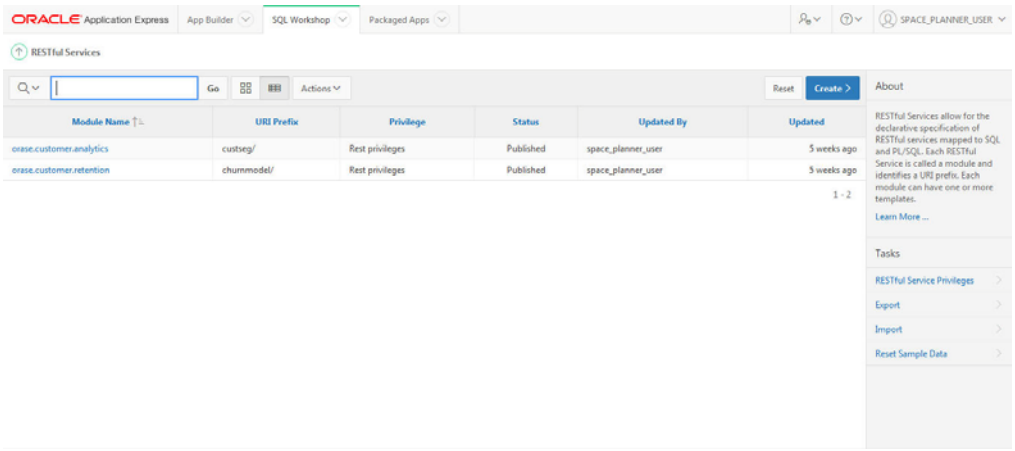
The SQL Scripts shown in Figure 13–7 can be used to store and run scripts. Use the Data Workshop to load and unload text, DDL, and spreadsheet data.

Figure 13–7 SQL Scripts



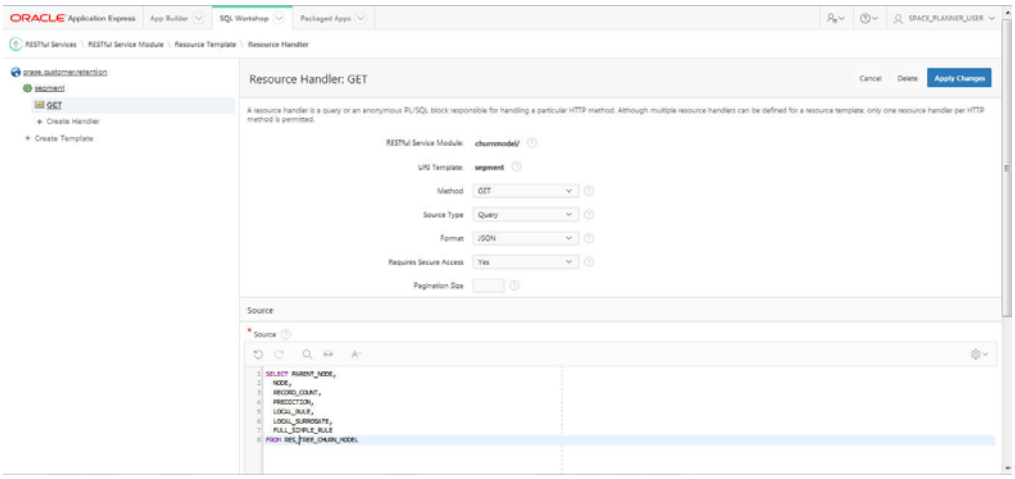
Use RESTful Services to manage web service definitions declaratively using SQL Query and to test the service modules from the web browser. Figure 13–8 shows a list of RESTful service modules.

Figure 13–8 RESTful Services



You can create a new RESTful Service module by clicking **Create** and setting the resource handler with a query or an anonymous PL/SQL block responsible for handling a particular HTTP method.

Figure 13–9 Resource Handler



APEX SQL Workshop Guide is available at
http://docs.oracle.com/database/apex-5.1/AEUTL/toc.htm

Application Builder

Application Builder is a development environment that you can use to build and deploy database-centric applications. You can create interactive grids, reports, and charts.

APEX App Builder User's Guide is available at
http://docs.oracle.com/database/apex-5.1/HTMDB/toc.htm

You can use wizards to declaratively assemble database applications organized by page to visualize and explore data, as a pre-process or post-process of data mining.

The following functionality is available.

Create

Launches the Create Application Wizard, which you can use to build a complete application containing multiple pages.

Import

Launches the Import Wizard, which you can use to import export files from the export repository.

Dashboard

Links to the App Builder Dashboard, which displays metrics about applications in the current workspace.

Workspace Utilities

Use the Workspace Utilities page to manage App Builder Defaults, news items displayed on the home page, workspace themes, and so on.

Manage Notebook

Data scientists can use the Innovation Workbench Notebook to create notebooks, which are collections of documentation, snippets of code, and visualizations. These notebooks are bundled with key python modules for machine learning, data mining, natural language processing, network analysis, and optimization solvers. Interactive graph visualization is available that supports highlighting, expansion, and applying filters. Custom paragraphs for advanced graph visualization concepts such as visual graph creation and large graph visualization concepts are supported. This powerful widget concept allows the analyst to curate data into a form that is easier to understand, highlights trends and outliers, and shares reports with business.

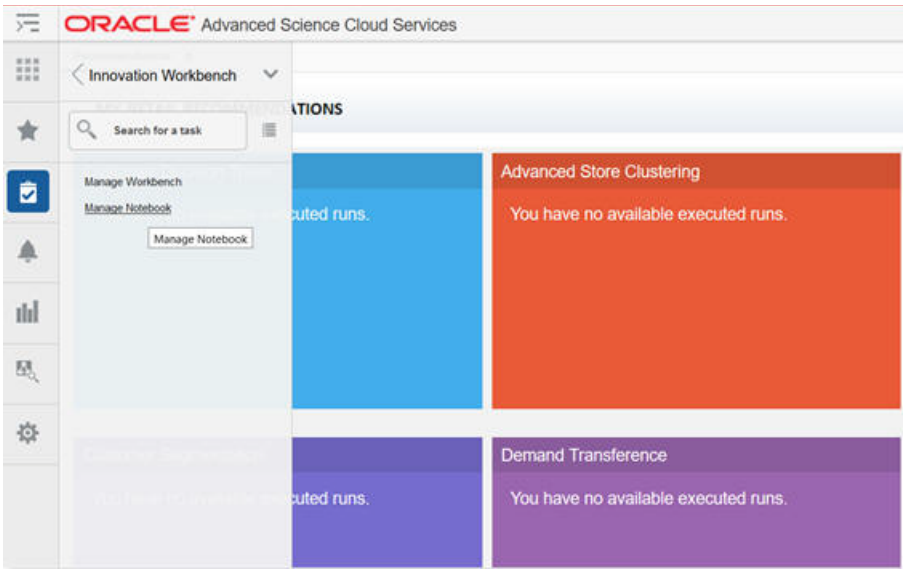
Features

The key features of Manage Notebook include:

- Create, import, export, and clone python notebooks
- Browse, filter, and search for Notebooks.
- Add comments from the user interface (UI)
- Visualize data in the form of a Graph Visualization.
- Use collection of Python Libraries bundled with the product
  - Machine Learning packages such as (Tensorflow, Mxnet and Keras) Analyze
  - Optimization Solvers such as Gurobi

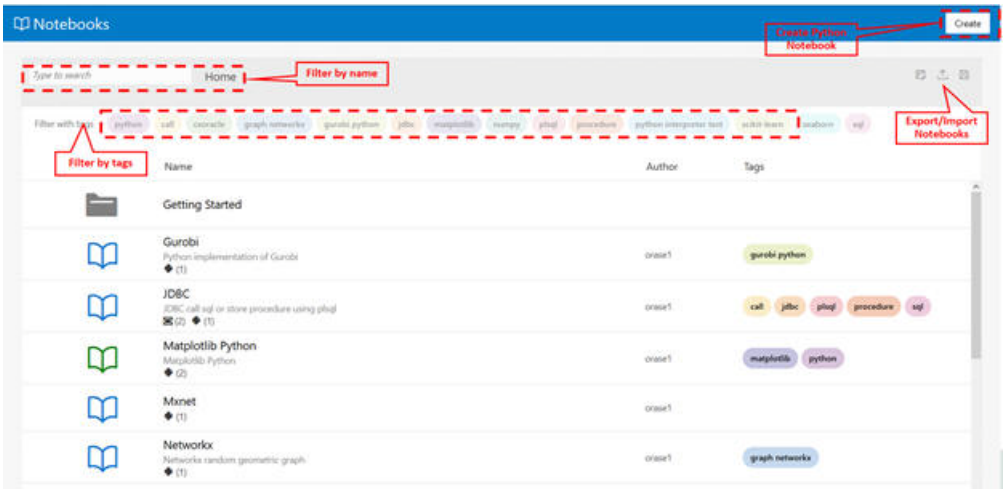
To access Python Notebook, click Innovation Workbench **Manage Notebook**. This opens a new tab in the browser and all the notebooks in the workspace are displayed.

Figure 13–10 Manage Notebook



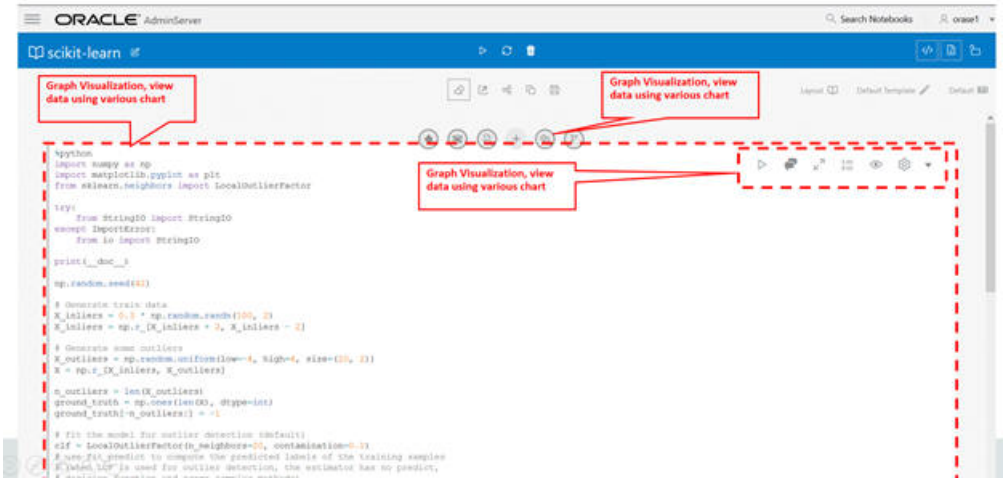
From the workspace, you can import or export notebooks, create new workbooks, or copy notebooks. You can also browse, filter, and search for Notebooks using either a name or tags.

Figure 13–11 Notebook Functionality



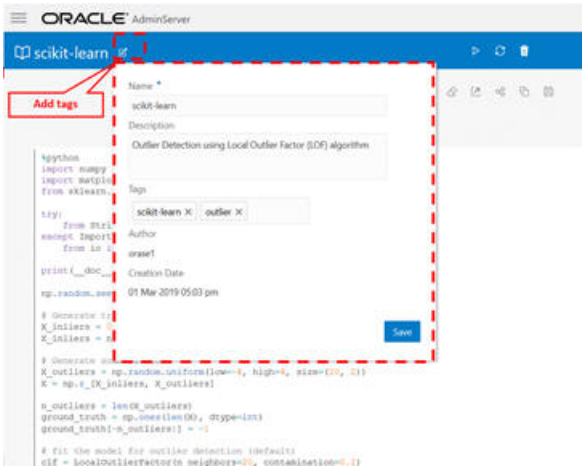
You can create a notebook and add paragraphs using the toolbar. These paragraphs can be Python, JDBC, Markdown, or graph analytics paragraphs. You can execute paragraphs individually and see the results at the bottom of each paragraph. You can also execute all the paragraphs from the toolbar provided at the top.

Figure 13–12 Notebook - Adding Paragraphs



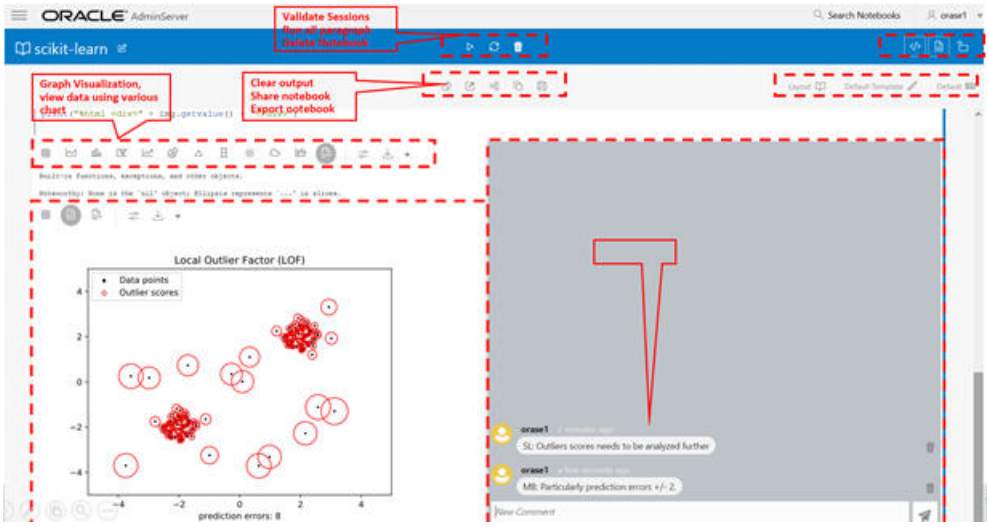
You can associate tags with the Notebook; these tags can be used for filtering and browsing the notebooks.

Figure 13–13 Notebook Tags



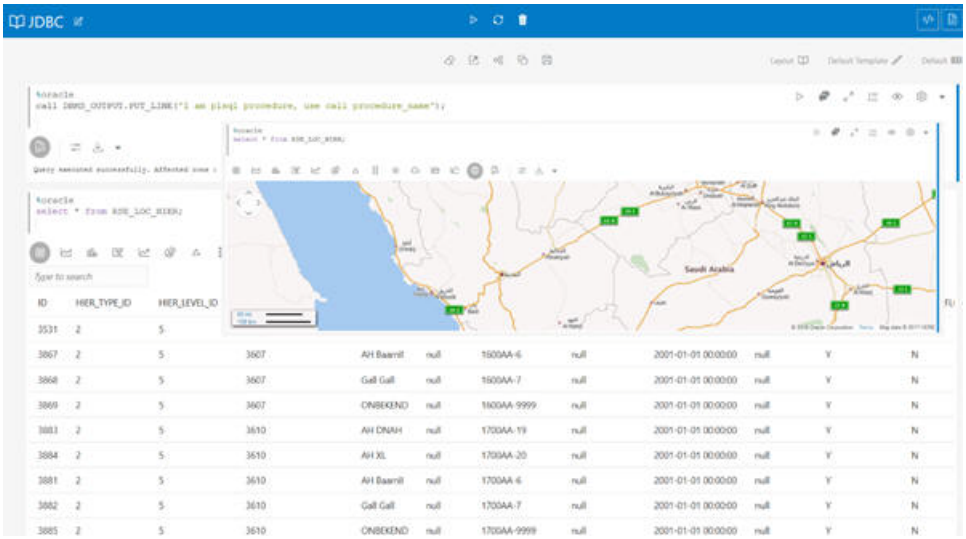
You can also add comments to the paragraphs and manage them from the user interface (UI). The comments can then be shared with others as a PDF. The Comments view shows all paragraph comments. You can visualize data in the form of a Graph Visualization and Exploration. You can explore graphs visually from within the notebook by viewing a summary of graphs or by using interactive exploration such as dragging, hovering, filtering, highlighting, and setting visualization properties. Different types of charts and graphs can be used for different types of datasets such as Line, Tag-Cloud, Tree Map, Scatter Plot, and so on.

Figure 13–14 Notebook Comments



You can visualize data in the form of a map and configure a column from a dataset to represent the latitude, longitude, title, and description.

Figure 13–15 Notebook Data as a Map



14

Oracle Digital Assistance

Oracle Digital Assistance (ODA) is a NLP and Machine Learning-based cloud service that enables application-specific interactions. It is a virtual user that assists end users with complex engagements using transactional data to review key business insights and exceptions. It also provides out-of-the-box knowledge sharing by responding to typical Q&A-style questions for customer service and support. ODA features are supported using voice along with conversational interface to distill end user intents, invoke actions, and provide reasonable responses. This enhances consulting services with minimal training and support.

Digital assistance is not a replacement for a web application but a channel that allows the user to complete context-driven tasks using a combination of text messages, voice and, simple UI.

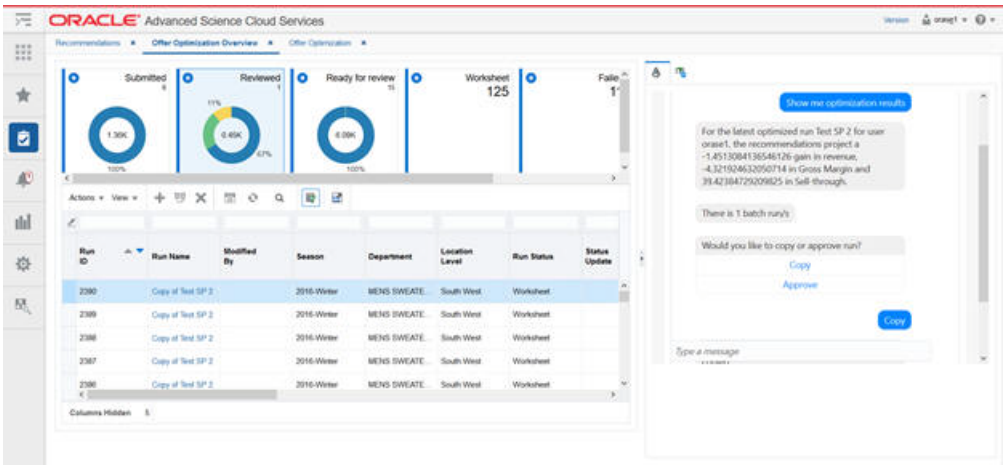
Digital assistance supports Transactional bots and Q&A bots.

This chapter describes how to use digital assistance. Digital Assistance in the Retail Science Platform is displayed in the Contextual area of each module.

Transactional Digital Assistance

Transactional digital assistance helps with business engagements by using transactional data to review business insights, rules, and exceptions. This assistant complete tasks and helps the end user navigate to specific tasks in application. The data source for these bots is the Retail Science Platform Database. They rely on NLP or Machine Learning to determine user intent.

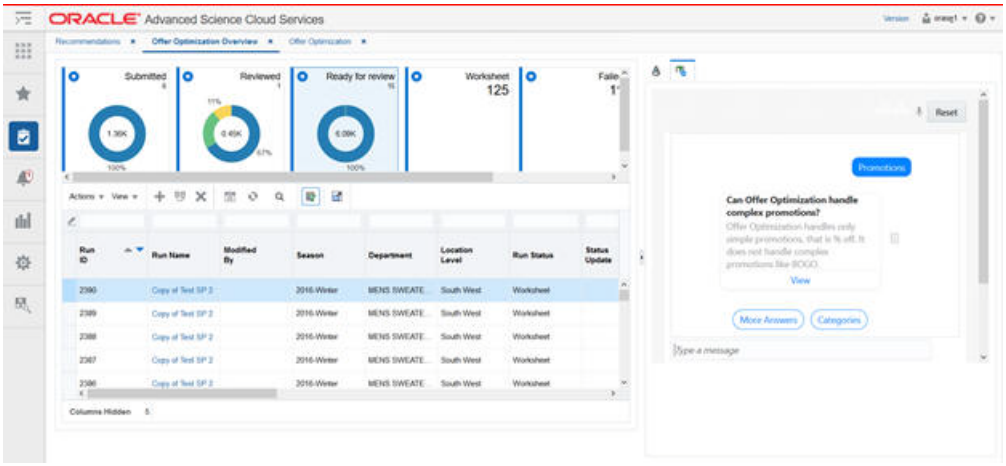
Figure 14–1 Transactional Digital Assistance



Q&A Digital Assistance

Q&A digital assistance helps answer general interest questions by returning one or more question and answer pairs. It helps the user find FAQs or other knowledge-based documents. The data source of these bots is pre-loaded CSV. They rely on Elastic Search to determine user intent.

Figure 14–2 Q&A Digital Assistance



Digital Assistance appears in the Contextual Area and has the components described in Table 14–1.

Figure 14–3 Digital Assistance Contextual Area

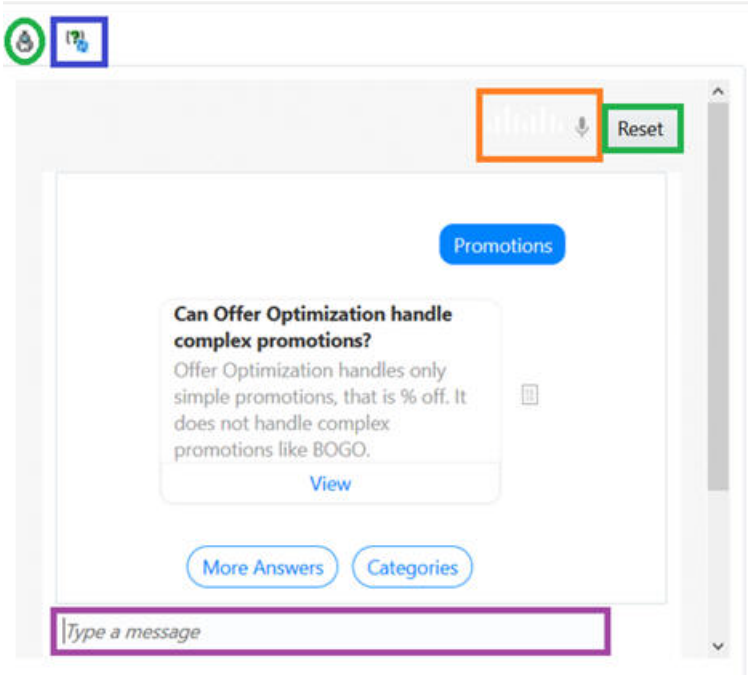
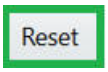
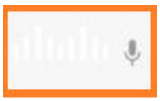
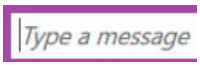


Table 14–1 Interface Components

Icon	Description
	Access transaction bots by clicking the bot icon in the contextual area.
	Access Q&A bots by clicking the question icon in the contextual area.
	Reset the context of the bots by clicking the reset icon.
	Access voice component by clicking the microphone icon. If voice is enabled, the bar color next to the microphone icon will be become orange and will animate. When user speaks, the content is typed in the message window.
	User can type message and click Enter.

Both of these bots can have voice enabled by clicking on microphone icon. To enable voice, complete the steps described in "Voice Component."

Voice Component

Voice support for ODA is only available in Chrome.

Enabling the microphone in Chrome

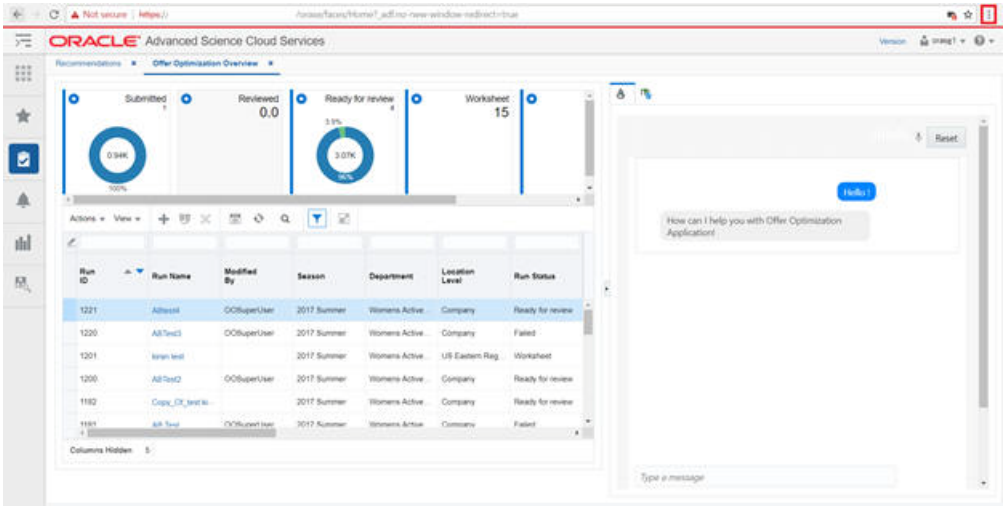
Chrome requires permission to enable the Microphone for all URLs (for example, https://<hostname>:<port>).

The process for adding or enabling Microphone differs in different versions of Chrome. Here are the recommended steps. The https mode is recommended.

https://<hostname>:<port>/orase/faces/Home

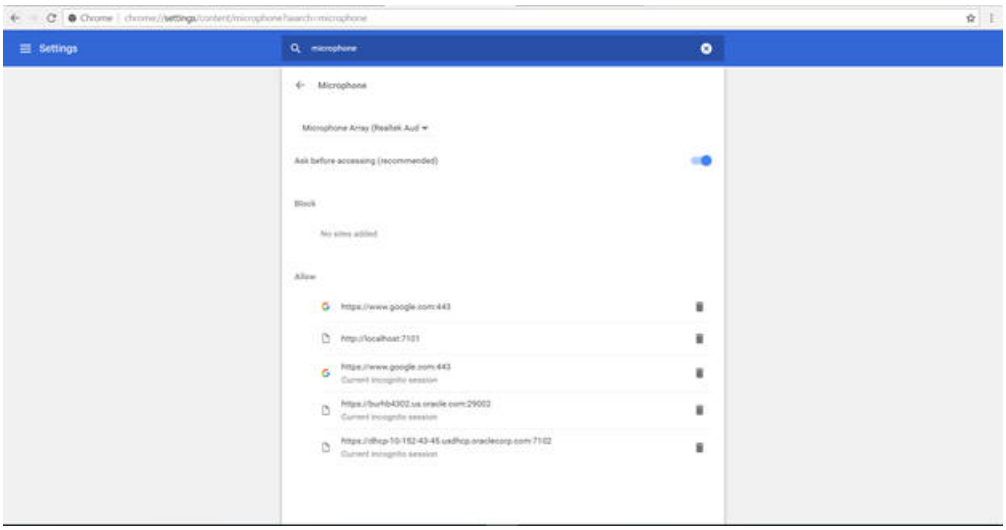
1. Click dot in the top right corner to check the settings. This opens a new tab in the browser.

Figure 14–4 Click Dots



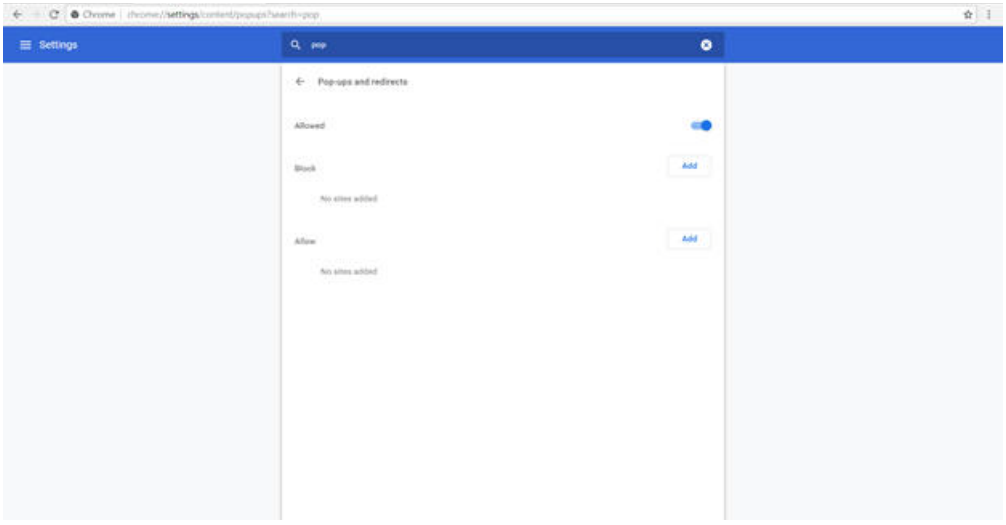
2. Type **microphone** in the search text box.

Figure 14–5 Search for Microphone



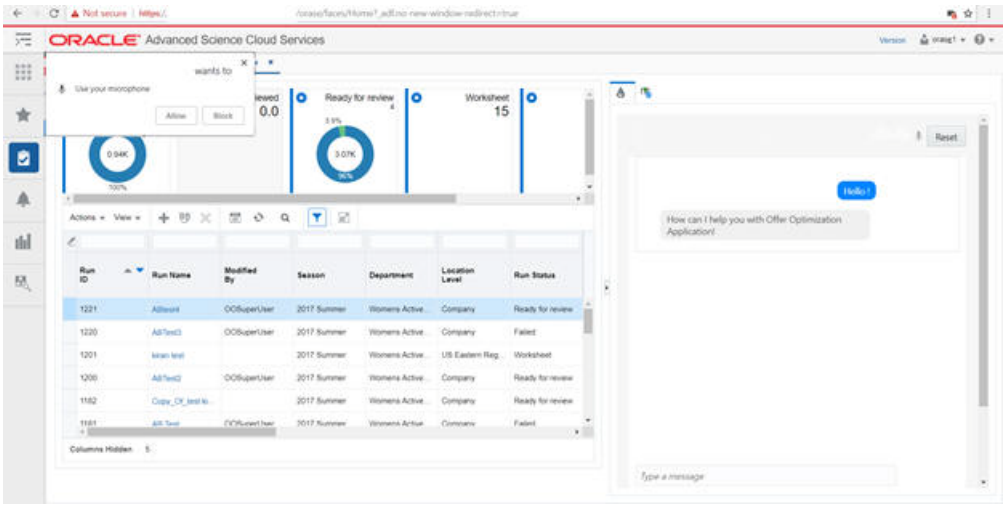
3. If the application URL is not added to allow section, then follow steps 4 through 7. No setup is required for voice enablement. Make sure that pop-ups are not blocked for the URL.

Figure 14–6 Allow Pop-Ups



4. Click **Microphone**. If the URL is allowed, you will see voice animation next to the yellow button when you speak. You are prompted to click Allow to enable voice. Click **Allow** or complete the following steps.

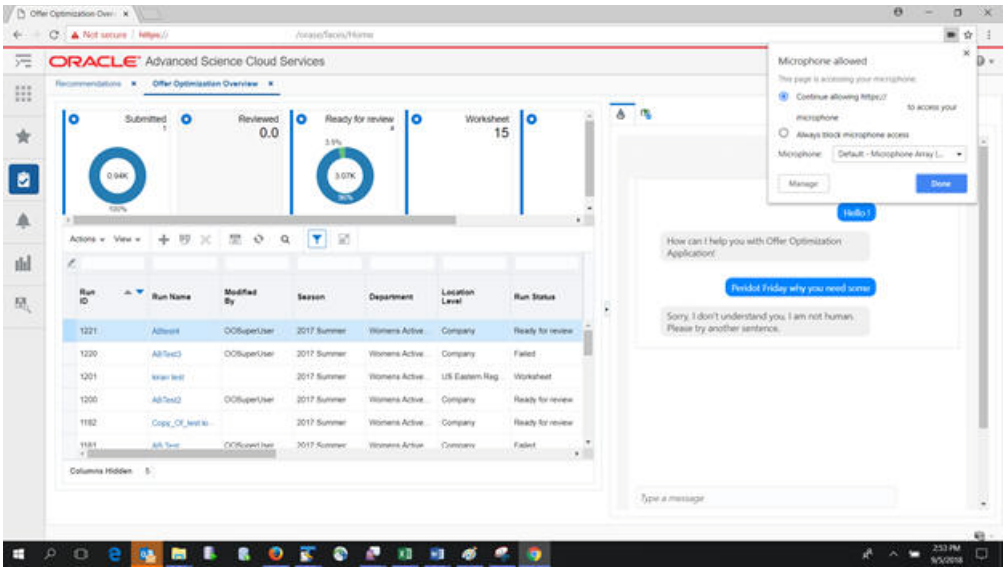
Figure 14–7 Click Microphone



5. The browser enables the button in the top right corner for permission. When asked if the URL should access the microphone, click **Allow**.

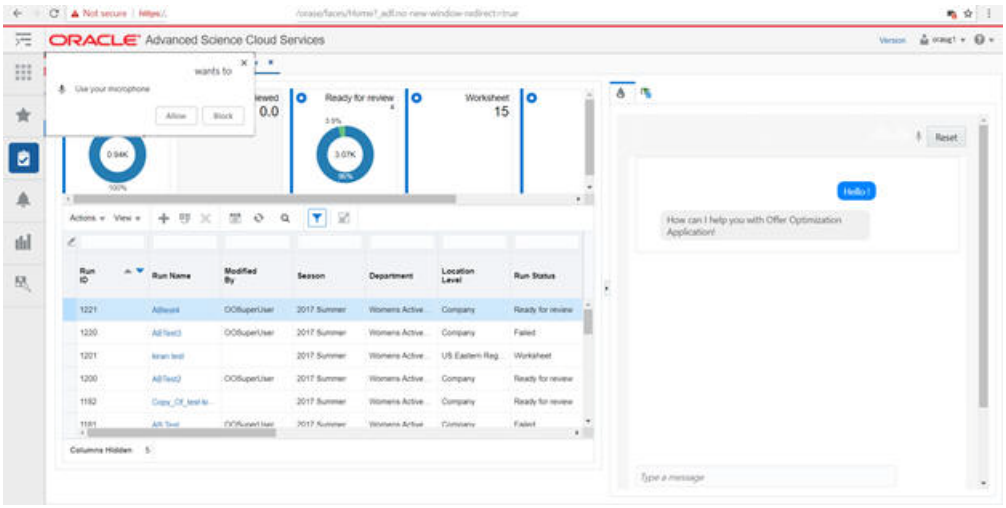
6. If above steps do not work, try using incognito mode by pressing Ctrl + Shift + N. When asked if <URL> want to access your microphone, reload the page.

Figure 14–8 Incognito Mode



7. Make sure the pop-up is not blocked and click Allow. Voice will be enabled.

Figure 14–9 Voice Enabled



Glossary

Assortment Cluster

Store cluster definitions that are used in Category Manager, which are defined at the product category level.

Bay

A physical unit of fixtures in which the configuration is the same from left to right. A fixture configuration can change across bays but not within a bay. Retailers often lay out POGs on a discrete number of bays. So bays can be considered a unit of macro space like linear ft/cm.

Capping

Products are stacked with units of different orientation. The bottom unit is in the main orientation, and the units above use a capping orientation. This applies only to shelves. Only simple capping is supported.

Current/Historical POG Assortment

The list of products from a specific POG that is imported from a third party POG tool. This assortment is POG-specific; the stores that received this POG have this assortment. A master list for a POG set can be created by finding the union of all products that are on the POGs in the set.

Display Style, Display Style Type

This provides the discrete unit of display for a given product. Display style types include single/unit, tray, case, and pallet. Display style types can also describe peg vs. shelf. These are set up in MSM.

For a given product, one or more display styles is established in MSM. Each one has its own set of size and spacing data, as well as a number of total units. For example, a specific SKU may have two display styles: single/unit and tray. Unit then corresponds to a single unit of the product. Tray is then 12 units. The number of units per tray depends on the product.

Note that MSM uses the term unit for a single unit display. In order to describe the number of units, ASO uses the term single/unit instead.

Empty Planogram

An empty POG is a POG skeleton. The bays and fixtures are defined but no products have been placed on the fixtures.

Facing

A facing is the smallest unit of micro space that is allocated to a product. On a shelf, facings are counted horizontally. In other words, it is the number of units across. On a pegboard or a freezer, facings are the number of total units that the customer can see when looking at the fixture.

Finalized Assortment

This is the final approved assortment.

Finger Space Above

The product-specific space that is required between the top of a product's facings and the bottom of the shelf above. This accounts for the space required to remove the product from the fixture. When available, the specified space is added to the top of the top unit in the product's facings in the *z* direction. This applies to pegboards, freezer chests, and shelves. In MSM, this is also called finger space above.

Finger Space Behind

The product-specific space that is required on the front and back of units of a product. When available, the specified space is added on either side of the unit in the *y* direction; half of the quantity goes on the front and half of the quantity goes on the back. This applies to shelves only. In MSM, this is called a gap in the *y* direction.

Finger Space Beside

The product-specific space that is required on the left side and right side of units of a product. When available, the specified space is added on either side of the product's facings; half of the quantity goes on either side in the *x* direction. For example, if there are two facings of product A with finger space beside equal to 1 inch, then horizontally it looks as follows: one-half inch, product A, one-half inch, one-half inch, product A, one-half inch. This applies to pegboards, freezer chests, and shelves. In MSM, this quantity is called a gap in the *x* direction.

Fixture

A fixture is defined as equipment such as a set of shelves, a pegboard, or a freezer chest. A single shelf is not a fixture, nor is a single peg a fixture.

Fixture Configuration

This provides details about the dimensions of shelving (that is, the number of shelves and the height, width, and depth of each one), freezer chests (that is, the height, width, and depth) or pegboard (in this case, the height, width, and depth equal the peg length, which may also have peg hole spacing details).

Fixture Type

This is the attribute that defines what kind of fixture is in a POG: shelf, pegboard, or freezer chest.

Inter Product Gap

The product-specific space that is required to separate adjacent products on a shelf. This accounts for space that holds product dividers. When available, the specified space is added to the left and right of the entire block of a product; half of the quantity goes on either side of the run of that product in the *x* direction. For example, if there are two facings of product A with inter product gap equal to 1 inch, then horizontally it looks as follows: one-half inch, product A, product A, one-half inch. This applies to pegboards, freezer chests, and shelves. In MSM, this is called finger space beside.

Item Performance Index (IPI)

This is a value generated by Category Manager that provides indication of the relative importance of products within an assortment. A value of 1 means the product is average, a value of 1.3 means it is 130% of average, and a value of 0.5 means it is 50% of average. ASO uses these values optionally as attributes or weights in the optimization objective function.

Mandatory Items

These are items from the preliminary assortment that cannot be dropped.

Optimization Level

This is the location level for micro optimization runs and is set for a specific optimization run. The options are cluster or store.

Optimization Location

This is a specific location node defined for optimization. It can be a store or a store cluster.

Order Point (OP) and Order Up to Level (OUTL)

These are proxies for the replenishment triggers and are used in the inventory and sales model. These values are not directly input from client feeds, but are calculated from input data.

Orientation, Legal Orientation, Default Orientation

This describes how a product's package faces the front of the fixture. The orientation has two pieces of information, the part of the package that is to the front (front, back, top, bottom, left, right) and how it is rotated in degrees (0, 90, 180, 270), for a total of 24 possible orientations. How a product is oriented determines how much space it takes in the *x*,*y*,*z* directions on a fixture.

Legal orientations are all orientations that can be used for the product.

The default orientation is the legal orientation that is normally used for the product.

Planogram (POG)

The layout of product on shelves or pegboards, which has an underlying fixture assumption (for example, number of shelves or overall size).

A third-party POG typically includes a visual depiction of the product layout with images. POGs are the underlying data and not the pictures.

Planogram (POG) set

Current and historical POGs in the same POG category or subcategory and with the same seasonal attribute.

Preliminary Category Manager Assortment

A product list created in Category Manager by product category, this varies by assortment cluster. The preliminary assortment is a list of products that are eligible. A preliminary assortment has a master list that is the union of all of the cluster-specific preliminary assortments.

Preliminary assortments from Category Manager are forward looking and can include new products or new product and store assignments.

POG Sub-Category, POG Category, POG Department

These are levels on the POG hierarchy, and are used to organize POGs. For example, the leaf to root path in the POG hierarchy is Grocery -> Crisps and Snacks -> Crackers.

Product Stacking Height Limit

This defines a product-specific limit on how high a given product can be stacked or nested in the *y* direction on any shelf. This applies only to shelves.

Season/Seasonal Attribute

This refers to a specific year-independent time period for a Category Manager assortment and a POG set. Examples include spring, holiday, back to school, and year-round.

Smart Start Process

The Smart Start Process (SSP) is a back-end process that assigns shelves to a Partial Shelf Fixture Planogram (PSFP), based on a combination of default values, user selections, and algorithms. SSP supports one or more algorithms that add shelves to PSFPs only. SSP does not perform fixture optimization.

Space Optimized Assortment

This is the product list after products have been dropped because of limited space. This varies by optimization location.

Stacking/Nesting, Nesting Height

Multiple units of a product can be stacked on top of one another in the *y* direction on a specific facing on a shelf. When the units are stacked, the total height equals the number of units x the height of a single unit. Nesting is defined as stacked units that fit inside each other so that the total height is less than the number of units x the height of a single unit. The nesting height is the incremental space that the nested units require. The total height of a facing with nested units is determined as follows: total height equals the height of a single unit + the number of nested units x the nesting height. This applies only to shelves.

Top Shelf Stacking Height Limit

This defines the limit on how high products can be stacked or nested on the top shelf in the *y* direction. This is not required for intermediate shelves in a fixture because of the physical limit created by the shelf above. This applies only to shelves.

Units of Measure and Currency

This document will use "ft/cm" when talking about units of space, and will use "value" when talking about currency

Virtual Planogram (VPOG)

This refers to a POG that was created by ASO.